

Detecting Human-Object Contact in Images

Yixin Chen^{1†} Sai Kumar Dwivedi² Michael J. Black² Dimitrios Tzionas³

¹Beijing Institute of General Artificial Intelligence, China

[†]Work done while interning at MPI-IS²

²Max Planck Institute for Intelligent Systems, Tübingen, Germany

³University of Amsterdam, the Netherlands

Abstract

Humans constantly contact objects to move and perform tasks. Thus, detecting human-object contact is important for building human-centered artificial intelligence. However, there exists no robust method to detect contact between the body and the scene from an image, and there exists no dataset to learn such a detector. We fill this gap with HOT (“Human-Object conTact”), a new dataset of human-object contacts in images. To build HOT, we use two data sources: (1) We use the PROX dataset of 3D human meshes moving in 3D scenes, and automatically annotate 2D image areas for contact via 3D mesh proximity and projection. (2) We use the V-COCO, HAKE and Watch-n-Patch datasets, and ask trained annotators to draw polygons around the 2D image areas where contact takes place. We also annotate the involved body part of the human body. We use our HOT dataset to train a new contact detector, which takes a single color image as input, and outputs 2D contact heatmaps as well as the body-part labels that are in contact. This is a new and challenging task, that extends current foot-ground or hand-object contact detectors to the full generality of the whole body. The detector uses a part-attention branch to guide contact estimation through the context of the surrounding body parts and scene. We evaluate our detector extensively, and quantitative results show that our model outperforms baselines, and that all components contribute to better performance. Results on images from an online repository show reasonable detections and generalizability. Our HOT data and model are available for research at <https://hot.is.tue.mpg.de>.

1. Introduction

Contact is an important part of people’s everyday lives. We constantly contact objects to move and perform tasks. We walk by contacting the ground with our feet, we sit by contacting chairs with our buttocks, hips and back, we grasp and manipulate tools by contacting them with our hands. Therefore, estimating contact between humans and objects is useful for human-centered AI, especially for applications



Figure 1. Our contact detector, trained on HOT (“Human-Object conTact”) dataset, estimates contact between humans and scenes from an image taken in the wild. Contact is important for interacting humans, yet, standard in-the-wild datasets unfortunately lack such information. Our contact dataset and detector are a step towards providing this in the wild. Images are from pexels.com.

such as AR/VR [1, 15, 26, 30], activity recognition [22, 39, 44], affordance detection [13, 25, 34, 67], fine-grained human-object interaction detection [27, 37, 51, 57], imitation learning [38, 48, 63], populating scenes with avatars [18, 62, 64], and sanitization of spaces and objects.

In contrast to off-the-shelf detectors for segmenting humans in images, or estimating their 2D joints or 3D shape and pose, there exists no general detector of contact. Some work exists for detecting part-specific contact, e.g., hand-object [35, 45] or foot-ground [41, 49] contact, while other work estimates contact only in constrained environments [20, 46] with limited generalization. What we need, instead, is a contact detector for the *entire body* that estimates detailed, body-part-related, contact *maps* in arbitrary images. To train this, we need data, but no suitable dataset exists at the moment. We address these limitations with a novel dataset and model for detecting contact between whole-body humans and objects in color images taken in the wild.

Annotating contact is challenging, as contact areas are ipso facto occluded. Think of a person standing on the floor; the sole of the shoe, and the floor area it contacts, can not be observed. A naive approach is to instrument a human with contact sensors, however, this is intrusive, cumbersome to set up and does not scale. Instead, we use

two alternative data sources, with different but complementary properties: (1) We use the PROX [17] dataset, which has pseudo ground-truth 3D human meshes for real humans moving in 3D scanned scenes. We *automatically* annotate contact areas, by computing the proximity between the 3D meshes. (2) We use the V-COCO [16], HAKE [27], and Watch-n-Patch [54] datasets, which contain images taken in the wild. We then hire professional annotators, and train them to annotate contact areas as 2D polygons in images. Although *manual* annotation is only approximate, 2D annotations are important because they allow scaling to large, varied, and natural datasets. This improves generalization. Note that in both cases we also annotate the body part that is involved in contact, corresponding to the body parts of the SMPL(-X) [32, 36] human model.

We thus present HOT (“Human-Object conTact”), a new dataset of images with human-object contact; see examples in Fig. 2. The first part of HOT, called “HOT-Generated” (Fig. 2b), has automatic annotations, but lacks variety for human subjects and scenes. The second part, called “HOT-Annotated” (Fig. 2a), has manual annotations, but has a huge variety of people, scenes and interactions. HOT has 35, 750 images with 163, 534 contact annotations.

We then train a new contact detector on our HOT dataset. Given a single color image as input, we want to know, if contact takes place in the image, the area in which it occurs, as well as the body part that is involved. Specifically, we detect 2D heatmaps in an image, encoding the contact location and likelihood, and classify each pixel in contact to one of SMPL(-X)’s body parts. However, training directly with HOT annotations leads to “bleeding” heatmaps and false detections. We observe that humans reason about contact by looking at body parts and their proximity to objects in their local vicinity. Therefore, we use a body-part-driven attention module that significantly boosts performance.

We evaluate our detector on withheld parts of the HOT dataset. Quantitative evaluation and ablation studies show that our model outperforms the baselines, and that all components contribute to detection performance. Our body-part attention module is the key component; a visual analysis shows that it attends to meaningful image locations, i.e., on body parts and their vicinity. Qualitative results show reasonable detections on in-the-wild images. By applying our detector on datasets unseen during training, we show that the model generalizes reasonably well; see Fig. 1. Then, we show that our general-purpose full-body contact detector performs on par with existing part-specific contact detectors for the foot [41] or hand [35], meaning it could serve as a drop-in replacement for these. Moreover, we show that our contact detector helps contact-driven 3D human pose estimation on PROX data [17]. Finally, we show that our HOT dataset helps a state-of-the-art (SOTA) 3D body-scene contact estimator [20] generalize to in-the-wild images.




Figure 2. Images and contact annotations for our HOT dataset. We show examples for both its parts, i.e., “HOT-Generated” (Sec. 3.2) and “HOT-Annotated” (Sec. 3.3). Contact annotations include the involved body part (c), shown color coded on a SMPL-X mesh.

In summary, HOT takes a step towards automatic contact detection between humans and objects in color images and our contributions are three-fold: (1) We introduce the task of full-body human-object contact detection in images. (2) To facilitate machine learning for this, we introduce the HOT dataset with 2D contact area heatmaps and the associated human part labels as annotations, using both auto-generated and manual annotations. (3) We develop a new contact detector that incorporates a body-part attention module. Experiments and ablations show the benefits of the proposed model and its components. Our data and code are available at <https://hot.is.tue.mpg.de>.

2. Related Work

Contact Modeling. Prior work on modeling the contact between humans and the world can be classified as either “body-object” or “hand-object.”

Body-object contact: Several works model the contact between the human body and object [8, 9, 21, 23, 56, 58]. Li et al. [28] reconstruct the 3D motion of a person interacting with an object by estimating the 3D pose of the person and object, the joint-level contact, and forces and torques actuated by the human limbs. Rempe et al. [40, 41] estimate joint-level foot-ground contact from a video, and use it to constrain the human pose with trajectory optimization. More recently, [20, 46] propose to directly estimate 3D body-scene contact, but the lack of data variety limits the model’s generalization to in-the-wild scenarios, even when a 3D scene is used as additional input [46]. Others use Human-Object Interaction (HOI) relationships to reconstruct [2, 5, 29, 53] or

generate [18, 62, 64] 3D human and object pose by encouraging contact and penalizing inter-penetrations. Prior work uses contact information as prior knowledge, but it is often oversimplified as (1) body-ground contact at the skeleton-joint level, (2) hand-object contact at a rough bounding-box level, or (3) manually-annotated contacts of other human parts in constrained environments. In this paper, we seek to automatically estimate contact heatmaps for the whole body in a bottom-up manner directly from a 2D image. We also predict the associated human body part label, which provides a more systematic understanding of human-object contact.

Hand-object contact: People interact with objects using their hands, so contact plays an important role in hand-object interaction and grasping. Contact information is often captured as a byproduct in grasp datasets [3, 11, 31, 47] through hand-object proximity or thermal information. Hand-object grasp reconstruction also employs contact to refine the hand and object pose estimation [4, 14, 19, 50, 52]. In addition, some works [35, 45, 60] detect hands and classify their physical contact state into self-contact, person-person contact, and person-object contact. Although they consider the relationship between hands and other objects in the scene, they detect only a rough bounding box or boundary for the hand, instead of a finer-grained contact area. In this paper, we take a step further to estimate general-purpose full-body contact from 2D images at a finer scale.

Human-Object Interaction (HOI). The goal of HOI understanding [37, 51, 57] is to infer the interaction relationships between humans and objects. While both humans and objects are located in the image, often in the form of 2D bounding boxes, the literature rarely focuses on how the interaction takes place, whether the interaction requires contact, and which human part is involved in the contact. This limits the applicability of current HOI detections for downstream scene understanding tasks. Recently, Li et al. [27] provide more detailed body-part state annotations in the context of HOI, and offer action labels (e.g., hold, paddle) and the involved human body part (e.g., hand, foot). However, they do not annotate 2D contact areas in images, and their predefined human parts are not fine-grained enough to capture everyday HOI scenarios. In contrast, our new dataset contains 2D contact areas that are also associated with the involved human body parts following the part segmentation of the popular SMPL(-X) [32, 36] statistical 3D body model.

Affordance Learning. Contact and HOI are closely related to object affordances, which reflect the functional aspects of an object. Recent work explores object affordance learning from human actions and object manipulation [13, 25, 34, 67]. More specifically, Fang et al. [12] and Nagarajan et al. [33] learn to predict the interaction region with the corresponding action label on a target object from

human demonstration videos. Savva et al. [43, 44] capture physical contact and visual attention links between 3D geometry and human body parts from RGB-D videos. Deng et al. [10] collect a 3D visual affordance dataset with potential interaction areas on 3D objects for various actions. Affordance learning is object-centric; it does not capture much about the human actor. In contrast, detecting interaction areas (e.g. contact heatmaps) reflects how people interact with objects, and considers both the human and the object.

3. Human-Object conTact (HOT) Dataset

To facilitate research in contact estimation, we introduce HOT, a new dataset with 2D contact areas and the associated human part labels as annotations. Annotating and detecting contact in images is challenging, as contact depends on the scene and its objects, the humans, the camera view and the occlusions arising from all these factors. To create a well-varied dataset, we collect images from two different sources and gather contact annotations. Below we discuss the creation of HOT and provide a comprehensive analysis.

3.1. Data Sources

“HOT-Generated”: First we collect data from the PROX dataset [17], which contains people reconstructed as 3D SMPL-X [36] meshes interacting with static 3D scenes; this involves actions like sitting, walking, lying down, etc. Recent work [40, 61] improves on the quality of reconstructed meshes in PROX, facilitating the automatic generation of contact heatmaps by simply using 3D proximity metrics between the 3D human mesh and the static 3D scene mesh. We sub-sample frames from the “qualitative set” of PROX, and form the “HOT-Generated” part of HOT.

“HOT-Annotated”: Another source for images with human-object contact is HOI datasets like V-COCO [16] and HAKE [27]. As they are collected from Flickr, these datasets contain very diverse HOI interactions in complex and cluttered scenes. Existing HOI datasets contain activity labels and bounding boxes for humans and objects, but boxes are too coarse for understanding contact. Thus, we select a subset from the V-COCO [16] and HAKE [27] datasets and use these to gather new contact annotations. To keep the task tractable, we first remove images with indirect human-object interaction, heavily cropped humans, motion blur, distortion or extreme lighting conditions. Other interesting datasets are indoor action recognition datasets like Watch-n-Patch [54] that contain several daily activities like “fetch-from-fridge”, “put book back”, etc. We sample image frames from video clips where human subjects and objects are clearly visible. We then combine images selected from V-COCO [16], HAKE [27] and Watch-n-Patch [54], and form the “HOT-Annotated” part of HOT.

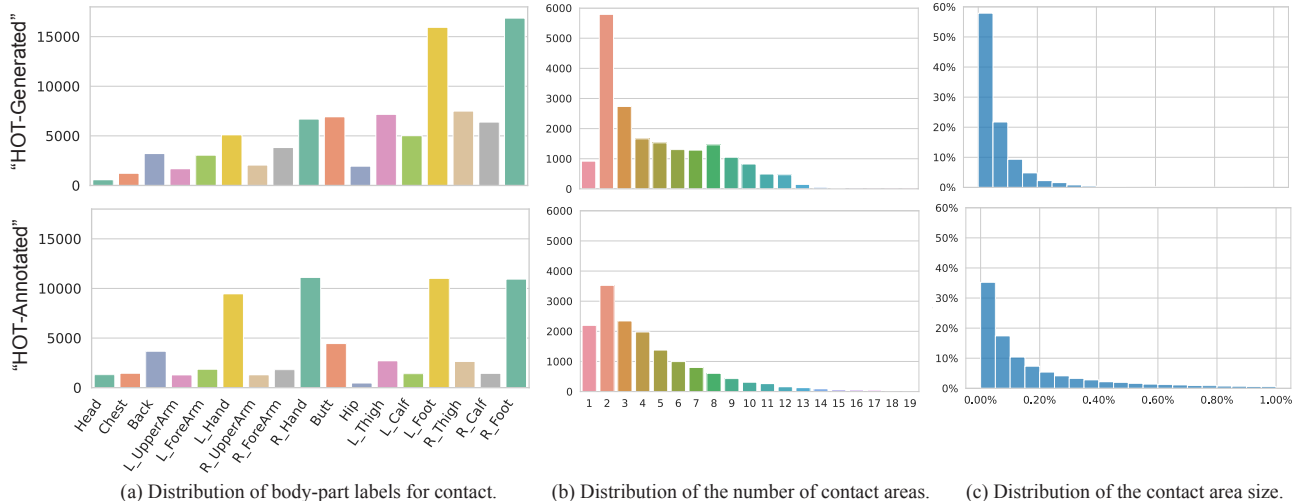


Figure 3. Data statistics. **(a)** Number of contact areas (Y-axis) for each body part (X-axis). **(b)** Number of images (Y-axis) with a certain number of contact areas (X-axis). **(c)** Percentage of contact areas (Y-axis) that occupy a certain percentage of image pixels (X-axis).

3.2. Contact Generation for “HOT-Generated”

The PROX dataset [17] captures people interacting with static scenes. PROX represents the human pose and shape with the SMPL-X [36] body model, which captures the body surface including the hands and face. The SMPL-X model represents the human body with pose parameters θ , and shape parameters β , and outputs a posed 3D mesh, $\mathcal{M}_b \in \mathbb{R}^{10475 \times 3}$. Each vertex, $v \in \mathbb{R}^3$, has a surface normal, n^v , and a human part label, c , associated with it. We divide the SMPL-X model into 17 parts c_i , with $i \in \{1, 2, \dots, 17\}$. The body parts are based on the original part segmentation of SMPL-X and, for simplicity, we merge parts that human annotators cannot easily differentiate; e.g., parts of the back across the spine. See Fig. 2c for the color-coded segmentation of SMPL-X into corresponding body-part labels.

For each frame, given the reconstructed SMPL-X mesh, \mathcal{M}_b , and the scene mesh, \mathcal{M}_s , we calculate human-to-scene mesh distances. Then, all human vertices with a distance to the scene below a threshold, and with compatible normals to the scene ones, are annotated as contact vertices. Finally, for the contact vertices we find the respective triangles on the 3D body mesh, and render these separately per body part to get dense 2D contact areas in the image space. In this way, we automatically create pseudo ground truth for contact. Examples are shown in Fig. 2b. For more details on the above steps, see Sup. Mat.

3.3. Contact Annotation for “HOT-Annotated”

We gather contact annotations for in-the-wild images without paired 3D human and scene meshes. Determining contact areas between a human and an object in an image is non-trivial, as contact areas are always occluded. While Amazon Mechanical Turk is popular for annotation collec-

tion, the diverse background of its annotators complicates the training, annotation, and quality control for our novel task. Thus, we hired a professional company with annotators that are already trained for similar tasks. The annotation has two steps: (1) drawing a polygon around the image areas containing human-object contacts, and (2) assigning a human body-part label to it. See Fig. 2a for annotation examples.

We take a number of steps to ensure good quality and consistency for the annotations. In particular, we first have a *trial annotation* with 3 rounds for 400 images; we provide task instructions, collect annotations, provide feedback to annotators, and iterate. Then, during the *real annotation*, 12 people perform the initial annotation, and we have two rounds of quality checks (QC); a different cohort of 4 people perform the first QC round, and another 2 people perform the second QC round. We use semantic-segmentation annotation tools adapted for our task. For more details on the annotation protocol, repeatability and quality check, see Sup. Mat.

Compared to the automatic annotations of “HOT-Generated”, the manual ones of “HOT-Annotated” are only approximate. However, capturing 3D people in scenes and accurately reconstructing them in 3D is hard and does not scale. Thus, manual 2D annotations are important because they allow scaling to large, varied, and natural datasets with images taken in the wild. For data-driven models, this helps improve their generalization and make them robust.

3.4. HOT Dataset Analysis

The HOT dataset has a total of 35,287 images and 162,267 contact area annotations, along with a body-part label for each area. Specifically, for “HOT-Annotated” we collect 5,235 images and 20,273 contact areas for V-COCO [16], 9,522 images and 45,645 contact areas for HAKE [27], and 325 images and 1,170 contact areas for

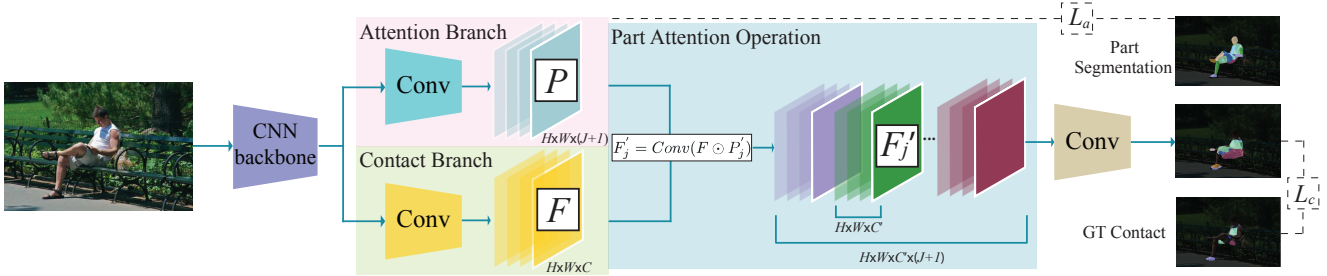


Figure 4. Architecture for our HOT contact detector. Our model takes as input a single color image, and as output it gives 2D contact heatmaps and a pixel-level classification label for the body part associated with contact. For a detailed explanation of the model, see Sec. 4.

the Watch-n-Patch [54] dataset. For “HOT-Generated”, we auto-generate 95, 179 contact areas in 20, 205 images using the PROX dataset. More statistics of “HOT-Annotated” and “HOT-Generated” are shown in Fig. 3.

Figure 3a shows the distribution of body-part labels for contact. We see that “HOT-Annotated” has noticeably more contacts than “HOT-Generated” for both hands. The reason is that PROX captures humans interacting with static scenes, i.e., without grasping and moving objects with their hands, while “HOT-Annotated” contains a lot of images with interactions between hands and objects.

Figure 3b shows the number of contact areas per image. We see that “HOT-Annotated” has generally more contacts per image than “HOT-Generated”. This is partially because HOI datasets contain images of multiple interacting persons, while PROX only has a single person in every image.

Figure 3c shows the distribution of contact area size. We observe that the areas are generally smaller for “HOT-Generated” than for “HOT-Annotated”. This is potentially because images in PROX are captured with the camera away from the body to include more scene context, whereas images in “HOT-Annotated” are taken in the wild, including close-ups, as well as more object grasps.

To further analyze the disparities between “HOT-Annotated” and “HOT-Generated”, we had two trained annotators annotate 200 random images from “HOT-Generated”, and we treat these manual annotations as ground truth for evaluating the automatically generated ones. The agreement for body-part contact labels is 83.7%, while for pixel contact labels, it is 52.4%. This can be attributed to PROX’s noisy SMPL-X reconstructions due to ambiguities arising from the monocular cameras and strong occlusions. Another contributing factor is the approximate nature of the manual annotations. We provide more discussions in the transfer experiments in Sec. 5.1.

4. HOT Contact Detector

To estimate contact areas in images, humans use the global context of the image, but also focus on regions around body parts to examine if there is contact. Based on these insights, we design our contact detector to extract global features with special attention to human body parts.

4.1. Architecture

Figure 4 shows the architecture of our model. Given an image, we use a CNN backbone to extract image features. Then, we use a decoder with two branches: an “attention branch” for inferring an attention mask for each body part and a “contact branch” for extracting contact features.

Attention branch: We denote this as $P \in \mathbb{R}^{H \times W \times (J+1)}$; J is the number of body parts, with an extra channel for the background, and H/W are the feature map’s height/width. The j th channel $P_j \in \mathbb{R}^{H \times W}$ represents the likelihood that each pixel is associated with contact of the j th body part. This guides the model to focus around different human parts in the feature space F of the contact branch. By applying a channel-wise softmax normalization $\sigma(\cdot)$ on P , we get the attention mask $P' = \sigma(P)$, with $P' \in \mathbb{R}^{H \times W \times (J+1)}$.

Contact branch: We denote this as $F \in \mathbb{R}^{H \times W \times C}$, with the same spatial dimensions $H \times W$ as the attention branch P , but with a different number of channels $C = 512$.

Part attention operation: We use the attention mask P'_j to extract part-related features from the contact branch F :

$$F'_j = \text{Conv}(F \odot P'_j), \quad \text{with } F'_j \in \mathbb{R}^{H \times W \times C'}, \quad (1)$$

where \odot is the element-wise product between all channels in F and the attention mask P'_j . We concatenate F'_j for all J parts and the background along the 3rd dimension as $F' \in \mathbb{R}^{H \times W \times C^*}$, where $C^* = C'(J+1)$ and $C' \neq C$, and pass it through a convolutional layer for per-pixel inference.

Supervision: We supervise the attention branch with part-segmentation maps (see “dataset splits” in Sec. 5.1 for the data source), and the contact branch with our contact annotations. Both branches classify pixels as being either a human part or the background; for the contact branch “background” means “no contact.” The overall loss is:

$$L = \lambda_a L_a + \lambda_c L_c, \quad (2)$$

where L_a is a cross-entropy loss between the estimated attention maps and ground-truth part-segmentation maps, L_c is a cross-entropy loss between the estimated and ground-truth contact maps, and λ_a and λ_c are steering weights.

4.2. Implementation Details

During training, body-part supervision for the attention branch is applied only in the initial stages, following Ko-




Figure 5. Attention visualization. The attention maps from the attention branch, supervised by human part segmentation, guide features to attend to each body-part area, but also to surrounding areas, for reasoning about the nearby scene context and the body-scene interaction.

cabas et al. [24]; λ_a is set to 0 at later stages. We use a pre-trained dilated ResNet-50 [59] as image encoder backbone. For the attention branch we use 3×3 convolutional layers with batch-norm and ReLU as image decoder. For the contact branch, we apply 3×3 convolutional layers with batch-norm and ReLU on the part-specific features. Since the background dominates the ground truth labels for both human part segmentation and contact estimation, we assign a smaller weight for the background label in the cross-entropy loss. For more details, see [Sup. Mat.](#)

5. Experiments

5.1. Contact Detection

Dataset splits: For “HOT-Annotated”, we randomly split the collected images into a training, validation and test set, resulting in 10,482 images for training, 2,300 for validation and 2,300 images for testing. For “HOT-Generated”, we split the training and testing set based on the scene; this results in 14,144 images for training, 3,031 for validation and 3,030 for testing. To supervise the attention branch, we obtain pseudo ground truth for human part segmentation by rendering part-segmented SMPL-X meshes into per-part masks in image space; we use LEMO’s [61] SMPL-X fits for the PROX dataset [17] and use FrankMocap [42] to estimate SMPL-X meshes for the images of “HOT-Annotated”.

Evaluation protocol: We adopt the evaluation protocol of Zhou et al. [66]; this is originally for semantic segmentation. We add one metric for contact prediction to evaluate whether the model distinguishes between contact and non-contact, i.e., “background”. We use the following metrics:

- *Semantic contact accuracy (SC-Acc.)*: The proportion of pixels that are both correctly classified as in contact and associated with the correct body-part label.
- *Contact accuracy (C-Acc.)*: The proportion of correctly classified pixels for binary contact labels; this ignores the body-part label, in contrast to “SC-Acc.”
- *Mean IoU (mIoU)*: The Intersection-over-Union (IoU) between the predicted and the ground-truth contact pixels, averaged over all the body-part labels.
- *Weighted IoU (wIoU)*: mIoU weighted by the pixel ratio

of each contact label.

To study the influence of the “HOT-Annotated” and “HOT-Generated” sets of the HOT dataset, we report performance by training and testing models separately on these, as well as on their combination that we denote as “Full Set”. For the “Full Set”, we randomly choose images from the “HOT-Generated” so that the number of training and testing images from both sets are the same.

Baselines: There exists no model for full-body contact detection in images, other than ours. Thus, we evaluate contact detection for two models, ResNet+UperNet [55] and ResNet+PPM [65], originally proposed for segmentation.

Ablations: We evaluate two variants of our proposed model to ablate the contribution of the attention branch: “Ours_{wo/att}” without the attention branch and “Ours_{pure.att}” without supervision for the attention branch, which functions as an unsupervised pure soft-attention module.

Results & discussion: Quantitative results for contact detection are shown in [Tab. 1](#), and qualitative results are shown in [Fig. 6](#). Below we summarize key findings.

1. Our model outperforms state-of-the-art (SOTA) methods [55,65] developed for semantic segmentation (retrained for our task). This is due to the different nature of semantic scene understanding and contact estimation. The former relies on dense pixel annotations of the entire scene and global contextual features. The latter relies on sparser annotations and needs an attention mechanism to focus around humans.
2. Our attention mechanism guides our model to learn better features that improve contact estimation. “Ours_{pure.att}”, which uses unsupervised pure soft-attention, outperforms “Ours_{wo/att}” which has no attention branch. By adding supervision on human-part segmentation in early training stages, the attention focuses on areas around each human part. Intuitively, this helps reasoning about contact by using both human-body and surrounding-object information; see [Fig. 5](#) for a visualization of the learned attention maps.
3. Learning on “HOT-Generated” is more difficult than on “HOT-Annotated”. This is partially because, even though we generate contact annotations from relatively “clean” SMPL-X fits by LEMO [61], which reasons about temporal continuity and occlusion, these are still a bit noisy.




Figure 6. Qualitative results of our HOTA contact detector for withheld images of our HOTA dataset. For each input image we show the following triplet: {Raw input image, Ground-truth contact, Predicted contact}.

Model	"HOT-Annotated"				"HOT-Generated"				"Full Set"			
	SC-Acc.↑	C-Acc↑	mIoU↑	wIoU↑	SC-Acc.↑	C-Acc↑	mIoU↑	wIoU↑	SC-Acc.↑	C-Acc↑	mIoU↑	wIoU↑
ResNet+UperNet [55]	35.1	62.6	0.195	0.227	21.1	42.7	0.080	0.116	32.5	62.4	0.187	0.214
ResNet+PPM [65]	34.6	61.1	0.201	0.233	21.2	41.1	0.075	0.119	31.5	58.4	0.176	0.212
Ours _{wo/att}	24.1	42.8	0.148	0.187	12.0	24.6	0.051	0.099	19.4	29.3	0.130	0.155
Ours _{pure_att}	33.8	58.4	0.189	0.237	20.3	40.1	0.077	0.113	30.4	55.9	0.163	0.206
Ours _{Full}	40.7	70.7	0.215	0.260	30.4	54.3	0.139	0.167	36.4	66.3	0.209	0.251

Table 1. Evaluation of contact detection accuracy on the HOTA dataset.

Train	Test	SC-Acc.↑	C-Acc. ↑	mIoU↑	wIoU↑
HOT-Gen	HOT-Gen	30.4	54.3	0.139	0.167
HOT-Ann	HOT-Gen	28.4	51.8	0.122	0.203
Full Set	HOT-Gen	34.3	59.2	0.140	0.205
HOT-Gen	HOT-Ann	2.46	6.37	0.019	0.042
HOT-Ann	HOT-Ann	40.7	70.7	0.215	0.260
Full Set	HOT-Ann	47.4	79.2	0.232	0.273

Table 2. Transfer across "HOT-Generated" ↔ "HOT-Annotated".

Fine-grained contact detection is sensitive to strong occlusions during interactions, motion blur, the low resolution for people observed by indoor-monitoring cameras, and the imperfect "hallucination" of SOTA pose estimation methods [40, 61] for resolving these ambiguities. This shows the value of "HOT-Annotated", i.e. the collection of a high-quality dataset of in-the-wild images with rich manual contact annotations, and points to important future work.

4. We conduct transfer experiments across "HOT-Annotated" and "HOT-Generated" and the results are shown in Tab. 2. The model trained on "HOT-Annotated" generalizes well to "HOT-Generated", but not vice versa. This is mainly because the former is captured in the wild and has rich variation, while the latter is captured in constrained settings. It is also noteworthy that combining both datasets ("Full Set") boosts

performance, suggesting that automatic contact annotations like "HOT-Generated" are beneficial for this task.

We discuss the failure cases and report our model's performance under different settings, i.e., contact for different body parts and various contact area sizes; see Sup. Mat.

5.2. Full-Body vs Part-Specific Contact

To evaluate the robustness of our general-purpose full-body HOTA contact detector, we compare it against two existing part-specific contact detectors, as shown in Fig. 7:

(i) **Foot contact:** "ContactDynamics" [41] estimates *joint-level* foot-ground contact from a *video*. We evaluate our model and "ContactDynamics" against the ground-truth foot contact from PROX's "quantitative set". Our detector achieves a similar performance (HOTA **59.2%** vs "ContactDynamics" 58.6%). Thus, it could be a drop-in replacement contact detector for 3D body pose estimation [41].

(ii) **Hand contact:** "ContactHands" [35] detects hands and classifies them into "self-contact", "person-person contact", and "person-object contact". We evaluate our HOTA detector and "ContactHands" on hand-object contact on a subset of the "HOT-Annotated" test set. We report contact recognition accuracy under an IoU threshold of 0.4; our detector achieves similar performance (HOTA **63.5%** vs [35] 62.2%).




Figure 7. Comparison of our general-purpose full-body contact detector (HOT) against existing part-specific detectors, namely “ContactDynamics” [41] for joint-level foot-ground contact, and “ContactHands” [35] for bounding-box-level hand contact.

Method	No Contact	PROX [17]	Predicted Contact	GT Contact
V2V ↓	183.3	174.0	172.3	163.0

Table 3. Contact-driven human pose estimation performance.

Train	Test	precision ↑	recall ↑	f1 ↑
RICH [20]	RICH	0.699	0.744	0.708
RICH+HOT-pGT	RICH	0.675	0.761	0.684
RICH	HOT	0.439	0.192	0.231
RICH+HOT-pGT	HOT	0.684	0.701	0.636

Table 4. 3D dense contact estimation performance.

Discussion: As shown in Fig. 7, “ContactDynamics” simply classifies foot joints as in contact or not, and “ContactHands” detects hands as bounding boxes, while we generalize to the full body and detect richer heatmaps. In Sup. Mat., we provide more details and some preliminary results when testing our model on self-contact and human-human contact. The fact that our full-body contact detector performs on par with existing part-expert detectors holds promise for developing a general purpose contact detector for human-object and human-human interactions in future work.

5.3. HOT Contact Detection vs Heuristic Contact

The PROX dataset [17] is widely used for developing and evaluating HOI methods. Its human meshes have been reconstructed with an optimization method that fits SMPL-X to images, paired with an a-priori known (i.e., pre-scanned) 3D scene. The human meshes look physically plausible, as the method encourages (manually annotated) “likely contact” body vertices that lie close to the scene to contact it while not penetrating it; this resolves pose errors.

We replace PROX’s manually annotated “likely contact” vertices with the ones of the body parts that our detector suggests are in contact, given the input image. We call this setup “Predicted Contact” and evaluate this on PROX’s “quantitative set” via the Vertex-to-Vertex (V2V) error. We also evaluate a baseline with “No Contact” constraints. For a fair comparison, for all baselines we use the same optimization process as in PROX [17]. Results in Tab. 3 show that our “Predicted Contact” is on par with “PROX”, indicating that detecting contact in images is promising for replacing PROX’s handcrafted heuristics. We also simulate a perfect




Figure 8. Qualitative results of 3D contact estimation on HOT.

contact detector using PROX’s ground truth (“GT Contact”). This shows that there is room and merit for improving image-based contact detection in future work.

5.4. HOT for 3D Contact on Bodies

The recent RICH dataset and BSTRO model [20] focus on dense 3D contact estimation on the human body from an image. To show the usefulness of our HOT dataset for this task, we “lift” our 2D annotations to coarse 3D ones, by annotating the respective 3D SMPL parts as in contact, and treat these as pseudo ground-truth (HOT-pGT). We then employ the BSTRO model [20], extend its training dataset with HOT-pGT, and re-train. We report performance in Tab. 4 and Fig. 8. The model trained on RICH data alone fails on HOT images, which are taken in the wild, in contrast to RICH images. Interestingly, adding HOT-pGT for training improves the 3D contact estimation accuracy for in-the-wild images, while not hurting for RICH data. This shows that our HOT dataset can help for 3D contact estimation and related applications. For details, see Sup. Mat.

6. Conclusion

We focus on human-object contact detection for images. To this end, we collect the HOT dataset and develop the HOT contact detector with human-part guided attention. Our detector outperforms baseline models and generalizes reasonably well for in-the-wild images. One limitation is that we build our method upon “simple” convolutional models; however our key insight (human part attention) is general and agnostic to the model’s architecture. We believe that this new task, dataset and model fill a gap in the literature and hope they can inspire more future endeavors into this topic, utilizing more complex models like transformers [6, 7] and exploring a wider range of applications.

Acknowledgments: We thank Chun-Hao Paul Huang for his valuable help with the RICH data and the training code of BSTRO [20]. We thank Lea Müller, Mohamed Hassan, Muhammed Kocabas, Shashank Tripathi, and Hongwei Yi for the insightful discussions. We thank Benjamin Pellkofer for IT help, and Nicole Overbaugh and Johanna Werminghausen for administrative help. This work was partially funded by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039B. DT’s work was mostly performed at the MPI-IS.

Disclosure: https://files.is.tue.mpg.de/black/CoI_CVPR_2023.txt

References

- [1] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems (RAS)*, 57(5):469–483, 2009. [1](#)
- [2] Bharat Lal Bhatnagar, Xianghui Xie, Ilya A Petrov, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. BEHAVE: Dataset and method for tracking human object interactions. In *Computer Vision and Pattern Recognition (CVPR)*, pages 15935–15946, 2022. [2](#)
- [3] Samarth Brahmabhatt, Chengcheng Tang, Christopher D. Twigg, Charles C. Kemp, and James Hays. ContactPose: A dataset of grasps with object contact and hand pose. In *European Conference on Computer Vision (ECCV)*, volume 12358, pages 361–378, 2020. [3](#)
- [4] Zhe Cao, Ilija Radosavovic, Angjoo Kanazawa, and Jitendra Malik. Reconstructing hand-object interactions in the wild. In *International Conference on Computer Vision (ICCV)*, pages 12417–12426, 2021. [3](#)
- [5] Yixin Chen, Siyuan Huang, Tao Yuan, Siyuan Qi, Yixin Zhu, and Song-Chun Zhu. Holistic++ scene understanding: Single-view 3D holistic scene parsing and human pose estimation with human-object interaction and physical commonsense. In *International Conference on Computer Vision (ICCV)*, pages 8648–8657, 2019. [2](#)
- [6] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1290–1299, 2022. [8](#)
- [7] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. In *Conference on Neural Information Processing Systems (NeurIPS)*, volume 34, pages 17864–17875, 2021. [8](#)
- [8] Henry M. Clever, Zackory Erickson, Ariel Kapusta, Greg Turk, Karen Liu, and Charles C. Kemp. Bodies at rest: 3D human pose and shape estimation from a pressure image using synthetic data. In *Computer Vision and Pattern Recognition (CVPR)*, pages 6215–6224, 2020. [2](#)
- [9] Yudi Dai, YiTai Lin, XiPing Lin, Chenglu Wen, Lan Xu, Hongwei Yi, Siqi Shen, Yuexin Ma, and Cheng Wang. SLOPER4D: A scene-aware dataset for global 4D human pose estimation in urban environments. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. [2](#)
- [10] Shengheng Deng, Xun Xu, Chaozheng Wu, Ke Chen, and Kui Jia. 3D AffordanceNet: A benchmark for visual object affordance understanding. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1778–1787, 2021. [3](#)
- [11] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J. Black, and Otmar Hilliges. ARCTIC: A dataset for dexterous bimanual hand-object manipulation. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. [3](#)
- [12] Kuan Fang, Te-Lin Wu, Daniel Yang, Silvio Savarese, and Joseph J. Lim. Demo2vec: Reasoning object affordances from online videos. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2139–2147, 2018. [3](#)
- [13] 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. In *European Conference on Computer Vision (ECCV)*, volume 7576, pages 732–745, 2012. [1, 3](#)
- [14] Patrick Grady, Chengcheng Tang, Christopher D. Twigg, Minh Vo, Samarth Brahmabhatt, and Charles C. Kemp. ContactOpt: Optimizing contact to improve grasps. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1471–1481, 2021. [3](#)
- [15] Shivam Grover, Kshitij Sidana, and Vanita Jain. Pipeline for 3D reconstruction of the human body from AR/VR headset mounted egocentric cameras. *arXiv:2111.05409*, 2021. [1](#)
- [16] Saurabh Gupta and Jitendra Malik. Visual semantic role labeling. *arXiv:1505.04474*, 2015. [2, 3, 4](#)
- [17] Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J. Black. Resolving 3D human pose ambiguities with 3D scene constraints. In *International Conference on Computer Vision (ICCV)*, pages 2282–2292, 2019. [2, 3, 4, 6, 8](#)
- [18] Mohamed Hassan, Partha Ghosh, Joachim Tesch, Dimitrios Tzionas, and Michael J. Black. Populating 3D scenes by learning human-scene interaction. In *Computer Vision and Pattern Recognition (CVPR)*, pages 14708–14718, 2021. [1, 3](#)
- [19] Yana Hasson, Gul Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J Black, Ivan Laptev, and Cordelia Schmid. Learning joint reconstruction of hands and manipulated objects. In *Computer Vision and Pattern Recognition (CVPR)*, pages 11807–11816, 2019. [3](#)
- [20] Chun-Hao P Huang, Hongwei Yi, Markus Höschle, Matvey Safroshkin, Tsvetelina Alexiadis, Senya Polikovskiy, Daniel Scharstein, and Michael J. Black. Capturing and inferring dense full-body human-scene contact. In *Computer Vision and Pattern Recognition (CVPR)*, pages 13274–13285, 2022. [1, 2, 8](#)
- [21] Siyuan Huang, Zan Wang, Puhao Li, Baoxiong Jia, Tengyu Liu, Yixin Zhu, Wei Liang, and Song-Chun Zhu. Diffusion-based generation, optimization, and planning in 3D scenes. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. [2](#)
- [22] Baoxiong Jia, Yixin Chen, Siyuan Huang, Yixin Zhu, and Song-chun Zhu. Lemma: A multi-view dataset for learning multi-agent multi-task activities. In *European Conference on Computer Vision (ECCV)*, volume 12371, pages 767–786, 2020. [1](#)
- [23] Nan Jiang, Tengyu Liu, Zhexuan Cao, Jieming Cui, Yixin Chen, He Wang, Yixin Zhu, and Siyuan Huang. CHAIRS: Towards full-body articulated human-object interaction. *arXiv preprint*, 2022. [2](#)
- [24] Muhammed Kocabas, Chun-Hao P. Huang, Otmar Hilliges, and Michael J. Black. PARE: Part attention regressor for 3D human body estimation. In *International Conference on Computer Vision (ICCV)*, pages 11127–11137, 2021. [6](#)
- [25] Hema S. Koppula and Ashutosh Saxena. Physically grounded spatio-temporal object affordances. In *European Conference on Computer Vision (ECCV)*, volume 8691, pages 831–847, 2014. [1, 3](#)

- [26] Vikash Kumar and Emanuel Todorov. Mujoco HAPTIX: A virtual reality system for hand manipulation. In *International Conference on Humanoid Robots (HUMANOIDS)*, pages 657–663, 2015. [1](#)
- [27] Yong-Lu Li, Liang Xu, Xinpeng Liu, Xijie Huang, Yue Xu, Shiyi Wang, Hao-Shu Fang, Ze Ma, Mingyang Chen, and Cewu Lu. PaStaNet: Toward human activity knowledge engine. In *Computer Vision and Pattern Recognition (CVPR)*, pages 382–391, 2020. [1](#), [2](#), [3](#), [4](#)
- [28] 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 *Computer Vision and Pattern Recognition (CVPR)*, pages 8640–8649, 2019. [2](#)
- [29] Zhi Li, Soshi Shimada, Bernt Schiele, Christian Theobalt, and Vladislav Golyanik. MoCapDeform: Monocular 3D human motion capture in deformable scenes. In *International Conference on 3D Vision (3DV)*, pages 1–11, 2022. [2](#)
- [30] Jeffrey I Lipton, Aidan J Fay, and Daniela Rus. Baxter’s homunculus: Virtual reality spaces for teleoperation in manufacturing. *Robotics and Automation Letters (RA-L)*, 3(1):179–186, 2017. [1](#)
- [31] Shaowei Liu, Hanwen Jiang, Jiarui Xu, Sifei Liu, and Xiaolong Wang. Semi-supervised 3D hand-object poses estimation with interactions in time. In *Computer Vision and Pattern Recognition (CVPR)*, pages 14687–14697, 2021. [3](#)
- [32] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *Transactions on Graphics (TOG)*, 34(6):248:1–248:16, 2015. [2](#), [3](#)
- [33] Tushar Nagarajan, Christoph Feichtenhofer, and Kristen Grauman. Grounded human-object interaction hotspots from video. In *International Conference on Computer Vision (ICCV)*, pages 8688–8697, 2019. [3](#)
- [34] Tushar Nagarajan and Kristen Grauman. Learning affordance landscapes for interaction exploration in 3D environments. In *Conference on Neural Information Processing Systems (NeurIPS)*, volume 33, pages 2005–2015, 2020. [1](#), [3](#)
- [35] Supreeth Narasimhaswamy, Trung Nguyen, and Minh Hoai. Detecting hands and recognizing physical contact in the wild. In *Conference on Neural Information Processing Systems (NeurIPS)*, volume 33, pages 7841–7851, 2020. [1](#), [2](#), [3](#), [7](#), [8](#)
- [36] 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 *Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019. [2](#), [3](#), [4](#)
- [37] Siyuan Qi, Wenguan Wang, Baoxiong Jia, Jianbing Shen, and Song-Chun Zhu. Learning human-object interactions by graph parsing neural networks. In *European Conference on Computer Vision (ECCV)*, volume 11213, pages 401–417, 2018. [1](#), [3](#)
- [38] Ilija Radosavovic, Xiaolong Wang, Lerrel Pinto, and Jitendra Malik. State-only imitation learning for dexterous manipulation. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 7865–7871, 2020. [1](#)
- [39] Nishant Rai, Haofeng Chen, Jingwei Ji, Rishi Desai, Kazuki Kozuka, Shun Ishizaka, Ehsan Adeli, and Juan Carlos Niebles. Home action genome: Cooperative compositional action understanding. In *Computer Vision and Pattern Recognition (CVPR)*, pages 11184–11193, 2021. [1](#)
- [40] Davis Remppe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J. Guibas. HuMoR: 3D human motion model for robust pose estimation. In *International Conference on Computer Vision (ICCV)*, pages 11488–11499, 2021. [2](#), [3](#), [7](#)
- [41] Davis Remppe, Leonidas J. Guibas, Aaron Hertzmann, Bryan Russell, Ruben Villegas, and Jimei Yang. Contact and human dynamics from monocular video. In *European Conference on Computer Vision (ECCV)*, volume 12350, pages 71–87, 2020. [1](#), [2](#), [7](#), [8](#)
- [42] Yu Rong, Takaaki Shiratori, and Hanbyul Joo. FrankMocap: A monocular 3D whole-body pose estimation system via regression and integration. In *International Conference on Computer Vision Workshops (ICCVw)*, pages 1749–1759, 2021. [6](#)
- [43] Manolis Savva, Angel X. Chang, Pat Hanrahan, Matthew Fisher, and Matthias Nießner. SceneGrok: Inferring action maps in 3D environments. *Transactions on Graphics (TOG)*, 33(6):1–10, 2014. [3](#)
- [44] Manolis Savva, Angel X. Chang, Pat Hanrahan, Matthew Fisher, and Matthias Nießner. PiGraphs: Learning interaction snapshots from observations. *Transactions on Graphics (TOG)*, 35(4):1–12, 2016. [1](#), [3](#)
- [45] Dandan Shan, Jiaqi Geng, Michelle Shu, and David F Fouhey. Understanding human hands in contact at internet scale. In *Computer Vision and Pattern Recognition (CVPR)*, pages 9869–9878, 2020. [1](#), [3](#)
- [46] Soshi Shimada, Vladislav Golyanik, Zhi Li, Patrick Pérez, Weipeng Xu, and Christian Theobalt. HULC: 3D human motion capture with pose manifold sampling and dense contact guidance. In *ECCV*, volume 13682, pages 516–533, 2022. [1](#), [2](#)
- [47] Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. GRAB: A dataset of whole-body human grasping of objects. In *European Conference on Computer Vision (ECCV)*, volume 12349, pages 581–600, 2020. [3](#)
- [48] Anand Thobbi and Weihua Sheng. Imitation learning of hand gestures and its evaluation for humanoid robots. In *International Conference on Information and Automation (ICIA)*, pages 60–65, 2010. [1](#)
- [49] Shashank Tripathi, Lea Müller, Chun-Hao P. Huang, Taheri Omid, Michael J. Black, and Dimitrios Tzionas. 3D human pose estimation via intuitive physics. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. [1](#)
- [50] Tze Ho Elden Tse, Zhongqun Zhang, Kwang In Kim, Ales Leonardis, Feng Zheng, and Hyung Jin Chang. S2 contact: Graph-based network for 3D hand-object contact estimation with semi-supervised learning. In *European Conference on Computer Vision (ECCV)*, volume 13661, pages 568–584, 2022. [3](#)
- [51] Tiancai Wang, Rao Muhammad Anwer, Muhammad Haris Khan, Fahad Shahbaz Khan, Yanwei Pang, Ling Shao, and

- Jorma Laaksonen. Deep contextual attention for human-object interaction detection. In *International Conference on Computer Vision (ICCV)*, pages 5694–5702, 2019. 1, 3
- [52] Yufei Wang, David Held, and Zackory Erickson. Visual haptic reasoning: Estimating contact forces by observing deformable object interactions. *Robotics and Automation Letters (RA-L)*, 7(4):11426–11433, 2022. 3
- [53] Zhenzhen Weng and Serena Yeung. Holistic 3D human and scene mesh estimation from single view images. In *Computer Vision and Pattern Recognition (CVPR)*, pages 334–343, 2021. 2
- [54] Chenxia Wu, Jiemi Zhang, Silvio Savarese, and Ashutosh Saxena. Watch-n-Patch: Unsupervised understanding of actions and relations. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4362–4370, 2015. 2, 3, 5
- [55] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *European Conference on Computer Vision (ECCV)*, volume 11209, pages 418–434, 2018. 6, 7
- [56] Xianghui Xie, Bharat Lal Bhatnagar, and Gerard Pons-Moll. CHORE: Contact, human and object reconstruction from a single RGB image. In *European Conference on Computer Vision (ECCV)*, volume 13662, pages 125–145, 2022. 2
- [57] Bingjie Xu, Yongkang Wong, Junnan Li, Qi Zhao, and Mohan S. Kankanhalli. Learning to detect human-object interactions with knowledge. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2019–2028, 2019. 1, 3
- [58] Hongwei Yi, Chun-Hao P. Huang, Shashank Tripathi, Lea Hering, Justus Thies, and Michael J. Black. MIME: Human-aware 3D scene generation. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [59] Fisher Yu, Vladlen Koltun, and Thomas Funkhouser. Dilated residual networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 472–480, 2017. 6
- [60] Lingzhi Zhang, Shenghao Zhou, Simon Stent, and Jianbo Shi. Fine-grained egocentric hand-object segmentation: Dataset, model, and applications. In *European Conference on Computer Vision (ECCV)*, volume 13689, pages 127–145, 2022. 3
- [61] Siwei Zhang, Yan Zhang, Federica Bogo, Pollefeys Marc, and Siyu Tang. Learning motion priors for 4D human body capture in 3D scenes. In *International Conference on Computer Vision (ICCV)*, pages 11323–11333, 2021. 3, 6, 7
- [62] Siwei Zhang, Yan Zhang, Qianli Ma, Michael J. Black, and Siyu Tang. PLACE: Proximity learning of articulation and contact in 3D environments. In *International Conference on 3D Vision (3DV)*, pages 642–651, 2020. 1, 3
- [63] Tianhao Zhang, Zoe McCarthy, Owen Jow, Dennis Lee, Xi Chen, Ken Goldberg, and Pieter Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In *International Conference on Robotics and Automation (ICRA)*, pages 5628–5635, 2018. 1
- [64] Yan Zhang, Mohamed Hassan, Heiko Neumann, Michael J. Black, and Siyu Tang. Generating 3D people in scenes without people. In *Computer Vision and Pattern Recognition (CVPR)*, pages 6194–6204, 2020. 1, 3
- [65] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2881–2890, 2017. 6, 7
- [66] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ADE20K dataset. *International Journal of Computer Vision (IJCV)*, 127(3):302–321, 2019. 6
- [67] Yixin Zhu, Chenfanfu Jiang, Yibiao Zhao, Demetri Terzopoulos, and Song-Chun Zhu. Inferring forces and learning human utilities from videos. In *Computer Vision and Pattern Recognition (CVPR)*, pages 3823–3833, 2016. 1, 3