

PoseExaminer: Automated Testing of Out-of-Distribution Robustness in Human Pose and Shape Estimation


Qihao Liu¹Adam Kortylewski^{2,3}Alan Yuille¹¹Johns Hopkins University²Max Planck Institute for Informatics³University of Freiburg

Abstract


Human pose and shape (HPS) estimation methods achieve remarkable results. However, current HPS benchmarks are mostly designed to test models in scenarios that are similar to the training data. This can lead to critical situations in real-world applications when the observed data differs significantly from the training data and hence is out-of-distribution (OOD). It is therefore important to test and improve the OOD robustness of HPS methods. To address this fundamental problem, we develop a simulator that can be controlled in a fine-grained manner using interpretable parameters to explore the manifold of images of human pose, e.g. by varying poses, shapes, and clothes. We introduce a learning-based testing method, termed PoseExaminer, that automatically diagnoses HPS algorithms by searching over the parameter space of human pose images to find the failure modes. Our strategy for exploring this high-dimensional parameter space is a multi-agent reinforcement learning system, in which the agents collaborate to explore different parts of the parameter space. We show that our PoseExaminer discovers a variety of limitations in current state-of-the-art models that are relevant in real-world scenarios but are missed by current benchmarks. For example, it finds large regions of realistic human poses that are not predicted correctly, as well as reduced performance for humans with skinny and corpulent body shapes. In addition, we show that fine-tuning HPS methods by exploiting the failure modes found by PoseExaminer improve their robustness and even their performance on standard benchmarks by a significant margin. The code are available for research purposes at <https://github.com/qihao067/PoseExaminer>.

1. Introduction

In recent years, the computer vision community has made significant advances in 3D human pose and shape (HPS) estimation. But despite the high performance on standard benchmarks, current methods fail to give reliable predictions for configurations that have not been trained on



(a) An example of failure modes on articulated pose found by PoseExaminer



(b) Training with PoseExaminer

Figure 1. PoseExaminer is an automatic testing tool used to study the performance and robustness of HPS methods in terms of articulated pose, shape, global rotation, occlusion, etc. It systematically explores the parameter space and discovers a variety of failure modes. (a) illustrates three failure modes in PARE [18] and (b) shows the efficacy of training with PoseExaminer.

or in difficult viewing conditions such as when humans are significantly occluded [18] or have unusual poses or clothing [34]. This lack of robustness in such out-of-distribution (OOD) situations, which typically would not fool a human observer, is generally acknowledged and is a fundamental open problem for HPS methods that should be addressed.

The main obstacle, however, to testing the robustness of HPS methods is that test data is limited, because it is expensive to collect and annotate. One way to address this problem is to diagnose HPS systems by generating large synthetic datasets that randomly generate images of humans in varying poses, clothing, and background [4, 34]. These studies suggest that the gap between synthetic and real do-

mains in HPS is sufficiently small for testing on synthetic to be a reasonable strategy, as we confirm in our experiments.

These methods, however, only measure the *average performance* on the space of images. In sensitive domains that use HPS (e.g. autonomous driving), the *failure modes*, where performance is low, can matter more. Most importantly, despite large scale, these datasets lack diversity in some dimensions. For example, they mostly study pose for common actions such as standing, sitting, walking, etc.

Inspired by the literature on adversarial machine learning [5, 11], recent approaches aim to systematically explore the parameter space of simulators for weaknesses in the model performance. This has proven to be an effective approach for diagnosing limitations in image classification [45], face recognition [43], and path planning [39]. However, they are not directly applicable to testing the robustness of the high-dimensional regression task of pose estimation. For example, [43, 45] were designed for simple binary classification tasks and can only optimize a limited number of parameters, and [39] only perturbs an initial real-world scene to generate adversarial examples, which does not enable the full exploration of the parameter space.

In this work, we introduce PoseExaminer (Fig. 1), a learning-based algorithm to automatically diagnose the robustness of HPS methods. It efficiently searches through the high-dimensional continuous space of a simulator of human images to find failure modes. The strategy in PoseExaminer is a multi-agent reinforcement learning approach, in which the agents collaborate with one another to search the latent parameter space for model weaknesses. More specifically, each agent starts at different random initialized seeds and explores the latent parameter space to find failure cases, while at the same time avoiding exploring the regions close to the other agents in the latent space. This strategy enables a highly parallelizable way of exploring the high-dimensional continuous latent space of the simulator. After converging to a local optimum, each agent explores the local parameter space to find a connected failure region that defines a whole subspace of images where the pose is incorrectly predicted (*i.e.* failure mode). We demonstrate that very large subspaces of failures exist even in the best HPS models. We use the size of these failure subspaces together with the success rate of the agents as a new measure of out-of-distribution robustness.

Our experiments on four state-of-the-art HPS models show that PoseExaminer successfully discovers a variety of failure modes that provide new insights about their real-world performance. For example, it finds large subspaces of realistic human poses that are not predicted correctly, and reduced performance for humans with skinny and corpulent body shapes. Notably, we find that the failure modes found in synthetic data generalize well to real images: In addition to using PoseExaminer as a new benchmark, we also

find that fine-tuning SOTA methods using the failure modes discovered by PoseExaminer enhances their robustness and even the performance on 3DPW [52] and AIST++ [25, 50] benchmarks. We also note that computer graphics rendering pipelines become increasingly realistic, which will directly benefit the quality of our automated testing approach.

In short, our contributions are three-fold:

- We propose PoseExaminer, a learning-based algorithm to automatically diagnose the robustness of human pose and shape estimation methods. Compared to prior work, PoseExaminer is the first to efficiently search for a variety of failure modes in the high-dimensional continuous parameter space of a simulator using a multi-agent reinforcement learning framework.
- We introduce new metrics, which have not been possible to measure before, for quantifying the robustness of HPS methods based on our automated testing framework. We perform an in-depth analysis of current SOTA methods, revealing a variety of diverse failure modes that generalize well to real images.
- We show that the failure modes discovered by PoseExaminer can be used to significantly improve the real-world performance and robustness of current methods.

2. Related work

Human Pose and Shape Estimation. Current methods typically follow one of two paradigms: *Regression-based methods* [12, 17, 19, 32, 36, 48, 51] directly estimate 3D human parameters from RGB images. However, due to the difficulty of collecting highly accurate 3D ground truth, there aren't enough labeled training examples. One solution is to use weak supervision such as 2D keypoints, silhouettes, and body part segmentation in the training either as auxiliary consistency loss [17, 36, 48, 51, 55] or intermediate representations [1, 32, 36]. In addition, pseudo-groundtruth [15] labels are demonstrated to be effective in improving the generalization of regression-based methods like HMR [17]. *Optimization-based methods* [3, 35, 54] require parametric human models like SMPL [2, 26, 35]. SMPLify [3] is the first automated method to fit 2D keypoint detection to the SMPL model. After that, a lot of work has been done to utilize more information during the fitting procedure [13, 16, 24]. Also, the optimization-based model can be used to provide better supervision for training: SPIN [20] combines HMR [17] with SMPLify [3] in a training loop. At each step, HMR initializes SMPLify and then the latter provides better supervision for the former.

These methods make great success on standard benchmarks, but the difficulty in obtaining annotations also causes limited test sets with small diversity, making it difficult to evaluate how close the field is to fully robust

and general solutions [34]. Recognizing this issue, recent work [4, 34] builds synthetic datasets with various subjects for training and testing, yielding promising results. However, these fixed datasets mainly focus on common poses like sitting and standing, and they only care about the average performance, while we consider the failure modes under a large variety of poses. More importantly, PoseExaminer takes a big step forward and automatically diagnoses the performance and robustness in a highly efficient way.

Testing CV with Synthetic Data. Although there is a relative paucity of work in this area, testing computer vision models with synthetic data is not a novel idea [14, 21, 22, 31, 37, 38, 44, 58]. Previous work mainly focuses on training CV models with synthetic images [7–10, 23, 30, 40, 41]. In contrast to this body of work, we propose to search the parameter space of a human simulator in order to test the HPS model in an adversarial manner. However, different from the traditional adversarial attacks [5, 11, 28, 33, 42, 47], we focus on finding OOD examples that are realistic and in some cases, common in real life, but still fool the models. More importantly, our method can find not only a single isolated adversarial example, but also subspaces of parameters that lie in the latent space of a human simulator.

Recently, there is very interesting work that systematically explores the parameter space of simulators for weaknesses in the model performance [39, 43, 45]. However, [39] only perturbs an initial real-world scene to generate adversarial examples, which does not enable the full exploration of the parameter space, and [43, 45] are designed for simple binary classification tasks and only optimize a limited number of parameters. We focus on human pose and shape, a more difficult high-dimensional regression task. More importantly, in comparison to [43, 45], we simultaneously search the entire parameter space for as many various failure modes as possible and determine the exact boundary of them, and we also demonstrate that the failure modes discovered by PoseExaminer are beneficial to solving real-world problems by revealing the limitations of current approaches and improving their robustness and performance.

3. Adversarial Examiner for HPS

PoseExaminer aims to search the entire parameter space for as many different failure modes as possible, and then determine the boundaries of all failure subspaces. For our purpose, we first need to define the parameter spaces \mathcal{P} of a simulator to control different features of human images (Sec. 3.1). Then we need an efficient policy to search over the parameter spaces for a variety of failure modes. To do so, we divide the search process into two phases: We first localize as many different failure cases as possible (Sec. 3.2), and then find the boundaries for all failure cases to form subspaces (Sec. 3.3). An illustration is provided in Fig. 2.




Figure 2. **PoseExaminer model pipeline.** Multiple RL agents collaborate to search for worst cases generated by a policy π_{ω} and find subspaces boundaries defined by ϕ_{up} and ϕ_{low} . The human simulator is conditioned on parameters generated by the agents. m images are then generated for each agent and are used to test a given HPS method. The prediction error of the HPS model serves as the reward signal to update the policy parameters of agents.

3.1. Human Image Generation

To represent human shape and pose, we use SMPL [27] parameters. SMPL represents the body pose and shape by Θ , which consists of the pose $\theta \in \mathbb{R}^{72}$ and shape $\beta \in \mathbb{R}^{10}$ parameters. Here we only consider the pose of human body, thus we omit the last two joints (*i.e.* the hands). To better test the HPS methods, we divide the rest of the pose parameters into two parts: global rotation $\theta_g \in \mathbb{R}^3$ and the articulated pose $\theta_a \in \mathbb{R}^{63}$. The former represents the pose vector of the pelvis and the latter represents the vector of the remaining 21 joints. Given these parameters, the SMPL model is a differentiable function that outputs a posed 3D mesh $\mathcal{M}(\theta, \beta) \in \mathbb{R}^{6890 \times 3}$. We use PyTorch3D to render 3D human mesh. During rendering, we also consider the skin color $c_{skin} \in \mathbb{R}^3$ which consists of the RGB value(scale) of each vertex, and clothing ψ_{text} that is used to generate UV maps to texture the mesh surface.

In addition to the parameters related to humans, we also consider the background ψ_{bg} which controls the background images, the occluder $\{\psi_{occ}, t_{occ}\}$, in which ψ_{occ} controls the shape and texture and $t_{occ} \in \mathbb{R}^3$ consists of the position (u_{occ}, v_{occ}) and rotation angle r_{occ} of the patch. Finally, we also consider the lighting $l \in \mathbb{R}^4$ of which the elements control the xyz position and the intensity.

Note that our preliminary experiments suggest that common and noiseless backgrounds, clothing, and textures of occluders are not the key factors affecting performance compared to other factors (see Supp.), and generating realistic images of them via a generative model with parameters $\psi_{(\cdot)}$ is time-consuming. Therefore, following PARE [18] that directly uses occluding patches, we collect real images including 150 backgrounds of 10 scenarios, 40 high-quality UV maps, and 12 occluders and use $\psi_{(\cdot)}$ as a dictionary.

In this work, we mainly focus on single-person HPS estimation, and because most methods require cropped images with the target person in the center, we directly use a per-

spective camera with focal lengths 60 mm and fixed scale and translation parameters. We can study the viewpoints by changing the global rotation θ_g .

3.2. Finding the Worst-case Poses

Now we have defined all the parameters required to generate images, we want to efficiently search the space, finding parameters such that the samples generated with them fool a given algorithm and the errors are above the threshold T . We take the articulated pose θ_a as an example in the following two subsections. Other parameters are much easier to study. We want to emphasize that our solution is general and can handle all objectives mentioned above. We analyze all parameters in the experiments section.

PoseExaminer has no knowledge of the given HPS method H and the simulator S is non-differentiable, so we use reinforcement learning (RL) to search for the adversarial parameters. We adopt a policy gradient method. To efficiently handle the high-dimensional pose space and learn as many different failure modes as possible, we employ n RL agents to search cooperatively. Another challenge is to ensure the poses we find are physically possible. Thus, we use VPoser [35], a VAE-based pose prior, to constrain the pose and let the agents search the 32-dimensional latent space in this phase. Other advanced priors can be used here. The prior allows our model to start from different simple poses and learn valid and relatively common adversarial examples. The limitation of using a prior that trained on a fixed dataset such as AMASS [29] is solved in phase 2 (Sec. 3.3).

Specifically, for agent i , we define a policy π_{ω^i} parameterized by ω that we can sample parameters $z^i \sim \pi_{\omega^i}(z^i)$. Then we use the pre-trained decoder f of VPoser to get the articulated pose $\theta_a^i = f(z^i)$, and render an image I^i via simulator S , $I^i = S(\theta_a^i, \Psi^i)$, where Ψ^i denotes the other controllable parameters in Sec. 3.1. After that we can test the method H on I^i and get the error. To avoid the ambiguity along depth, we use 2D MPJPE err_{2D}^i so that the agents only focus on finding poses that cause large error even after being projected to 2D. We define a reward $R(z^i) = c - err_{2D}^i$ where c is a constant. Then we can optimize the parameters ω^i by maximizing

$$L(\omega^i) = \mathbb{E}_{z^i \sim \pi_{\omega^i}}[R(z^i)] + \mathbb{1}_{\{i \neq b\}} \gamma \mathbb{E}[D(\pi_{\omega^i}, \pi_{\omega^b})] \quad (1)$$

where $\mathbb{1}_{\{\cdot\}}$ denotes the indicator function, $\gamma (> 0)$ denotes a weighting factor, and $D(\pi_{\omega}, \pi_{\omega'})$ is the distance measure between two policies π_{ω} and $\pi_{\omega'}$ (here we use the $L2$ distance between z and z').

The gradient of the second term in Eq. 1 is easy to obtain by computing the mean distance between agents. Gradient of the first term is computed following the REINFORCE rule

$$\nabla_{\omega^i} J(\omega^i) = \mathbb{E}_{z^i \sim \pi_{\omega^i}}[\nabla_{\omega^i} \log(\pi_{\omega^i}) R(z^i)] \quad (2)$$

and an unbiased, empirical estimate of Eq. 2 is:

$$\nabla_{\omega^i} \hat{J}(\omega^i) = \frac{1}{K} \sum_{k=1}^K \nabla_{\omega^i} \log(\pi_{\omega^i})(R(z^i) - b) \quad (3)$$

in which K is the number of parameters sampled and b is a baseline and is updated every iteration by $b = (1 - \tau)b + \tau R(z^i)$. Please see Supp. for the pseudo-code.

3.3. Finding Failure Subspaces

After learning multiple failure cases in phase 1, we search the space around these points for failure subspaces. However, VPoser has several limitations. For example, the distances between individual poses are not preserved in the VAE latent space, implying that the space is not necessarily continuous or smooth mathematically. More importantly, it is trained on a finite dataset, resulting in a biased distribution. Therefore, in this phase, we directly search the pose space using a hierarchical searching policy. We also introduce two new metrics to evaluate the subspaces and boundaries we discover (Sec. 4.1), and demonstrate the effectiveness of our searching policy.

Assume agent i has discovered an adversarial point z^i , then we have the corresponding adversarial pose $\theta_a^i = f(z^i)$, $\theta_a^i \in \mathbb{R}^{63}$. We define two non-negative variables $\phi_{up}^i \in \mathbb{R}^{63}$ and $\phi_{low}^i \in \mathbb{R}^{63}$ that denote the upper bound and the lower bound respectively. To ensure searching efficiency and accuracy, for each iteration, we select one joint, e.g. joint j , according to the kinematic tree, from parent to child nodes, and only consider the space around the upper bound or the lower bound of the three rotation directions of the selected joint (i.e. 3 out of 63×2 dimensions). Given a step size $\delta \in \mathbb{R}^{63}$, we can sample m poses $\theta_{a,j}^i$ either from $\mathcal{U}(\theta_a^i + \phi_{up}^i, \theta_a^i + \phi_{up}^i + \delta)$ to optimize upper bound, or from $\mathcal{U}(\theta_a^i - \phi_{low}^i - \delta, \theta_a^i - \phi_{low}^i)$ to optimize lower bound. The subscript j means only the elements corresponding to joint j are different in these samples. After that, we test the method H and compute the errors. To ensure the poses are physically possible, we consider the limits of joint angles and use Bounding Volume Hierarchies [49] to avoid self-collisions. We update the boundary when the minimum error of all samples is higher than the adversarial threshold T and the pose is valid. Note that the latter is computed at joint level. For computational efficiency, we use a highly parallelized implementation of BVH following VPoser [35]. The step size decreases linearly w.r.t the minimum error. Pseudo-code and details of this phase are provided in Supp.

4. Benchmarking and Improving Robustness with PoseExaminer

PoseExaminer contributes to the field of HPS in two aspects: (1) It provides a highly efficient way to systematically study the performance, robustness, and weaknesses of

a given method (Sec. 4.1). (2) It can be used to improve the robustness and performance of existing methods (Sec. 4.2).

4.1. Quantifying Robustness with PoseExaminer

Using PoseExaminer to diagnose HPS methods is intuitive. All parameters in Sec. 3.1 can be studied separately or jointly by optimizing one or more types of parameters while fixing the others. Note that by using different sets of hyperparameters of PoseExaminer, such as the searching boundary of policy π_ω , the adversarial threshold T , the limits of joint angles, and so on, we can get examiners with different difficulty levels. For a fair comparison, we select one set of hyperparameters for all benchmark experiments. Please refer to the Supp. for details.

Metrics. We introduce several new metrics in this work. Two of them are used to assess the quality of failure modes and their boundaries learned by PoseExaminer, and the others are used to analyze the robustness of HPS methods.

To assess the quality of boundaries, we sample $n \times 200$ examples from the failure subspaces and test the given method on these examples (n is the number of agents). **Percentage of non-adversarial examples (Pnae)** represents the percentage of examples with 3D MPJPE smaller than the adversarial threshold T . We also consider **minimum of the mean per joint position error over failure modes (minMPJPE)**. Usually minMPJPE is slightly smaller than the adversarial threshold T . Pnae ensures that the boundaries are ‘small’ enough and no non-adversarial examples are included, while minMPJPE ensures that the boundaries are ‘large/accurate’ enough to be close to the threshold.

To evaluate the robustness of a given method, we consider the failure modes PoseExaminer learns. **Region Size** denotes the L_2 distance between the upper and lower bounds of the subspaces, *i.e.* $\frac{1}{n} \sum_{i=1}^n \|\phi_{up}^i + \phi_{low}^i\|_2$ (here ϕ_{up}^i and ϕ_{low}^i are non-negative). **Success Rate** refers to the percentage of agents that find failure modes in specific time steps. We also consider the **mean MPJPE over failure modes (meanMPJPE)** and the **maximum MPJPE over failure modes (maxMPJPE)** that describe the average and worst performance under failure modes. A method is more robust if these metrics are lower.

4.2. Improving Robustness Using Failure Subspaces

With the failure modes being discovered, the method can be improved by fine-tuning it using the detected modes. An intuitive way is to directly sample poses from the learned subspaces and add them to training data. However, as mentioned previously, current in-distribution (ID) datasets are limited and only cover a small space, while PoseExaminer searches a large space and finds many failure modes with different types and difficulty levels. Due to model capacity and the forgetting problem, directly using a strict examiner could harm the performance on standard benchmarks and

simple poses (see Sec. 5.5). To make the fine-tuning procedure stable and achieve good results on both ID and OOD datasets, we use curriculum learning and provide an efficient way to use PoseExaminer to improve algorithms.

Specifically, we fine-tune the model following an easy-to-difficult schedule, using a pre-defined ordered list of hyperparameters that initializes examiners taking into account different factors and difficulty levels. For each loop, after finding multiple failure modes, we sample m cases per failure mode and add them to an adversary set \mathcal{F} , then we fine-tune the given method on a weighted combination of data from the original training set \mathcal{T} and the adversary set \mathcal{F} . For each loop, ϵ percent of data are sampled from \mathcal{F} and $(1 - \epsilon)$ percent of data are from \mathcal{T} . Pseudo-code of fine-tuning with PoseExaminer is provided in the Supp.

We apply a discount factor α to the learning rate during fine-tuning. Since we render all images with a fixed camera, we remove the loss terms of 2D keypoints and camera parameters and only consider the terms directly related to the SMPL parameters and 3D keypoints. For other details of the hyperparameters, data augmentations, and training procedures, we strictly follow the original methods.

5. Experiments

We then highlight the new capabilities our PoseExaminer enables. Sec. 5.1 shows that the failure modes discovered by PoseExaminer generalize well to real images. Then we use our method to diagnose four HPS methods and provide in-depth analyses in Sec. 5.2 and 5.3. The failure modes are then used to improve the current arts in Sec. 5.4 and 5.5.

Implementation Details. We use a multi-variable Gaussian policy $\pi_\omega(\cdot) = \mathcal{N}(\mu_\pi, \sigma_\pi^2)$ in phase one. The variance is fixed $\sigma_\pi^2 = 0.05 \times I$ and the mean μ_π is learned by a single layer perceptron. We set $T = 90$ as the adversarial threshold. $c = 50$ and $\gamma = 0.2$. In the method analyses in Sec. 5.2 and 5.3, we employ $n = 320$ agents and each agent samples $m = 50$ failure cases in each iteration during training, and $m = 200$ to compute the metrics in Sec 4.1. In the training experiments in Sec. 5.4 and 5.5, we set $n = 40$ and $m = 500$ to generate the adversary set \mathcal{F} . Adam optimizer with a learning rate of 0.2 is used. Each agent is trained for at most 300 iterations in phase one and 600 in total.

Models. We consider four HPS methods: HMR [17], SPIN [20], HMR-EFT [15], and PARE [18]. In the method analyses in Sec. 5.2 and 5.3, we directly use the official pre-trained model and/or PyTorch-based implementations provided in MMHuman3D [6]. In the training experiments in Sec. 5.4 and 5.5, we fine-tune these methods strictly following their original implementation details, with a discount factor $\alpha = 0.05$ applied to the original learning rate. We set sample rate $\epsilon = 0.1$, *i.e.* 10% of the training data is from the adversary set \mathcal{F} and 90% is from the original training sets.

Datasets. The 3DPW test split [52] and AIST++ [25, 50]

	real 3DPW			sync 3DPW		
	MPJPE	PA-MPJPE	PVE	MPJPE	PA-MPJPE	PVE
SPIN [20]	96.9	59.2	135.1	93.9 (-3.0)	60.9 (+1.7)	130.9 (-4.2)
PARE [18]	82.0	50.9	97.9	79.6 (-2.4)	51.7 (+0.8)	94.5 (-3.4)
PARE(w. 3DPW) [18]	74.5	46.5	88.6	75.8 (+1.3)	49.4 (+2.9)	89.5 (+0.9)

Table 1. **Measuring the performance gap of SPIN and PARE on real and synthetic images.** We can observe that both models perform very well on synthetic data, which has exactly the same poses as the real dataset, indicating that the domain gap is small and even favorable towards the synthetic data.

are used for evaluation. 3DPW contains subjects performing common actions such as sitting, walking, etc, and is thus classified as an ID dataset. AIST++ is a more complex dataset in which professional dancers dance in front of a white background, surrounded by nine cameras. It contains a large variety of poses and is used as an ODD dataset. We convert the videos to images with 60 FPS following [25].

A disadvantage of the AIST++ dataset is that it only contains pseudo labels, and some have completely incorrect annotations. We filter out these images by ensuring (1) the consistency of the SMPL, 3D keypoints, and 2D keypoints annotations, (2) the smoothness of the depth estimates, and (3) the smoothness of the 3D joints on the face. Then we manually check 1000 randomly selected images to ensure the quality of the annotations. (Please see Supp. for details)

The remaining annotations are sufficiently accurate and can be used for evaluation because they are regressed from images captured by nine cameras surrounding the subjects. We randomly select 40K images from the clean AIST++ images and named them cAIST. Also, we use the Vposer encoder to obtain the latent distribution of each pose in the clean AIST++ images, select poses that have more than 15 dimensions latent parameters larger than 1, and randomly select 40K poses from them to form a subset named cAIST-ext, which mainly contains extreme and OOD poses.

Evaluation Metrics. In addition to the new metrics provided in Sec 4.1, we also report Procrustes-aligned mean per joint position error (PA-MPJPE) and mean per joint position error (MPJPE) in *mm*. We only report per vertex error (PVE) for 3DPW because AIST++ does not provide shape annotations and it uses mean shape instead.

5.1. Generalization from Synthetic to Real

In this section, we demonstrate that (1) with careful design, the performance gap between real and synthetic is small on both ID and OOD datasets, proving that evaluation on synthetic is worthwhile, and (2) the failure modes discovered by PoseExaminer generalize well to real images.

Performance Gap between Real and Synthetic. To generate synthetic versions of the 3DPW and cAIST-EXT datasets, we need to obtain the original background image, and then render ground-truth human meshes on top of it. We first use a video instance segmentation method [53] to




Figure 3. **Failure modes generalize well to real images.** We compare the failure modes discovered by PoseExaminer (b) with images from the cAIST dataset (a), and visualize the results on two groups of images that have similar poses. We observe that similar poses across synthetic and real domains result in similar errors. More visualizations of the failure modes are provided in Supp.

remove the people from the image, and then use image inpainting [46] to fill in the gaps left by the removal of the people. After that, for each video, we select UV textures that look the most like the removed people in order to render the synthetic human onto the background image. We also consider the overlapped masks when running the segmentation method, and add an occluding patch that looks similar to the original occluder in the image to restore the occlusions. (Please see the Supp. for figures of each step.)

We evaluate the official pretrained models of SPIN and PARE on both synthetic and real versions of 3DPW and cAIST-EXT datasets and provide results in Tab. 1 and Supp. We can see that all models achieve similar performance on the corresponding synthetic and real versions, no matter on the ID or OOD dataset, or on common (simple) or extreme (complex) poses. PVE on synthetic 3DPW is much smaller than on real. We think this is because the shape annotations are more difficult to get and thus less accurate, whereas synthetic data has perfect shape annotations. Even so, the performance gap between real and synthetic are small.

Failure Modes Generalize Well to Real Images. After running PoseExaminer and finding failure modes of SPIN and PARE, we compare them with the real images in the cAIST-EXT dataset and find several real images with people in very similar poses. We visualize the results in Fig. 3. The HPS algorithms perform similarly in these examples, and more importantly, Fig. 3 suggests that a method may generate the same error patterns on the same failure mode, even across the domains of synthetic and real, which **demonstrates that PoseExaminer is able to provide insights into the model’s real-world behaviors.**

5.2. Robustness towards OOD Poses

Articulated pose is a very important source of variability. It has a very large and high-dimensional space. Efficiently

	Model robustness				Subspace quality	
	Succ. Rate ↓	Region Size ↓	meanMPJPE ↓	maxMPJPE ↓	minMPJPE	Pnae
HMR [17]	97.5%	8.521	225.03	448.37	83.55	0.4%
SPIN [20]	95.0%	5.417	169.64	352.69	84.14	0.9%
HMR-EFT [15]	84.1%	4.675	161.35	337.46	81.84	1.2%
PARE [18]	75.3%	3.948	150.06	300.76	81.81	1.3%

Table 2. **Robustness towards adversarial examination of articulated pose.** We compare success rate, region size in radian, and MPJPE. PARE outperforms other algorithms on all these metrics, but compared to its good results on 3DPW, the high success rate and large region size indicate that it still has some weaknesses which could not be detected by standard testing.

searching the articulated pose space while avoiding invalid poses is one of our main challenges. Meanwhile, despite being a very critical factor, little work has been done on the OOD poses. In this section, we investigate the robustness of four methods to articulated pose $\theta_a \in \mathbb{R}^{63}$, which also proves the efficiency and effectiveness of PoseExaminer.

Since we only consider articulated poses here, we use the simplest set of parameters for other factors (*i.e.*, mean body shape, white background image, standard clothing, no occlusion, and no global rotation so that the subject is facing the camera). The results are shown in Tab. 2

Effectiveness of PoseExaminer. The success rate in Tab. 2 indicates that PoseExaminer is able to localize failure cases very efficiently. Even with the most robust method (PARE), 75.3% of agents (241 agents) discover failure cases in no more than 300 iterations. Then we evaluate the boundary quality. We can see that the minMPJPE is slightly smaller than the threshold $T = 90$, indicating that the agent has already found a large enough boundary for each failure case. Meanwhile, Pnae shows that there is only about 1% of poses in these subspaces result in errors smaller than T , indicating that the boundary is accurate enough that only a few non-adversarial examples are included in the subspaces.

In short, **we show that with a total of 600 iterations, PoseExaminer is able to efficiently search the complex high-dimensional space, discover a variety of failure modes, and accurately determine their boundaries.** Due to space constraints, we only provide two failure modes of SPIN and PARE in Fig. 3 and three modes of PARE in Fig. 1. More examples are provided in the Supp.

Testing HPS Methods on OOD Poses. From Tab. 2 we can find that current methods are not robust to OOD poses even when only considering articulated pose and excluding all other challenges such as occlusion, orientation, uncommon body shape, and so on. PARE is the most robust method, but it still has many failure modes and fails to give accurate estimates to some realistic but possibly less common poses.

5.3. Robustness towards Body Shape and Occlusion

As previously stated, our method can be used to study any controllable factor of the rendering pipeline. We also

	Without occluder		With small occluder	
	Succ. Rate ↓	Region Size ↓	Succ. Rate ↓	Region Size ↓
HMR-EFT [15]	84.1%	4.675	100%	7.745 (+3.070)
PARE [18]	75.3%	3.948	94.5%	5.522 (+1.574)

Table 3. **Robustness towards occlusion.** PARE is more robust to occlusion compared with HMR-EFT.




Figure 4. **Body shape bias on PARE.** (a) shows the relations between MPJPE and aBMI of all samples during optimization and (b) shows the human meshes of one agent in the first and final iterations and the estimates given by PARE. We observe a large drop in performance for the current SOTA method on humans with skinny and corpulent body shapes.

analyze all other factors in this work, but we only include results on body shape and occlusion in this section due to space limits. Results for global rotation, clothing, background, lighting, and gender are provided in the Supp.

Shape. To study the robustness to body shape, we consider two experiments: (1) only optimizing shape parameters $\beta \in \mathbb{R}^{10}$ while the pose $\theta_a \in \mathbb{R}^{63}$ are randomly initialized and fixed, and (2) optimizing both shape and articulated pose simultaneously. We initialize and freeze other parameters following Sec. 5.2. We find that all agents eventually converge on either a skinny or a corpulent human body. To quantify the results, we borrow the concept of Body Mass Index (BMI) that measures leanness or corpulence based on height and weight, and compute approximate Body Mass Index (aBMI): $aBMI = volume/height^2$. Given a shape parameter, we render a human mesh with the canonical pose, and then we can read the height and use [56] to compute its volume and get the aBMI of this shape. Fig. 4 shows the relations between MPJPE and aBMI of all samples during optimization, as well as some visualization results.

Occlusion. Occlusion is a popular problem in HPS [18, 57]. We optimize the pose parameters $\theta_a \in \mathbb{R}^{63}$ and occlusion parameters $t_{occ} \in \mathbb{R}^3$ simultaneously and compared the robustness of PARE and HMR-EFT. Results are shown in Tab. 3. PARE outperforms HMR-EFT in robustness, but in some hard poses, even minor occlusion may cause large errors. (See Supp. for visualizations.) This implies that current methods solve occlusion partly based on the pose distribution of training data. The problem becomes worse when a person with an OOD pose is occluded.

	Robustness (PoseExaminer)		ID dataset (3DPW)			OOD dataset (cAIST)		Extreme poses (cAIST-EXT)	
	Succ. Rate↓	Region Size↓	MPJPE↓	PA-MPJPE↓	PVE↓	MPJPE↓	PA-MPJPE↓	MPJPE↓	PA-MPJPE↓
SPIN [20]	95.0%	5.417	94.11	57.54	111.12	108.09	68.59	133.65	81.03
+ PE (Ours)	76.6% (-18.4)	3.964 (-1.435)	88.66 (-5.45)	54.34 (-3.20)	103.94 (-7.18)	98.88 (-9.21)	65.28 (-3.31)	120.98 (-12.67)	76.82 (-4.21)
HMR-EFT [15]	84.1%	4.675	85.23	51.88	107.88	98.55	66.01	122.19	78.39
+ PE (Ours)	63.6% (-20.5)	3.429 (-1.246)	80.41 (-4.82)	49.15 (-2.73)	101.43 (-6.45)	88.53 (-10.02)	64.00 (-2.01)	112.78 (-9.41)	75.15 (-3.24)
PARE [18]	75.3%	3.948	81.81	50.78	102.27	99.15	62.43	117.45	72.68
+ PE (Ours)	48.4% (-26.9)	2.953 (-0.995)	77.46 (-4.35)	48.01 (-2.77)	94.86 (-7.41)	87.44 (-11.71)	59.80 (-2.63)	109.82 (-7.63)	70.73 (-1.95)

Table 4. **Fine-tuning with PoseExaminer.** ‘PE’ is short for PoseExaminer. Despite being synthetic, fine-tuning on the failure modes discovered by PoseExaminer can significantly improve the robustness and real-world performance of all methods on all benchmarks.

	Robustness (PoseExaminer)		ID dataset (3DPW)			OOD dataset (cAIST)		Extreme poses (cAIST-EXT)	
	Succ. Rate↓	Region Size↓	MPJPE↓	PA-MPJPE↓	PVE↓	MPJPE↓	PA-MPJPE↓	MPJPE↓	PA-MPJPE↓
PARE [18]	75.3%	3.948	81.81	50.78	102.27	99.15	62.43	117.45	72.68
PARE + PE (random)	69.5%	3.542	79.77	50.20	97.80	95.10	62.40	115.54	73.25
PARE + PE (easy)	65.4%	3.373	74.44	47.86	92.94	105.09	59.95	129.77	71.84
PARE + PE (hard)	46.7%	<u>3.005</u>	81.15	48.54	100.08	86.92	59.03	107.13	70.15
PARE + PE (mixed)	48.4%	2.956	<u>77.46</u>	48.01	<u>94.86</u>	<u>87.44</u>	<u>59.80</u>	<u>109.82</u>	<u>70.73</u>

Table 5. **Ablation study.** ‘Random’ is short for ‘random sampling’. PoseExaminer with different difficulty levels leads to different results on different real-world benchmarks. Best in **bold**, second best underlined.

5.4. Improving HPS Methods with PoseExaminer

To evaluate the efficacy of fine-tuning with PoseExaminer, we fine-tune the pretrained models of SPIN, HMR-EFT, and PARE in Tab. 4. Training with the failure modes discovered by PoseExaminer is very effective. It significantly improves robustness and performance on both ID and OOD datasets. Note that the failure modes we learned here mainly focus on hard and extreme poses, we are surprised to see that performance on the 3DPW dataset is improved as well. In fact, we can use a simpler examiner to further boost the performance on 3DPW (See Sec. 5.5).

5.5. Ablation Study

Training with Different Settings of PoseExaminer. As mentioned above, fine-tuning PARE with PoseExaminer of different difficulty levels yields different results. We define two levels of difficulty here: easy and hard. We set a more strict restriction to the easy examiner so that it can only search a very small parameter space, whereas the hard examiner searches a much larger parameter space. We also use curriculum learning to combine these two examiners and name it ‘mixed’ (*i.e.* the PoseExaminer used in Tab. 4). The results are provided in Tab. 5. Fine-tuning PARE with easy PoseExaminer significantly improves performance on 3DPW but degrades performance on cAIST, while hard PoseExaminer benefits more on the cAIST dataset. This suggests that a model with larger capability will benefit more from fine-tuning with PoseExaminer.

Comparison to Random Sampling. We compare our RL-based searching with random sampling in Tab. 5. For random sampling, instead of sampling parameters from the discovered failure modes, we directly sample the same number of parameters from the entire space. Random sampling also

improves performance on all benchmarks. However, PoseExaminer (‘mixed’) improves the performance and robustness by a much larger margin, especially on OOD datasets, demonstrating the efficacy of training on failure modes.

6. Conclusion

We have introduced PoseExaminer, a learning-based testing method that automatically diagnoses HPS algorithms by searching over the parameter space of human pose images for failure modes. These failure modes are relevant in real-world scenarios but are missed by current benchmarks. PoseExaminer enables a highly efficient way to study and quantify the robustness of HPS methods towards various factors such as articulated pose, thereby providing new insights for existing methods. Furthermore, despite being synthetic, fine-tuning on the failure modes discovered by PoseExaminer can significantly improve the robustness and real-world performance on all benchmarks.

Acknowledgements. AK acknowledges support via his Emmy Noether Research Group funded by the German Science Foundation (DFG) under Grant No. 468670075. QL acknowledges support from the NIH R01 EY029700 and Army Research Laboratory award W911NF2320008. This research is also based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via [2022-21102100005]. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The US. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [1] Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single rgb camera. In *CVPR*, pages 1175–1186, 2019. 2
- [2] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: Shape completion and animation of people. 2005. 2
- [3] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it SMPL: Automatic estimation of 3d human pose and shape from a single image. In *ECCV*, pages 561–578. Springer, 2016. 2
- [4] Zhongang Cai, Mingyuan Zhang, Jiawei Ren, Chen Wei, Daxuan Ren, Jiatong Li, Zhengyu Lin, Haiyu Zhao, Shuai Yi, Lei Yang, et al. Playing for 3d human recovery. *arXiv preprint arXiv:2110.07588*, 2021. 1, 3
- [5] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017. 2, 3
- [6] MMHuman3D Contributors. Openmmlab 3d human parametric model toolbox and benchmark. <https://github.com/open-mmlab/mmhuman3d>, 2021. 5
- [7] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *CoRL*, pages 1–16. PMLR, 2017. 3
- [8] Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *CVPR*, pages 4340–4349, 2016. 3
- [9] Baris Gecer, Binod Bhattarai, Josef Kittler, and Tae-Kyun Kim. Semi-supervised adversarial learning to generate photorealistic face images of new identities from 3d morphable model. In *ECCV*, pages 217–234, 2018. 3
- [10] Baris Gecer, Alexandros Lattas, Stylianos Ploumpis, Jiankang Deng, Athanasios Papaioannou, Stylianos Moschoglou, and Stefanos Zafeiriou. Synthesizing coupled 3d face modalities by trunk-branch generative adversarial networks. In *ECCV*, pages 415–433. Springer, 2020. 3
- [11] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015. 2, 3
- [12] Riza Alp Guler and Iasonas Kokkinos. HoloPose: Holistic 3D human reconstruction in-the-wild. In *CVPR*, pages 10884–10894, 2019. 2
- [13] Yinghao Huang, Federica Bogo, Christoph Lassner, Angjoo Kanazawa, Peter V. Gehler, Javier Romero, Ijaz Akhter, and Michael J. Black. Towards accurate marker-less human shape and pose estimation over time. In *3DV*, 2017. 2
- [14] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, pages 2901–2910, 2017. 3
- [15] Hanbyul Joo, Natalia Neverova, and Andrea Vedaldi. Exemplar fine-tuning for 3d human model fitting towards in-the-wild 3d human pose estimation. In *3DV*, pages 42–52. IEEE, 2021. 2, 5, 7, 8
- [16] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3D deformation model for tracking faces, hands, and bodies. In *CVPR*, pages 8320–8329, 2018. 2
- [17] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, pages 7122–7131, 2018. 2, 5, 7
- [18] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *ICCV*, pages 11127–11137, 2021. 1, 3, 5, 6, 7, 8
- [19] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, pages 2252–2261, 2019. 2
- [20] Nikos Kolotouros, Georgios Pavlakos, Michael J. Black, and Kostas Daniilidis. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *ICCV*, pages 2252–2261, 2019. 2, 5, 6, 7, 8
- [21] Adam Kortylewski, Bernhard Egger, Andreas Schneider, Thomas Gerig, Andreas Morel-Forster, and Thomas Vetter. Empirically analyzing the effect of dataset biases on deep face recognition systems. In *CVPRW*, pages 2093–2102, 2018. 3
- [22] A. Kortylewski, B. Egger, A. Schneider, T. Gerig, A. Morel-Forster, and T. Vetter. Analyzing and reducing the damage of dataset bias to face recognition with synthetic data. In *CVPR*, pages 2261–2268, 2019. 3
- [23] Adam Kortylewski, Andreas Schneider, Thomas Gerig, Bernhard Egger, Andreas Morel-Forster, and Thomas Vetter. Training deep face recognition systems with synthetic data. *arXiv preprint arXiv:1802.05891*, 2018. 3
- [24] 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 *CVPR*, pages 4704–4713, 2017. 2
- [25] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *ICCV*, pages 13401–13412, 2021. 2, 5, 6
- [26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. In *SIGGRAPH*, 2015. 2
- [27] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics*, 34(6):1–16, 2015. 3
- [28] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018. 3
- [29] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *ICCV*, pages 5442–5451, 2019. 4
- [30] Richard T Marriott, Sami Romdhani, and Liming Chen. A 3d gan for improved large-pose facial recognition. In *CVPR*, pages 13445–13455, 2021. 3

- [31] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR*, pages 4040–4048, 2016. 3
- [32] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter V. Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model-based human pose and shape estimation. 2018. 2
- [33] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, pages 506–519. ACM, 2017. 3
- [34] Priyanka Patel, Chun-Hao P Huang, Joachim Tesch, David T Hoffmann, Shashank Tripathi, and Michael J Black. Agora: Avatars in geography optimized for regression analysis. In *ICCV*, pages 13468–13478, 2021. 1, 3
- [35] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *CVPR*, pages 10975–10985, 2019. 2, 4
- [36] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3D human pose and shape from a single color image. In *CVPR*, pages 459–468, 2018. 2
- [37] Nicolas Pinto, James J DiCarlo, and David D Cox. Establishing good benchmarks and baselines for face recognition. In *Workshop on Faces In 'Real-Life' Images: Detection, Alignment, and Recognition*, 2008. 3
- [38] Weichao Qiu, Fangwei Zhong, Yi Zhang, Siyuan Qiao, Zihao Xiao, Tae Soo Kim, and Yizhou Wang. Unrealcv: Virtual worlds for computer vision. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1221–1224, 2017. 3
- [39] Davis Rempe, Jonah Philion, Leonidas J Guibas, Sanja Fidler, and Or Litany. Generating useful accident-prone driving scenarios via a learned traffic prior. In *CVPR*, pages 17305–17315, 2022. 2, 3
- [40] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV*, pages 102–118. Springer, 2016. 3
- [41] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, pages 3234–3243, 2016. 3
- [42] Nataniel Ruiz, Sarah Adel Bargal, and Stan Sclaroff. Disrupting deepfakes: Adversarial attacks against conditional image translation networks and facial manipulation systems. In *ECCV*, pages 236–251. Springer, 2020. 3
- [43] Nataniel Ruiz, Adam Kortylewski, Weichao Qiu, Cihang Xie, Sarah Adel Bargal, Alan Yuille, and Stan Sclaroff. Simulated adversarial testing of face recognition models. In *CVPR*, pages 4145–4155, 2022. 2, 3
- [44] Nataniel Ruiz, Barry-John Theobald, Anurag Ranjan, Ahmed Hussein Abdelaziz, and Nicholas Apostoloff. Morphgan: One-shot face synthesis gan for detecting recognition bias. In *BMVC*, 2021. 3
- [45] Michelle Shu, Chenxi Liu, Weichao Qiu, and Alan Yuille. Identifying model weakness with adversarial examiner. In *AAAI*, volume 34, pages 11998–12006, 2020. 2, 3
- [46] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *WACV*, pages 2149–2159, 2022. 6
- [47] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR*, 2014. 3
- [48] Jun Kai Vince Tan, Ignas Budvytis, and Roberto Cipolla. Indirect deep structured learning for 3D human shape and pose prediction. In *BMVC*, 2017. 2
- [49] Matthias Teschner, Stefan Kimmerle, Bruno Heidelberger, Gabriel Zachmann, Laks Raghupathi, Arnulph Fuhrmann, M-P Cani, François Faure, Nadia Magnenat-Thalmann, Wolfgang Strasser, et al. Collision detection for deformable objects. In *Computer graphics forum*, volume 24, pages 61–81. Wiley Online Library, 2005. 4
- [50] Shuhei Tsuchida, Satoru Fukayama, Masahiro Hamasaki, and Masataka Goto. Aist dance video database: Multi-genre, multi-dancer, and multi-camera database for dance information processing. In *ISMIR*, volume 1, page 6, 2019. 2, 5
- [51] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. In *NeurIPS*, pages 5236–5246, 2017. 2
- [52] Timo Von Marcard, Roberto Henschel, Michael J Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *ECCV*, pages 601–617, 2018. 2, 5
- [53] Junfeng Wu, Qihao Liu, Yi Jiang, Song Bai, Alan Yuille, and Xiang Bai. In defense of online models for video instance segmentation. In *ECCV*, pages 588–605. Springer, 2022. 6
- [54] Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. Monocular total capture: Posing face, body, and hands in the wild. In *CVPR*, pages 10965–10974, 2019. 2
- [55] Andrei Zanfir, Eduard Gabriel Bazavan, Hongyi Xu, Bill Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Weakly supervised 3d human pose and shape reconstruction with normalizing flows. In *ECCV*, pages 465–481, 2020. 2
- [56] Cha Zhang and Tsuhan Chen. Efficient feature extraction for 2d/3d objects in mesh representation. In *Proceedings 2001 International Conference on Image Processing (Cat. No. 01CH37205)*, volume 3, pages 935–938. IEEE, 2001. 7
- [57] Tianshu Zhang, Buzhen Huang, and Yangang Wang. Object-occluded human shape and pose estimation from a single color image. In *CVPR*, pages 7376–7385, 2020. 7
- [58] Yi Zhang, Weichao Qiu, Qi Chen, Xiaolin Hu, and Alan Yuille. Unrealstereo: Controlling hazardous factors to analyze stereo vision. In *3DV*, pages 228–237. IEEE, 2018. 3