

Resolving 3D Human Pose Ambiguities with 3D Scene Constraints

Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas and Michael J. Black
Max Planck Institute for Intelligent Systems

{mhassan, vchoutas, dtzionas, black}@tuebingen.mpg.de




Figure 1: Standard 3D body estimation methods predict bodies that may be inconsistent with the 3D scene even though the results may look reasonable from the camera viewpoint. To address this, we exploit the 3D scene structure and introduce *scene constraints* for *contact* and *inter-penetration*. From left to right: (1) RGB image (top) and 3D scene reconstruction (bottom), (2) overlay of estimated bodies on the original RGB image without (yellow) and with (gray) scene constraints, 3D rendering of both the body and the scene from (3) camera view, (4) top view and (5) side view.

Abstract

To understand and analyze human behavior, we need to capture humans moving in, and interacting with, the world. Most existing methods perform 3D human pose estimation without explicitly considering the scene. We observe however that the world constrains the body and vice-versa. To motivate this, we show that current 3D human pose estimation methods produce results that are not consistent with the 3D scene. Our key contribution is to exploit static 3D scene structure to better estimate human pose from monocular images. The method enforces Proximal Relationships with Object eXclusion and is called PROX. To test this, we collect a new dataset composed of 12 different 3D scenes and RGB sequences of 20 subjects moving in and interacting with the scenes. We represent human pose using the 3D human body model SMPL-X and extend SMPLify-X to estimate body pose using scene constraints. We make use of the 3D scene information by formulating two main constraints. The inter-penetration constraint penalizes intersection be-

tween the body model and the surrounding 3D scene. The contact constraint encourages specific parts of the body to be in contact with scene surfaces if they are close enough in distance and orientation. For quantitative evaluation we capture a separate dataset with 180 RGB frames in which the ground-truth body pose is estimated using a motion capture system. We show quantitatively that introducing scene constraints significantly reduces 3D joint error and vertex error. Our code and data are available for research at <https://prox.is.tue.mpg.de>.

1. Introduction

Humans move through, and interact with, the 3D world. The world limits this movement and provides opportunities (affordances) [20]. In fact, it is through contact between our feet and the environment that we are able to move at all. Whether simply standing, sitting, lying down, walking, or manipulating objects, our posture, movement, and behavior is affected by the world around us. Despite this, most work on 3D human pose estimation from images ignores

the world and our interactions with it.

Here we formulate human pose estimation differently, making the 3D world a first class player in the solution. Specifically we estimate 3D human pose from a *single RGB image* conditioned on the 3D scene. We show that the world provides constraints that make the 3D pose estimation problem easier and the results more accurate.

We follow two key principles to estimate 3D pose in the context of a 3D scene. First, from intuitive physics, two objects in 3D space cannot *inter-penetrate* and share the same space. Thus, we penalize poses in which the body interpenetrates scene objects. We formulate this “exclusion principle” as a differentiable loss function that we incorporate into the SMPLify-X pose estimation method [49].

Second, physical interaction requires *contact* in 3D space to apply forces. To exploit this, we use the simple heuristic that certain areas of the body surface are the most likely to contact the scene, and that, when such body surfaces are close to scene surfaces, and have the same orientation, they are likely to be in contact. Although these ideas have been explored to some extent by the 3D hand-object estimation community [38, 47, 51, 56, 67, 68] they have received less attention in work on 3D body pose. We formulate a term that implements this contact heuristic and find that it improves pose estimation.

Our method extends SMPLify-X [49], which fits a 3D body model “top down” to “bottom up” features (e.g. 2D joint detections). We choose this optimization-based framework over a direct regression method (deep neural network) because it is more straightforward to incorporate our physically-motivated constraints. The method enforces *Proximal Relationships with Object eXclusion* and is called *PROX*. Figure 1 shows a representative example where the human body pose is estimated with and without our environmental terms. From the viewpoint of the camera, both solutions look good and match the 2D image but, when placed in a scan of the 3D scene, the results without environmental constraints can be grossly inaccurate. Adding our constraints to the optimization reduces inter-penetration and encourages appropriate contact.

One may ask why such constraints are not typically used? One key reason is that to estimate and reason about contact and inter-penetration, one needs both a model of the 3D scene and a *realistic* model of the human body. The former is easy to obtain today with many scanning technologies but, if the body model is not accurate, it does not make sense to reason about contact and inter-penetration. Consequently we use the SMPL-X body model [49], which is realistic enough to serve as a “proxy” for the real human in the 3D scene. In particular, the feet, hands, and body of the model have realistic shape and degrees of freedom.

Here we assume that a rough 3D model of the scene is available. It is fair to ask whether it is realistic to perform

monocular human pose estimation but assume a 3D scene? We argue that it is for two key reasons. First, scanning a scene today is quite easy with commodity sensors. If the scene is static, then it can be scanned once, enabling accurate body pose estimation from a single RGB camera; this may be useful for surveillance, industrial, or special-effects applications. Second, methods to estimate 3D scene structure from a single image are advancing extremely quickly. There are now good methods to infer 3D depth maps from a single image [15], as well as methods that do more semantic analysis and estimate 3D CAD models of the objects in the scene [45]. Our work is complementary to this direction and we believe that monocular 3D scene estimation and monocular 3D human pose estimation should happen together. The work here provides a clear example of why this is valuable.

To evaluate PROX, we use three datasets: two *qualitative datasets* and a *quantitative dataset*. The qualitative datasets contain: 3D scene scans, monocular RGB-D videos and pseudo ground-truth human bodies. The pseudo ground-truth is extracted from RGB-D by extending SMPLify-X to use both RGB and depth data to fit SMPL-X.

In order to get true ground-truth for the quantitative dataset, we set up a living room in a marker-based motion capture environment, scan the scene, and collect RGB-D images in addition to the MoCap data. We fit the SMPL-X model to the MoCap marker data using MoSh++ [41] and this provides ground-truth 3D body shape and pose. This allows us to quantitatively evaluate our method.

Our datasets and code are available for research at <https://prox.is.tue.mpg.de>.

2. Related Work

Human pose estimation and 3D scene reconstruction have been thoroughly studied for decades, albeit mostly disjointly. Traditionally, human pose estimation methods [43] estimate bodies in isolation ignoring the surrounding world, while 3D reconstruction methods focus on acquiring the dense 3D shape of the scene only [76] or performing semantic analysis [7, 13, 54], assuming no humans are present. In this work we focus on exploiting and capturing human-world interactions.

The community has made significant progress on estimating human body pose and shape from images [18, 43, 53, 60]. Recent methods based on deep learning, extend 3D human pose estimation to complex scenes [32, 42, 48, 50] but the 3D accuracy is limited. To estimate human-scene interaction, however, more realistic body models are needed that include fully articulated hands such as in [31, 49].

Joint Human & World Models: Several works focus on improving 2D object detection, 2D pose, and action recognition by observing RGB imagery of people interacting with objects [5, 23, 35, 52, 72]. [14, 17, 24] use similar observations to reason about the 3D scene, *i.e.* rough 3D

reconstruction and affordances, however scene cues are not used as feedback to improve human pose. Another direction models human-scene interactions by hallucinating synthetic people either in real RGB images of scenes [29] for general scene labeling, or in synthetic 3D scenes to learn affordances [21, 33] or 3D object layout in the scene [30], or in real 3D scans of scenes [16] for scene synthesis. Here we exploit this 3D structure to better capture poses of humans in it. In the following we focus on the more recent works of [21, 33, 44, 61, 62] that follow this idea.

Several of these observe real human-world interactions in RGB-D videos [44, 61, 62]. [62] learns a joint probabilistic model over 3D human poses and 3D object arrangements, encoded as a set of human-centric prototypical interaction graphs (PiGraphs). The learned PiGraphs can then be used to generate plausible static 3D human-object interaction configurations from high level textual description. [44] builds on the PiGraphs dataset to define a database of “scenelets”, that are then fitted in RGB videos to reconstruct plausible dynamic interaction configurations over space-time. Finally, [61] employs similar observations to predict action maps in a 3D scene. However, these works capture noisy human poses and do not make use of scene constraints to improve them. They also represent human pose as a 3D skeleton, not a full 3D body.

Other works like [21, 33] use synthetic 3D scenes and place virtual humans in them to reason about affordances. [21] do this by using defined key poses of the body and evaluating human-scene distances and mesh intersections. These methods do not actually capture people in scenes. Our approach could provide rich training data for methods like these to reason about affordances.

Human & World Constraints: Other works employ human-world interactions more explicitly to establish physical constraints, *i.e.* either contact or collision constraints. Yamamoto and Yagishita [71] were the first to use scene constraints in 3D human tracking. They observed that the scene can constrain the position, velocity and acceleration of an articulated 3D body model. Later work adds object contact constraints to the body to effectively reduce the degrees of freedom of the body and make pose estimation easier [34, 58]. Brubaker et al. [11] focus on walking and perform 3D person tracking by using a kinematic model of the torso and the lower body as a prior over human motion and conditioning its dynamics on the 2D Anthropomorphic Walker [36]. Hasler et al. [25] reconstruct a rough 3D scene from multiple unsynchronized moving cameras and employ scene constraints for pose estimation. The above methods all had the right idea but required significant manual intervention or were applied in very restricted scenarios.

Most prior methods that have used world constraints focus on interaction with a ground plane [69] or simply constrain the body to move along the ground plane [74]. Most

interesting among these is the work of Vondrak et al. [69] where they exploit a game physics engine to infer human pose using gravity, motor forces, and interactions with the ground. This is a very complicated optimization and it has not been extended beyond ground contact.

Gupta et al. [22] exploit contextual scene information in human pose estimation using a GPLVM learning framework. For an action like sitting, they take motion capture data of people sitting on objects of different heights. Then, conditioned on the object height, they estimate the pose in the image, exploiting the learned pose model.

Shape2Pose [33] learns a model to generate plausible 3D human poses that interact with a given 3D object. First contact points are inferred on the object surface and then the most likely pose that encourages close proximity of relevant body parts to contact points is estimated. However, the approach only uses synthetic data. [73] establish contact constraints between the feet and an estimated ground plane. For this they first estimate human poses in multi-person RGB videos independently and fit a ground plane around the ankle joint positions. They then refine poses in a global optimization scheme over all frames incorporating contact and temporal constraints, as well as collision constraints, using a collision model comprised of shape primitives similar to [10, 47]. More recently, [39] introduced a method to estimate contact positions, forces and torques actuated by the human limbs during human-object interaction.

The 3D hand-object community has also explored similar physical constraints, such as [37, 47, 51, 56, 67, 68] to name a few. Most of these methods employ a collision model to avoid hand-object inter-penetrations with varying degrees of accuracy; using underlying shape primitives [38, 47] or decomposition in convex parts of more complicated objects [38], or using the original mesh to detect colliding triangles along with 3D distance fields [68]. Triangle intersection tests have also been used to estimate contact points and forces [56]. Most other work uses simple proximity checks [64, 67, 68] and employs an attraction term at contact points. Recently, [27] propose an end-to-end model that exploits a contact loss and inter-penetration penalty to reconstruct hands manipulating objects in RGB images.

In summary, past work focuses either on specific body parts (hands or feet) or interaction with a limited set of objects (ground or hand-held objects). Here, for the first time, we address the full articulated body interacting with diverse, complex and full 3D scenes. Moreover, we show how using the 3D scene improves monocular 3D body pose estimation.

3. Technical Approach

3.1. 3D Scene Representation

To study how people interact with a scene, we first need to acquire knowledge about it, *i.e.* to perform scene recon-

struction. Since physical interaction takes place through surfaces, we chose to represent the scene as a 3D mesh $M_s = (V_s, F_s)$, with $|V_s| = N_s$ vertices $V_s \in \mathbb{R}^{(N_s \times 3)}$ and triangular faces F_s . We assume a static 3D scene and reconstruct M_s with a standard commercial solution; the Structure Sensor [4] camera and the Skanect [3] software. We chose the scene frame to represent the world coordinate frame; both the camera and the human model are expressed w.r.t. this as explained in Sections 3.2 and 3.3, respectively.

3.2. Camera Representation

We use a Kinect-One camera [1] to acquire RGB and depth images of a person moving and interacting with the scene. We use a publicly available tool [2] to estimate the intrinsic camera parameters K_c and to capture synchronized RGB-D images; for each time frame t we capture a 512×424 depth image Z^t and 1920×1080 RGB image I^t at 30 FPS. We then transform the RGB-D data into point cloud P^t .

To perform human MoCap w.r.t. to the scene, we first need to register the RGB-D camera to the 3D scene. We assume a static camera and estimate the extrinsic camera parameters, *i.e.* the camera-to-world rigid transformation $T_c = (R_c, t_c)$, where $R_c \in SO(3)$ is a rotation matrix and $t_c \in \mathbb{R}^3$ is a translation vector. For each sequence a human annotator annotates 3 correspondences between the 3D scene M_s and the point cloud P^t to get an initial estimate of T_c , which is then refined using ICP [9, 75]. The camera extrinsic parameters (R_c, t_c) are fixed during each recording (Section 3.4),

The human body b is estimated in the camera frame and needs to be registered to the scene by applying T_c to it too. For simplicity of notation, we use the same symbols for the camera c and body b after transformation to the world coordinate frame.

3.3. Human Body Model

We represent the human body using SMPL-X [49]. SMPL-X is a generative model that captures how the human body shape varies across a human population, learned from a corpus of registered 3D body, face and hand scans of people of different sizes, genders and nationalities in various poses. It goes beyond similar models [6, 26, 40, 57] by holistically modeling the body with facial expressions and finger articulation, which is important for interactions.

SMPL-X is a differentiable function $M_b(\beta, \theta, \psi, \gamma)$ parameterized by shape β , pose θ , facial expressions ψ and translation γ . Its output is a 3D mesh $M_b = (V_b, F_b)$ for the human body, with $N_b = 10475$ vertices $V_b \in \mathbb{R}^{(N_b \times 3)}$ and triangular faces F_b . The shape parameters $\beta \in \mathbb{R}^{10}$ are coefficients in a lower-dimensional shape space learned from approximately 4000 registered CAESAR [55] scans. The pose of the body is defined by linear blend skinning with an underlying rigged skeleton, whose 3D joints $J(\beta)$

are regressed from the mesh vertices. The skeleton has 55 joints in total; 22 for the main body (including a global pelvis joint), 3 for the neck and the two eyes, and 15 joints per hand for finger articulation. The pose parameters $\theta = (\theta_b, \theta_f, \theta_h)$ are comprised of $\theta_b \in \mathbb{R}^{66}$ and $\theta_f \in \mathbb{R}^9$ parameters in axis-angle representation for the main body and face joints respectively, with 3 degrees of freedom (DOF) per joint, as well as $\theta_h \in \mathbb{R}^{12}$ pose parameters in a lower-dimensional pose space for finger articulation of both hands, captured by approximately 1500 registered hand scans [57]. The pose parameters θ and translation vector $\gamma \in \mathbb{R}^3$ define a function that transforms the joints along the kinematic tree $R_{\theta\gamma}$. Following the notation of [10] we denote posed joints with $R_{\theta\gamma}(J(\beta)_i)$ for each joint i .

3.4. Human MoCap from Monocular Images

To fit SMPL-X to single RGB images we employ SMPLify-X [49] and extend it to include human-world interaction constraints to encourage contact and discourage inter-penetrations. We name our method *PROX* for *Proximal Relationships with Object eXclusion*. We extend SMPLify-X to SMPLify-D, which uses both RGB and an additional depth input for more accurate registration of human poses to the 3D scene. We also extend PROX to use RGB-D input instead of RGB only; we call this configuration PROX-D.

Inspired by [49], we formulate fitting SMPL-X to monocular images as an optimization problem, where we seek to minimize the objective function

$$E(\beta, \theta, \psi, \gamma, M_s) = E_J + \lambda_D E_D + \lambda_{\theta_b} E_{\theta_b} + \lambda_{\theta_f} E_{\theta_f} + \lambda_{\theta_h} E_{\theta_h} + \lambda_{\alpha} E_{\alpha} + \lambda_{\beta} E_{\beta} + \lambda_{\mathcal{E}} E_{\mathcal{E}} + \lambda_{\mathcal{P}} E_{\mathcal{P}} + \lambda_{\mathcal{C}} E_{\mathcal{C}} \quad (1)$$

where θ_b, θ_f and θ_h are the pose vectors for the body, face (neck, jaw) and the two hands respectively, $\theta = \{\theta_b, \theta_f, \theta_h\}$ is the full set of optimizable pose parameters, γ denotes the body translation, β the body shape and ψ the facial expressions, as described in Section 3.3. $E_J(\beta, \theta, \gamma, K, J_{est})$ and $E_D(\beta, \theta, \gamma, K, Z)$ are data terms that are described below; E_J is the RGB data term used in all configurations, while E_D is the optional depth data term which is used whenever depth data is available. The terms $E_{\theta_b}(\theta_b)$, $E_{\theta_f}(\theta_f)$, $E_{\mathcal{E}}(\mathcal{E})$ and $E_{\beta}(\beta)$ are L_2 priors for the hand pose, facial pose, facial expressions and body shape, penalizing deviation from the neutral state. Following [10, 49] the term $E_{\alpha}(\theta_b) = \sum_{i \in \{elbows, knees\}} \exp(\theta_i)$ is a prior penalizing extreme bending only for elbows and knees, while $E_{\theta_b}(\theta_b)$ is a VAE-based body pose prior called VPoser introduced in [49]. The term $E_{\mathcal{C}}(\beta, \theta, \gamma, M_s)$ encourages contact between the body and the scene as described in Section 3.5. The term $E_{\mathcal{P}}(\theta, \beta, M_s)$ is a penetration penalty modified from [49] to reason about both self-penetrations and human-scene inter-penetrations, as described in Section 3.6. The



Figure 2: Annotated vertices that come frequently in contact with the world, highlighted with blue color.

terms E_J , E_{θ_b} , E_{θ_h} , E_α , E_β and weights λ_i are as described in [49]. The weights λ_i denote steering weights for each term. They were set empirically in an annealing scheme similar to [49].

For the *RGB data term* E_J we use a re-projection loss to minimize the weighted robust distance between 2D joints $J_{est}(I)$ estimated from the RGB image I and the 2D projection of the corresponding posed 3D joints $R_{\theta_\gamma}(J(\beta)_i)$ of SMPL-X, as defined for each joint i in Section 3.3. Following the notation of [10, 49], the data term is

$$E_J(\beta, \theta, \gamma, K, J_{est}) = \sum_{\text{joint } i} \kappa_i \omega_i \rho_J(\Pi_K(R_{\theta_\gamma}(J(\beta)_i) - J_{est,i})) \quad (2)$$

where Π_K denotes the 3D to 2D projection with intrinsic camera parameters K . For the 2D detections we rely on OpenPose [12, 63, 70], which provides body, face and hands keypoints jointly for each person in an image. To account for noise in the detections, the contribution of each joint in the data term is weighted by the detection confidence score ω_i , while κ_i are per-joint weights for annealed optimization, as described in [49]. Furthermore, ρ_J denotes a robust Geman-McClure error function [19] for down-weighting noisy detections.

The *depth data term* E_D minimizes the discrepancy between the visible body vertices $V_b^v \subset V_b$ and a segmented point cloud P^t that belongs only to the body and not the static scene. For this, we use the body segmentation mask from the Kinect-One SDK. Then, E_D is defined as

$$E_D(\beta, \theta, \gamma, K, Z) = \sum_{p \in P^t} \rho_D(\min_{v \in V_b^v} \|v - p\|) \quad (3)$$

where ρ_D denotes a robust Geman-McClure error function [19] for downweighting vertices V_b^v that are far from P^t .

3.5. Contact Term

Using the RGB term E_J without reasoning about human-world interaction might result in physically implausible poses, as shown in Figure 1; However, when humans

interact with the scene they come in *contact* with it, e.g. feet contact the floor while standing or walking. We therefore introduce the term E_C to encourage contact and *proximity* between body parts and the scene around contact areas.

To that end, we annotate a set of candidate contact vertices $V_C \subset V_b$ across the whole body that come frequently in contact with the world, focusing on the actions of walking, sitting and touching with hands. We annotate 1121 vertices across the whole body, as shown in Figure 2. We also explored choosing all body vertices as contact vertices but found that this choice is suboptimal, for evaluation see Sup. Mat. We define the contact vertices as: 725 vertices for the hands, 62 vertices for the thighs, 113 for the gluteus, 222 for the back, and 194 for the feet. E_C is defined as:

$$E_C(\beta, \theta, \gamma, M_s) = \sum_{v_C \in V_C} \rho_C(\min_{v_s \in V_s} \|v_C - v_s\|) \quad (4)$$

where ρ_C denotes a robust Geman-McClure error function [19] for down-weighting vertices in V_C that are far from the nearest vertices in V_s of the 3D scene M_s .

3.6. Penetration Term

Intuitive physics suggests that two objects can not share the same 3D space. However, human pose estimation methods might result in self-penetrations or bodies penetrating surrounding 3D objects, as shown in Figure 1. We therefore introduce a penetration term that combines $E_{\mathcal{P}_{self}}$ and $E_{\mathcal{P}_{inter}}$ that are defined below:

$$E_{\mathcal{P}}(\theta, \beta, \gamma, M_s) = E_{\mathcal{P}_{self}}(\theta, \beta) + E_{\mathcal{P}_{inter}}(\theta, \beta, \gamma, M_s) \quad (5)$$

For *self-penetrations* we follow the approach of [8, 49, 68], that follows local reasoning. We first detect a list of colliding body triangles \mathcal{P}_{self} using Bounding Volume Hierarchies (BVH) [66] and compute local conic 3D distance fields Ψ . Penetrations are then penalized according to the depth in Ψ . For the exact definition of Ψ and $E_{\mathcal{P}_{self}}(\theta, \beta)$ we refer the reader to [8, 68].

For body-scene *inter-penetrations* local reasoning at colliding triangles is not enough, as the body might be initialized deep inside 3D objects or even outside the 3D scene. To resolve this, we penalize all penetrating vertices using the signed distance field (SDF) of the scene M_s . The distance field is represented with a uniform voxel grid with size $256 \times 256 \times 256$, that spans a padded bounding box of the scene. Each voxel cell c_i stores the distance from its center $p_i \in \mathbb{R}^3$ to the nearest surface point $p_i^s \in \mathbb{R}^3$ of M_s with normal $n_i^s \in \mathbb{R}^3$, while the sign is defined according to the relative orientation of the vector $p_i - p_i^s$ w.r.t. n_i^s as

$$\text{sign}(c_i) = \text{sign}((p_i - p_i^s) \cdot n_i^s); \quad (6)$$

a positive sign means that the body vertex is outside the nearest scene object, while a negative sign means that it is



Figure 3: Reconstructed 3D scans of the 12 indoor scenes of our PROX dataset, as well as an additional scene for our quantitative dataset, shown at the bottom right corner.



Figure 4: Example RGB frames of our PROX dataset showing people moving in natural indoor scenes and interacting with them. We reconstruct in total 12 scenes and capture 20 subjects. Figure 3 shows the 3D reconstructions of our indoor scenes.

inside the nearest scene object and denotes penetration. In practice, during optimization we can find how each body vertex V_{b_i} is positioned relative to the scene by reading the signed distance $d_i \in \mathbb{R}$ of the voxel it falls into. Since the limited grid resolution influences discretization of the 3D distance field, we perform trilinear interpolation using the neighboring voxels similar to [28]. Then we resolve body-scene inter-penetration by minimizing the loss term

$$E_{\mathcal{P}_{inter}} = \sum_{d_i < 0} \|d_i n_i^s\|^2. \quad (7)$$

3.7. Optimization

We optimize Equation 1 similar to [49]. More specifically, we implement our model in PyTorch and use the Limited-memory BFGS optimizer (L-BFGS) [46] with strong Wolfe line search.

4. Datasets

4.1. Qualitative Datasets

The qualitative datasets, PiGraphs and PROX, contain: 3D scene scans and monocular videos of people interacting

with the 3D scenes. They do not include ground-truth bodies, thus we cannot evaluate our method quantitatively on these datasets.

4.1.1 PiGraphs dataset

This dataset was released as part of the work of Sava *et al.* [62]. The dataset has several 3D scene scans and RGB-D videos. It suffers from multiple limitations; the color and depth frames are neither synchronized nor spatially calibrated, making it hard to use both RGB and depth. The human poses are rather noisy and are not well registered into the 3D scenes, which are inaccurately reconstructed. The dataset has a low frame rate of 5 fps, it is limited to only 5 subjects and does not have ground-truth.

4.1.2 PROX dataset

We collected this dataset to overcome the limitations of the PiGraphs dataset. We employ the commercial Structure Sensor [4] RGB-D camera and the accompanying 3D reconstruction solution Skanect [3] and reconstruct 12 indoor scenes, shown in Figure 3. The scenes can be grouped to: 3 bedrooms, 5 living rooms, 2 sitting booths and 2 offices. We then employ a Kinect-One [1] RGB-D camera to capture 20 subjects (4 females and 16 males) interacting with these scenes. Subjects gave written informed consent to make their data available for research purposes. The dataset provides 100K synchronized and spatially calibrated RGB-D frames at 30 fps. Figure 4 shows example RGB frames from our dataset. We leverage the RGB-D videos to get pseudo ground-truth by extending SMPLify-X to SMPLify-D which fits SMPL-X to both RGB and depth data instead of RGB only.

4.2. Quantitative Dataset

Neither our PROX dataset nor PiGraphs [62] have ground-truth for quantitative evaluation. To account for this, we captured a separate *quantitative dataset* with 180 static RGB-D frames in sync with a 54 camera Vicon system. We placed markers on the body and the fingers. We placed everyday furniture and objects inside the Vicon area to mimic a living room, and performed 3D reconstruction of the scene, shown in the bottom right corner of Figure 3 with the Structure Sensor [4] and Skanect [3] similar to above. We then use MoSh++ [41] which is a method that converts MoCap data into realistic 3D human meshes represented by a rigged body model. Example RGB frames are shown in Figure 5 (left), while our mesh pseudo ground-truth is shown with aqua blue color.

Our datasets will be available for research purposes.

Eq. 1 terms				Error			
E_J	E_C	E_P	E_D	PJE	V2V	p.PJE	p.V2V
✓	✗	✗	✗	220.27	218.06	73.24	60.80
✓	✓	✗	✗	208.03	208.57	72.76	60.95
✓	✗	✓	✗	190.07	190.38	73.73	62.38
✓	✓	✓	✗	167.08	166.51	71.97	61.14
✓	✗	✗	✓	72.91	69.89	55.53	48.86
✓	✓	✓	✓	68.48	60.83	52.78	47.11
mm							
✓	✗	✗	✗	232.29	227.49	66.02	53.15
✓	✓	✓	✗	144.60	156.90	65.04	52.60
mm							

Table 1: Ablation study for Equation 1; each row contains the terms indicated by the check-boxes. Units in *mm*. PROX and PROX-D are shown in bold. **Table (a):** Evaluation on our *quantitative dataset* using mesh pseudo ground-truth based on Vicon and MoSh++ [41]. **Table (b):** Evaluation on chosen sequences of our *qualitative dataset* using pseudo ground-truth based on SMPLify-D. **Tables (a, b):** We report the mean per-joint error without/with procrustes alignment noted as “PJE” / “p.PJE”, and the mean vertex-to-vertex error noted as “V2V” / “p.V2V”.

5. Experiments

Quantitative Evaluation: To evaluate the performance of our method, as well as to evaluate the importance of different terms in Equation 1, we perform quantitative evaluation in Table 1. As performance metrics we report the mean per-joint error without and with procrustes alignment noted as “PJE” and “p.PJE” respectively, as well as the mean vertex-to-vertex error noted similarly as “V2V” and “p.V2V”. Each row in the table shows a setup that includes different terms as indicated by the check-boxes. Table 1 includes two sub-tables for different datasets. **Table 1 (a):** We employ our new *quantitative dataset* with mesh pseudo ground-truth based on Vicon and MoSh++ [41], as described in Section 4. The first row with only E_J is an RGB-only baseline similar to SMPLify-X [49], that we adapt to our needs by using a fixed camera and estimating body translation γ , and gives the biggest “PJE” and “V2V” error. In the second row we add only the contact term E_C , while in the third row we add only the penetration term E_P . In both cases the error drops a bit, however the drop is significantly bigger for the fourth row that includes both E_C and E_P ; this corresponds to PROX and achieves 167.08 mm “PJE” and 166.51 mm “V2V” error. This suggests that both E_C and E_P contribute to accuracy and are complementary. To inform the upper bound of performance, in the fifth row we employ an RGB-D baseline with E_J and E_D , which corresponds to SMPLify-D as described in Section 3.4. All terms of Equation 1 are employed in the last row; we call this configuration PROX-D. We observe that using scene constraints boosts the performance even when the depth is available. This gives the best overall performance, but

PROX (fourth row) achieves reasonably good performance with less input data, *i.e.* using RGB only. **Table 1 (b):** We chose 4 random sequences of our new PROX dataset. We generate pseudo ground-truth with SMPLify-D, which uses both RGB and depth. We show a comparison between the RGB-only baseline (first row) and PROX (second row) compared to the pseudo ground-truth of SMPLify-D. The results support the above finding that the scene constraints in PROX contribute significantly to accuracy.

The run time for all configurations is reported in the Sup. Mat.

Qualitative Evaluation: In Figure 5 we show qualitative results too for our *quantitative dataset*. Furthermore, in Figure 6 we show representative qualitative results on the *qualitative datasets*; our PROX dataset and PiGraphs dataset. In both figures, the lack of scene constraints (yellow) results in severe penetrations in the scene. Our method, PROX, includes scene constraints (light gray) and estimates bodies that are significantly more consistent with the 3D scene, *i.e.* with realistic contact and without penetrations. More qualitative results are available in the Sup. Mat.

6. Conclusion

In this work we focus on human-world interactions and capture the motion of humans interacting with a real static 3D scene in RGB images. We use a holistic model, SMPL-X [49], that jointly models the body with face and fingers, which are important for interactions. We show that incorporating interaction-based human-world constraints in an optimization framework (PROX) results in significantly more realistic and accurate MoCap. We also collect a new dataset of 3D scenes with RGB-D sequences involving human interactions and occlusions. We perform extensive quantitative and qualitative evaluations that clearly show the benefits of incorporating scene constraints into 3D human pose estimation. Our code, data and MoCap are available for research purposes.

Limitations and Future work: A limitation of the current formulation is that we do not model scene occlusion. Current 2D part detectors do not indicate when joints are occluded and may provide inaccurate results. By knowing the scene structure we could reason about what is visible and what is not. Another interesting direction would be the unification of the self-penetration and the body-scene interpenetration by employing the implicit formulation of [65] for the whole body. Future work can exploit recent deep networks to estimate the scene directly from monocular RGB images. More interesting directions would be to extend our method to dynamic scenes [59], human-human interaction and to account for scene and body deformation.

Acknowledgments: We thank Dorotea Lleshaj, Markus Höschle, Mason Landry, Andrea Keller and Tsvetelina Alexiadis for their help with the data collection. Jean-




Figure 5: Examples from our *quantitative dataset*, described in Section 5. From left to right: (1) RGB images, (2) rendering of the fitted model and the 3D scene from the camera viewpoint; aqua blue for the mesh pseudo ground-truth, light gray for the results of our method PROX, yellow for results without scene constraints, green for SMPLify-D, (3) top view and (4) side view. More results can be found in Sup. Mat.




Figure 6: Qualitative results of our method on two datasets; on our *qualitative dataset* (top set) and on the PiGraphs dataset [62] (bottom set). From left to right: (1) RGB images, (2) rendering from the camera viewpoint; light gray for the results of our method PROX, yellow for results without scene constraints, and green for SMPLify-D (applicable only for the top set), (3) rendering from a different view, that shows that the camera view is deceiving. More results can be found in Sup. Mat.

Claude Passy for helping with the data collection software. Nima Ghorbani for MoSh++. Benjamin Pellkofer for the IT support. Jonathan Williams for managing the website.

Disclosure: MJB has received research gift funds from Intel, Nvidia, Adobe, Facebook, and Amazon. While MJB

is a part-time employee of Amazon, his research was performed solely at, and funded solely by, MPI. MJB has financial interests in Amazon and Meshcapade GmbH.

References

- [1] Kinect for xbox one. [https://en.wikipedia.org/wiki/Kinect#Kinect_for_Xbox_One_\(2013\)](https://en.wikipedia.org/wiki/Kinect#Kinect_for_Xbox_One_(2013)). 4, 6
- [2] Monocle: Kinect data capture app. <https://github.com/bmabey/monocle>. 4
- [3] Skanect: 3d scanning. <https://skanect.occipital.com>. 4, 6
- [4] Structure sensor: 3d scanning, augmented reality and more. <https://structure.io/structure-sensor>. 4, 6
- [5] Eren Erdal Aksoy, Alexey Abramov, Florentin Wörgötter, and Babette Dellen. Categorizing object-action relations from semantic scene graphs. In *2010 IEEE International Conference on Robotics and Automation (ICRA)*, pages 398–405, 2010. 2
- [6] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: Shape Completion and Animation of PEople. *ACM Transactions on Graphics (TOG), (Proc. SIGGRAPH)*, 24(3):408–416, 2005. 4
- [7] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1534–1543, 2016. 2
- [8] Luca Ballan, Aparna Taneja, Juergen Gall, Luc Van Gool, and Marc Pollefeys. Motion capture of hands in action using discriminative salient points. In *The European Conference on Computer Vision (ECCV)*, pages 640–653, 2012. 5
- [9] Paul J. Besl and Neil D. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 14(2):239–256, 1992. 4
- [10] 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 *The European Conference on Computer Vision (ECCV)*, 2016. 3, 4, 5
- [11] Marcus A. Brubaker, David J. Fleet, and Aaron Hertzmann. Physics-based person tracking using the anthropomorphic walker. *International Journal of Computer Vision*, 87(1):140, Aug 2009. 3
- [12] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2D pose estimation using part affinity fields. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5
- [13] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [14] Vincent Delaitre, David F Fouhey, Ivan Laptev, Josef Sivic, Abhinav Gupta, and Alexei A Efros. Scene semantics from long-term observation of people. In *The European Conference on Computer Vision (ECCV)*, pages 284–298, 2012. 2
- [15] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in Neural Information Processing Systems*, pages 2366–2374, 2014. 2
- [16] Matthew Fisher, Manolis Savva, Yangyan Li, Pat Hanrahan, and Matthias Nießner. Activity-centric scene synthesis for functional 3d scene modeling. *ACM Transactions on Graphics (TOG)*, 34(6):179, 2015. 3
- [17] David F Fouhey, Vincent Delaitre, Abhinav Gupta, Alexei A Efros, Ivan Laptev, and Josef Sivic. People watching: Human actions as a cue for single view geometry. *International Journal of Computer Vision (IJCV)*, 110(3):259–274, 2014. 2
- [18] Darius M. Gavrilă. The visual analysis of human movement: A survey. *Computer Vision and Image Understanding (CVIU)*, 73(1):82 – 98, 1999. 2
- [19] Stuart Geman and Donald E. McClure. Statistical methods for tomographic image reconstruction. In *Proceedings of the 46th Session of the International Statistical Institute, Bulletin of the ISI*, volume 52, 1987. 5
- [20] James J Gibson. *The perception of the visual world*. Houghton Mifflin, 1950. 1
- [21] Helmut Grabner, Juergen Gall, and Luc Van Gool. What makes a chair a chair? In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1529–1536, 2011. 3
- [22] Abhinav Gupta, Trista Chen, Francine Chen, Don Kimber, and Larry S Davis. Context and observation driven latent variable model for human pose estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1 – 8, 2008. 3
- [23] Abhinav Gupta, Aniruddha Kembhavi, and Larry S Davis. Observing human-object interactions: Using spatial and functional compatibility for recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 31(10):1775–1789, 2009. 2
- [24] Abhinav Gupta, Scott Satkin, Alexei A Efros, and Martial Hebert. From 3d scene geometry to human workspace. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1961–1968, 2011. 2
- [25] Nils Hasler, Bodo Rosenhahn, Thorsten Thormahlen, Michael Wand, Jürgen Gall, and Hans-Peter Seidel. Markerless motion capture with unsynchronized moving cameras. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 224–231, June 2009. 3
- [26] Nils Hasler, Carsten Stoll, Martin Sunkel, Bodo Rosenhahn, and Hans-Peter Seidel. A statistical model of human pose and body shape. *Computer Graphics Forum*, 28(2):337–346, 2009. 4
- [27] Yana Hasson, Gül Varol, Dimitrios Tzionas, Igor Kaleyvatykh, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning joint reconstruction of hands and manipulated objects. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [28] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*. 2015. 6

- [29] Yun Jiang, Hema Koppula, and Ashutosh Saxena. Hallucinated humans as the hidden context for labeling 3d scenes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2993–3000, 2013. 3
- [30] Yun Jiang, Marcus Lim, and Ashutosh Saxena. Learning object arrangements in 3d scenes using human context. In *Proceedings of the 29th International Conference on Machine Learning*, pages 907–914, 2012. 3
- [31] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3D deformation model for tracking faces, hands, and bodies. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [32] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [33] Vladimir G Kim, Siddhartha Chaudhuri, Leonidas Guibas, and Thomas Funkhouser. Shape2pose: Human-centric shape analysis. *ACM Transactions on Graphics (TOG)*, 33(4):120, 2014. 3
- [34] Hedvig Kjellström, Danica Kragić, and Michael J Black. Tracking people interacting with objects. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 747–754, 2010. 3
- [35] Hema Swetha Koppula, Rudhir Gupta, and Ashutosh Saxena. Learning human activities and object affordances from rgb-d videos. *The International Journal of Robotics Research*, 32(8):951–970, 2013. 2
- [36] Arthur D Kuo. A simple model of bipedal walking predicts the preferred speed–step length relationship. *Journal of biomechanical engineering*, 123(3):264–269, 2001. 3
- [37] Nikolaos Kyriazis and Antonis Argyros. Physically plausible 3D scene tracking: The single actor hypothesis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9–16, 2013. 3
- [38] Nikolaos Kyriazis and Antonis Argyros. Scalable 3D tracking of multiple interacting objects. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3430–3437, 2014. 2, 3
- [39] Zongmian Li, Jiri Sedlar, Justin Carpentier, Ivan Laptev, Nicolas Mansard, and Josef Sivic. Estimating 3d motion and forces of person-object interactions from monocular video. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [40] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics (TOG), (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. 4
- [41] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2019. 2, 6, 7
- [42] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Transactions on Graphics (TOG)*, 36(4):44:1–44:14, July 2017. 2
- [43] Thomas B. Moeslund, Adrian Hilton, and Volker Krüger. A survey of advances in vision-based human motion capture and analysis. *Computer Vision and Image Understanding (CVIU)*, 104(2):90–126, 2006. 2
- [44] Aron Monszpart, Paul Guerrero, Duygu Ceylan, Ersin Yumer, and Niloy J Mitra. imapper: interaction-guided scene mapping from monocular videos. *ACM Transactions on Graphics (TOG)*, 38(4):92, 2019. 3
- [45] Muzammal Naseer, Salman Khan, and Fatih Porikli. Indoor scene understanding in 2.5/3d for autonomous agents: A survey. *IEEE Access*, 7:1859–1887, 2019. 2
- [46] Jorge Nocedal and Stephen J Wright. *Nonlinear Equations*. Springer, 2006. 6
- [47] Iason Oikonomidis, Nikolaos Kyriazis, and Antonis A. Argyros. Full dof tracking of a hand interacting with an object by modeling occlusions and physical constraints. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 2088–2095, 2011. 2, 3
- [48] 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. In *3DV*, Sept. 2018. 2
- [49] 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 *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 4, 5, 6, 7
- [50] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3D human pose and shape from a single color image. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [51] Tu-Hoa Pham, Nikolaos Kyriazis, Antonis A Argyros, and Abderrahmane Kheddar. Hand-object contact force estimation from markerless visual tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 40(12):2883–2896, Dec 2018. 2, 3
- [52] Hamed Pirsiavash and Deva Ramanan. Detecting activities of daily living in first-person camera views. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2847–2854, 2012. 2
- [53] Ronald Poppe. Vision-based human motion analysis: An overview. *Computer Vision and Image Understanding (CVIU)*, 108(1-2):4–18, 2007. 2
- [54] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 652–660, 2017. 2
- [55] Kathleen M. Robinette, Sherri Blackwell, Hein Daanen, Mark Boehmer, Scott Fleming, Tina Brill, David Hoeflerlin, and Dennis Burnsides. Civilian American and European Surface Anthropometry Resource (CAESAR) final report. Technical Report AFRL-HE-WP-TR-2002-0169, US Air Force Research Laboratory, 2002. 4

- [56] Grégory Rogez, James S. Supančič III, and Deva Ramanan. Understanding everyday hands in action from rgb-d images. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 3889–3897, 2015. 2, 3
- [57] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics (TOG)*, 36(6):245, 2017. 4
- [58] Bodo Rosenhahn, Christian Schmaltz, Thomas Brox, Joachim Weickert, Daniel Cremers, and Hans-Peter Seidel. Markerless motion capture of man-machine interaction. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, June 2008. 3
- [59] Martin Rünz and Lourdes Agapito. Co-fusion: Real-time segmentation, tracking and fusion of multiple objects. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4471–4478, 2017. 7
- [60] Nikolaos Sarafianos, Bogdan Boteanu, Bogdan Ionescu, and Ioannis A. Kakadiaris. 3d human pose estimation: A review of the literature and analysis of covariates. *Computer Vision and Image Understanding (CVIU)*, 152:1–20, 2016. 2
- [61] Manolis Savva, Angel X Chang, Pat Hanrahan, Matthew Fisher, and Matthias Nießner. Scenegrok: Inferring action maps in 3d environments. *ACM Transactions on graphics (TOG)*, 33(6):212, 2014. 3
- [62] Manolis Savva, Angel X Chang, Pat Hanrahan, Matthew Fisher, and Matthias Nießner. Pigraphs: learning interaction snapshots from observations. *ACM Transactions on Graphics (TOG)*, 35(4):139, 2016. 3, 6, 8
- [63] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multi-view bootstrapping. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5
- [64] Srinath Sridhar, Franziska Mueller, Michael Zollhöfer, Dan Casas, Antti Oulasvirta, and Christian Theobalt. Real-time joint tracking of a hand manipulating an object from rgb-d input. In *The European Conference on Computer Vision (ECCV)*, pages 294–310, 2016. 3
- [65] Jonathan Taylor, Vladimir Tankovich, Danhang Tang, Cem Keskin, David Kim, Philip Davidson, Adarsh Kowdle, and Shahram Izadi. Articulated distance fields for ultra-fast tracking of hands interacting. *ACM Transactions on Graphics (TOG)*, 36(6):244:1–244:12, Nov. 2017. 7
- [66] Matthias Teschner, Stefan Kimmerle, Bruno Heidelberger, Gabriel Zachmann, Laks Raghupathi, Arnulph Fuhrmann, Marie-Paule Cani, François Faure, Nadia Magnenat-Thalmann, Wolfgang Strasser, and Pascal Volino. Collision detection for deformable objects. In *Eurographics*, pages 119–139, 2004. 5
- [67] Aggeliki Tsoli and Antonis A. Argyros. Joint 3d tracking of a deformable object in interaction with a hand. In *The European Conference on Computer Vision (ECCV)*, 2018. 2, 3
- [68] Dimitrios Tzionas, Luca Ballan, Abhilash Srikantha, Pablo Aponte, Marc Pollefeys, and Juergen Gall. Capturing hands in action using discriminative salient points and physics simulation. *International Journal of Computer Vision (IJCV)*, 118(2):172–193, 2016. 2, 3, 5
- [69] Marek Vondrak, Leonid Sigal, and Odest Chadwicke Jenkins. Dynamical simulation priors for human motion tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 35(1):52–65, Jan 2013. 3
- [70] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [71] Masanobu Yamamoto and Katsutoshi Yagishita. Scene constraints-aided tracking of human body. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 151–156 vol.1, June 2000. 3
- [72] Bangpeng Yao and Li Fei-Fei. Modeling mutual context of object and human pose in human-object interaction activities. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17–24, 2010. 2
- [73] Andrei Zanfir, Elisabeta Marinoiu, and Cristian Sminchisescu. Monocular 3d pose and shape estimation of multiple people in natural scenes—the importance of multiple scene constraints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2148–2157, 2018. 3
- [74] Tao Zhao and Ram Nevatia. Tracking multiple humans in complex situations. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 26(9):1208–1221, Sep. 2004. 3
- [75] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3D: A modern library for 3D data processing. *arXiv:1801.09847*, 2018. 4
- [76] Michael Zollhöfer, Patrick Stotko, Andreas Görlitz, Christian Theobalt, Matthias Nießner, Reinhard Klein, and Andreas Kolb. State of the art on 3d reconstruction with rgb-d cameras. *Computer Graphics Forum*, 37(2):625–652, 2018. 2

Resolving 3D Human Pose Ambiguities with 3D Scene Constraints

Supplementary Material

Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas and Michael J. Black
Max Planck Institute for Intelligent Systems

{mhassan, vchoutas, dtzionas, black}@tuebingen.mpg.de

Our method enforces *Proximal Relationships with Object eXclusion* and is called *PROX*. The figures below show representative examples where the human body pose is estimated with (gray color) and without (yellow color) our environmental terms. From the viewpoint of the camera, both solutions look good and match the 2D image features but, when placed in a scan of the 3D scene, the results without environment constraints can be grossly inaccurate. Adding our constraints to the optimization reduces inter-penetration and encourages appropriate contact.

Why such constraints are not typically used? One key reason is that to estimate and reason about contact and inter-penetration, one needs *both* a model of the 3D *scene* and a realistic model of the *human body*. The former is easy to obtain today with many scanning technologies but, if the body model is not accurate, it does not make sense to reason about contact and inter-penetration. Consequently we use the SMPL-X body model [3], which is realistic enough to serve as a “proxy” for the real human in the 3D scene. In particular, the feet, hands, and body of the model have realistic shape and degrees of freedom.

Is it realistic to assume a 3D scene for refining pose? Here we assume that a rough 3D model of the scene is available; one could argue that this is a hard assumption. Reconstructing a 3D scene from a single RGB image is a hot research topic, but the problem is ill-posed and currently unsolved. Here we want to show in the first place that knowledge about the scene helps pose estimation. Our results support this hypothesis, and scanning a scene today is quite easy. Our next step is to relax this assumption, and move to the more difficult problem of exploiting recent deep networks to estimate the scene directly from monocular RGB images. There are now good methods to infer depth maps from a single image [1] as well as methods that do more semantic analysis and estimate 3D CAD models of the objects in the scene [2]. Our work is complementary to this direction and we believe that monocular 3D scene estimation and monocular 3D human pose estimation should happen together. The work here provides a clear example of why this is valuable.

Qualitative Results - Our Dataset

Figures A.1-A.3 show additional qualitative results for our method (light gray) on our PROX dataset and compare it to the RGB-only baseline (yellow). For each example we show from left to right: (1) RGB image, (2) renderings from different viewpoints.

Qualitative Results - PiGraphs

Figure A.4 shows additional qualitative results for our method (light gray) on the *PiGraphs* dataset [4] and compare it to the RGB-only baseline (yellow). Please note that [4] estimate just a 3D skeleton of only the major body joints. In contrast, we estimate a full 3D mesh, and include facial expressions and finger articulation. The mesh representation of our realistic human model helps to better reason about proximity to the world, contact and penetrations. For each example we show from left to right: (1) RGB image, (2) renderings from different viewpoints.

Computational Complexity

Table A.1 reports the average runtime for all our configurations (PROX in bold) for 10 randomly sampled frames. Compared to using RGB alone; PROX improved “V2V” by 24% with a runtime increase of 41%.

E_J	E_P	E_C	E_D	Run time	%
✓	✗	✗	✗	33.75	
✓	✓	✗	✗	46.91	
✓	✗	✓	✗	42.68	

E_J	E_P	E_C	E_D	Run time	%
✓	✓	✓	✗	47.64	
✓	✗	✗	✓	54.28	
✓	✓	✓	✓	73.08	

Table A.1: Runtime for all configurations of our approach.

Choice of Contact Vertices

We choose the body vertices that often come in contact with the 3D world. This choice is not exclusive. Table A.2 evaluates different sets of candidate contact vertices, namely our annotations and all vertices. Performance deteriorates in the latter case, while runtime increases by ~ 7 seconds. This suggests the importance of affordances and

semantics; future work can learn the likely contact vertices for different object classes in a data-driven fashion. To this end, the community first needs training data similar to the data generated by our work.

Contact vertices	PJE	V2V	p.PJE	p.V2V	
Selected of Fig. 2	208.03	208.57	72.76	60.95	mm
All selected	217.82	216.62	72.35	60.16	

Table A.2: Different sets of candidate contact vertices.

Failure Cases

Figures A.5-A.6 show failure cases of our method (light gray) on our PROX dataset. For each example we show from left to right: (1) RGB image, (2) OpenPose result overlaid on the RGB image, (3) result of our method. Figure A.5-top shows that our method still results in some penetration. Our assumption of a static scene is not always true; in this case the bed is deformable and its shape changes during interaction. In future work we plan to model deformations of the human body and the world. Figure A.5-bottom shows a failure of our inter-penetration term. In cases where initialization of body translation is not accurate enough, the optimizer might end up in a local minimum that is not always in agreement with the real pose in 3D space. Figure A.6 shows typical failure cases of OpenPose. In Figure A.6-top the left leg is not detected correctly, while in Figure A.6-middle and Figure A.6-bottom several body joints are flipped by OpenPose.

References

- [1] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in Neural Information Processing Systems*, pages 2366–2374, 2014. 1
- [2] Muzammal Naseer, Salman Khan, and Fatih Porikli. Indoor scene understanding in 2.5/3d for autonomous agents: A survey. *IEEE Access*, 7:1859–1887, 2019. 1
- [3] 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 *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1
- [4] Manolis Savva, Angel X Chang, Pat Hanrahan, Matthew Fisher, and Matthias Nießner. Pigraphs: learning interaction snapshots from observations. *ACM Transactions on Graphics (TOG)*, 35(4):139, 2016. 1, 6



Figure A.1: Qualitative results on our PROX dataset. The human body pose is estimated *with* (light gray) and *without* (yellow) our environmental terms. We show from left to right: (1) RGB images, (2) renderings from different viewpoints.



Figure A.2: Qualitative results on our PROX dataset. The human body pose is estimated *with* (light gray) and *without* (yellow) our environmental terms. We show from left to right: (1) RGB images, (2) renderings from different viewpoints.




Figure A.3: Qualitative results on our PROX dataset. The human body pose is estimated *with* (light gray) and *without* (yellow) our environmental terms. We show from left to right: (1) RGB images, (2) renderings from different viewpoints.



Figure A.4: Qualitative results on the PiGraphs [4] dataset. The human body pose is estimated *with* (gray color) and *without* (yellow color) our environmental terms. Please note that [4] estimate just a 3D skeleton of only the major body joints. We show from left to right: (1) RGB images, (2) renderings from different viewpoints.




Figure A.5: Representative failure cases on our PROX dataset. We show from left to right: (1) RGB image, (2) OpenPose result overlaid on the RGB image, (3) result of our method.




Figure A.6: Representative failure cases on our PROX dataset. We show from left to right: (1) RGB image, (2) OpenPose result overlaid on the RGB image, (3) result of our method.