

# Learning Complex 3D Human Self-Contact

Mihai Fieraru<sup>1</sup> Mihai Zanfir<sup>1</sup> Elisabeta Oneata<sup>1</sup>  
Alin-Ionut Popa<sup>1</sup> Vlad Olaru<sup>1</sup> Cristian Sminchisescu<sup>2,1</sup>

<sup>1</sup> Institute of Mathematics of the Romanian Academy, <sup>2</sup> Lund University  
{mihai.fieraru, mihai.zanfir, elisabeta.oneata, alin.popa, vlad.olaru}@imar.ro, cristian.sminchisescu@math.lth.se

## Abstract

Monocular estimation of three dimensional human self-contact is fundamental for detailed scene analysis including body language understanding and behaviour modeling. Existing 3d reconstruction methods do not focus on body regions in self-contact and consequently recover configurations that are either far from each other or self-intersecting, when they should just touch. This leads to perceptually incorrect estimates and limits impact in those very fine-grained analysis domains where detailed 3d models are expected to play an important role. To address such challenges we detect self-contact and design 3d losses to explicitly enforce it. Specifically, we develop a model for Self-Contact Prediction (*SCP*), that estimates the body surface signature of self-contact, leveraging the localization of self-contact in the image, during both training and inference. We collect two large datasets to support learning and evaluation: (1) HumanSC3D, an accurate 3d motion capture repository containing 1,032 sequences with 5,058 contact events and 1,246,487 ground truth 3d poses synchronized with images collected from multiple views, and (2) FlickrSC3D, a repository of 3,969 images, containing 25,297 surface-to-surface correspondences with annotated image spatial support. We also illustrate how more expressive 3d reconstructions can be recovered under self-contact signature constraints and present monocular detection of face-touch as one of the multiple applications made possible by more accurate self-contact models.

## Introduction

Most monocular 3d human reconstruction systems do not directly infer human self-contact, although its central role in correctly recognizing the subtleties of many iconic poses or gestures is widely acknowledged perceptually. Current modeling deficiencies result in contact regions either far from each other or self-intersecting in the final 3d reconstruction, when they should instead just touch (e.g. contact between one’s hand and chin). In turn, unpredictable reconstructions of self-contact decrease the appeal of using 3d representations for fine grained analysis of behavior and intent, particularly as many self-touch events are elicited frequently, and with little or no human awareness. Correctly tracking self-contact would be invaluable not just for behavior analy-

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Our self-contact prediction network (*SCP*) estimates the body regions in contact, their correspondences and the self-contact positioning in image space.

sis but in assessing hygiene and possible health implications during a pandemic.

To overcome some of the shortcomings of existing, self-contact agnostic, 3d reconstruction methods, we propose to represent self-contact explicitly and show how the resulting models can assist behavioural understanding in applications assessing face touching. Our models learn to predict the image location of contact in order to assist the detection of body regions in self-contact, as well as their signature, defined as the correspondences between regions on the surface of a human body model that touch. Conditioned on such detailed estimates self-contact can be recovered correctly in the 3d reconstruction. To train models and for large-scale quantitative evaluation, we collect and annotate two large scale datasets containing images of people in self-contact. HumanSC3D is an accurate 3d motion capture dataset containing 1,032 sequences with 5,058 contact events and 1,246,487 ground truth 3d poses synchronized with images captured from multiple views. We also collect FlickrSC3D, a dataset of 3,969 images, containing 25,297 annotations of body part region pairs in contact, defined on a 3d human surface model, together with their self-contact localisation in the image.

The main contributions of the paper are as follows:

- Introduce a first principled model to detect self-contact body regions and their signature. Our novel deep neural network *SCP* is assisted by an intermediate self-contact image localisation (branch) predictor, leveraged both in training, for local feature selection, and in testing, by enforcing consistency with the estimated 3d contact signature.
- Novel, task-specific, large scale, valuable community

datasets capturing people in self-contact, together with dense annotations of a 3d body model to capture the surface regions in contact, as well as image annotations associated to the observed points of contact. The data and models will be made available for research.

- Quantitative and qualitative demonstration of metrically more accurate and perceptually veridical 3d reconstructions based on self-contact signatures.
- A foundation for a large class of applications that would benefit from accurate 3d self-contact representations, such as, health monitoring of possible infections when hands touch parts of the face (mouth, nose, eyes) in hospitals or during a pandemic, or subtle behavioral understanding of gestures for robot-assisted therapy of children with autism, to name just a few.

## Related Work

Automatic 3d human pose and shape estimation from images and video has been increasingly studied in recent years and significant progress has been made (Mehta et al. 2018; Zangir et al. 2018; Li et al. 2019a; Su et al. 2019; Benzine et al. 2019; Kolotouros et al. 2019; Kanazawa et al. 2019; Kocabas, Athanasiou, and Black 2020). These methods focus on 3d pose, and to some extent shape estimation, and person’s relative placement with respect to the scene. However, the subtleties of 3d shape especially in conjunction with contact are still largely unexplored, with vast potential for improvement well beyond existing art. Challenges include human-object interaction, inter-human interactions or human self-contact. In this paper we present models and insights - methodological, experimental and logistic, in terms of data collection - focusing on human self-contact. In the rest of this section we review previous work on human contact and self-contact applications.

**Self-Contact.** Most of previous work on self-contact (Tzionas et al. 2016; Tzionas and Gall 2013; Taylor et al. 2017; Mueller et al. 2019) applies to the interaction of human extremities, such as hands. Tzionas et al. (2016) introduces a method for modelling 3d hand to hand or hand to object interactions based on RGB-D data. The hand reconstruction is done via an energy-based modeling which incorporates physics and collision constraints. However the shape of the hand is not estimated, and only a standard template is used. Mueller et al. (2019) propose a similar real-time system based on a RGB-D sensor that is also able to estimate the shape and pose of interacting hands. None of the above methods explicitly detects the regions in self-contact or predicts their signature. In contrast, we handle full bodies, not only hand regions, and do not require depth data. Others (Bogo et al. 2016; Zangir, Marinoiu, and Sminchisescu 2018; Pavlakos et al. 2019) use non-self intersection constraints to prevent inadmissible 3d human reconstructions. Avoiding self-collisions only, and in the absence of any semantics of self-contact and the underlying surface regions, however, makes it difficult to enforce self-contact of surfaces when these actually touch.

**Human - Object/Scene Contact.** Contact between humans and the environment has also been studied recently. Hassan

et al. (2019) propose an optimisation method for 3d human shape estimation which incorporates scene constraints (including a contact-aware loss function) in the form of depth information. Leveraging the same contact loss, (Zhang et al. 2020) learn how to plausibly place 3d people in 3d scenes. Contact between human feet and the ground is also modeled in (Zou et al. 2020), as it has been previously done by (Zangir, Marinoiu, and Sminchisescu 2018), and used to constrain 3d human reconstruction. Interaction with objects is also studied in (Hasson et al. 2019), who jointly model the reconstruction of human hands and interacting objects based on single view RGB data. Li et al. (2019b) estimates contact positions, forces and torques.

**Human - Human Contact.** Contact between people is typical in close interactions like business meetings, informal conversations, or other social events. Liu et al. (2013, 2011) scan participants and rig them to a 3d skeleton. Given a green background setup, the motion in various scenarios is recorded and later estimated based on an energy model. Yet, the interaction between participants is modelled solely by non-self intersection constraints. In our recent work (Fieraru et al. 2020) we model and learn contact regions between people by means of their contact signature. However, we did not cover self-contact and the image localization of contact is not annotated or estimated. Still, in this work we also consider as baseline (Fieraru et al. 2020)’s ISP prediction method, adapted to estimate self-contact.

**Applications.** Self-contact prediction can enable numerous applications, and we only reference a few here. Kwok, Gralton, and McLaws (2015) study hygiene and virus transmission, by monitoring how often students touch their face with their hands. Yet, data is gathered manually with multiple investigators annotating videotape recordings. Mueller, Martin, and Grunwald (2019) also study the locations and durations of facial self-touches, but using accelerometers and EMG. An automatic labeling approach such as our *SCP* has the potential to enable the automation of larger such quantitative studies using only RGB sensors. Gesture analysis can benefit from improved self-contact signature predictions. For example, applause detection (Manoj et al. 2011) can possibly be performed from soundless videos as the detection of frequent self-contact between one’s hands. Similarly, a self-contact signature such as covering ears with both hands can be used as a signal of patients’ noise-sensitivity, with applications in robot-assisted autism therapy (Rudovic et al. 2017; Marinoiu et al. 2018).

## Methodology

In this section, we describe our proposed model *SCP* to predict the image spatial support of self-contact, the self-contact segmentation and the self-contact signature, as well as our model for 3d human reconstruction under self-contact constraints. Building up on our earlier work on modeling contact between people (Fieraru et al. 2020), we define the self-contact segmentation and signature of a person in an image  $I$  by discretizing of the surface of the human body model into  $N_R$  regions. In our case, region-level self-contact **signature**  $C^R(I) \in \{0, 1\}^{N_R} \times \{0, 1\}^{N_R}$  is defined as  $C_{r_1, r_2}^R(I) = 1$  when region  $r_1$  is in contact with region

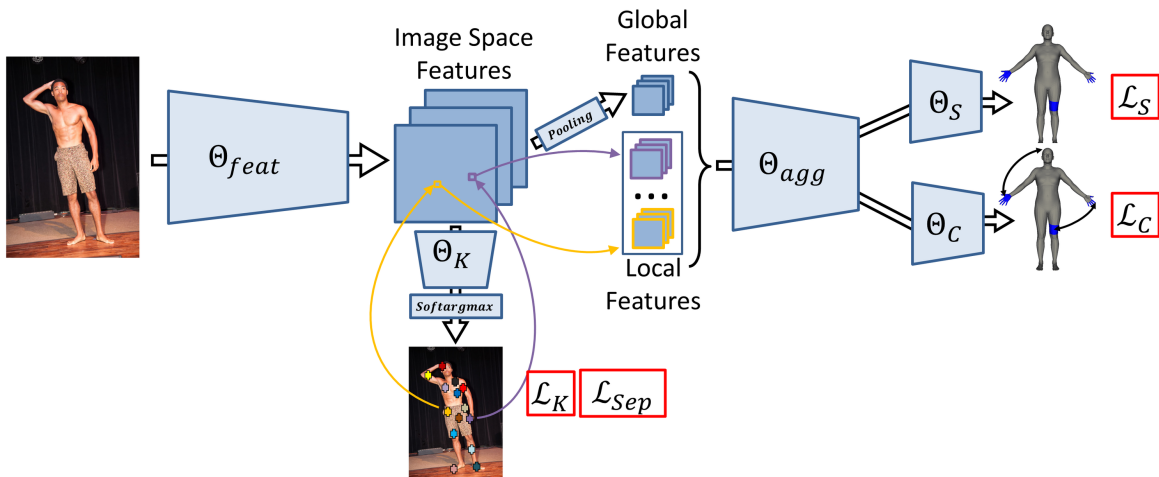


Figure 2: Our *SCP* architecture that estimates self-contact spatial support  $K$ , supervised by both  $\mathcal{L}_K$  (eq. 2) and  $\mathcal{L}_{sep}$  (eq. 3), and self-contact segmentation  $S$  and signature  $C$ , supervised by losses  $\mathcal{L}_S$  and  $\mathcal{L}_C$ . The input is an RGB image cropped around a person. *SCP* predicts the spatial support of self-contact  $K$  with  $\Theta_K$  and uses it to select local features (one for each body region). Merged with global features, these are processed by an aggregation layer  $\Theta_{agg}$  and specialization layers, for segmentation  $\Theta_S$  and signature prediction  $\Theta_C$ .

$r_2$ , where  $r_1 \neq r_2$  are surface regions of the same person and  $C_{r_1, r_2}^R(I) = 0$  otherwise. Note that  $C^R(I)$  is a symmetric matrix. Similarly, region-level self-contact **segmentation**  $S^R(I) \in \{0, 1\}^{N_R}$  is defined as  $S_r^R(I) = 1$  when  $r$  is in contact with any other surface region on the same body and  $S_r^R(I) = 0$  otherwise. In addition, we introduce the notion of **image support** of a region self-contact:

$$K^R(I) = \{(x_r, y_r) | S_r^R(I) = 1\} \quad (1)$$

where  $[x_r, y_r]$  is the coordinate of the center of region  $r$  projected in the image.

An overview of *SCP* is illustrated in fig. 2. *SCP* takes as input an RGB image cropped around a person and learns to extract image space features  $\Theta_{feat}$  using the backbone of the ResNet50 (He et al. 2016) architecture, up to the 16th convolutional layer.

### Self-Contact Image Support

One way in which the image support of self-contact can be leveraged is by informing the selection of local features needed for downstream tasks. To this end, after the feature encoder  $\Theta_{feat}$ , we extract a set of  $N_R$  heatmaps using  $\Theta_K$  and apply the softargmax operation to obtain a set of  $N_R$  image coordinates  $\{(\tilde{x}_r, \tilde{y}_r) | r \in \{1, \dots, N_R\}\}$ .

For regions that belong to the image support of self-contact ground truth  $K^R(I)$ , we apply the Euclidean loss  $\mathcal{L}_K$  to guide the discovered landmarks towards the ground-truth image support. Note that corresponding surface regions in self-contact have the same spatial support.

$$\mathcal{L}_K = \frac{1}{|K^R(I)|} \sum_{\substack{(x_r, y_r) \\ \in K^R(I)}} \|(x_r, y_r) - (\tilde{x}_r, \tilde{y}_r)\|_2^2 \quad (2)$$

For pairs of regions not in self-contact, we impose a separation constraint  $\mathcal{L}_{sep}$  to guide them towards different image

areas. We adopt the loss function proposed in (Zhang et al. 2018), that highly penalizes small distances between points and vanishes quickly as distance between them increases.

$$\mathcal{L}_{sep} = \sum_{\substack{(r_1, r_2) \\ C_{r_1, r_2}^R(I) = 0}} \exp\left(\frac{\|(\tilde{x}_{r_1}, \tilde{y}_{r_1}) - (\tilde{x}_{r_2}, \tilde{y}_{r_2})\|_2^2}{-2\sigma_{sep}^2}\right) \quad (3)$$

One can observe that  $\mathcal{L}_{sep}$  is a weakly-supervised loss, since it does not require ground truth support  $K^R(I)$ , but only ground truth self-contact signatures  $C^R(I)$ .

### Self-Contact Segmentation and Signature

The network can use the  $N_R$  landmarks as putative locations where local image features can be extracted. For each region  $r$ , we sample the image space features at location  $(\tilde{x}_r, \tilde{y}_r)$ . Since  $(\tilde{x}_r, \tilde{y}_r)$  are continuous, we bilinearly interpolate between features at the 4 nearby discrete coordinates on the image space grid. Each of the  $N_R$  local features are then concatenated with global features (obtained by a holistic pooling operation on the image space features) and fed to an aggregation module  $\Theta_{agg}$ . This is implemented as a series of two fully connected layers that progressively reduce the dimensionality of the  $N_R$  features.

For self-contact segmentation and signature tasks, we draw inspiration from the two specialization layers and the underlying losses introduced in (Fieraru et al. 2020).  $\Theta_S$  and  $\Theta_C$  are fully connected layers, and  $\mathcal{L}_S, \mathcal{L}_C$  are sigmoid cross-entropy losses, with the positive and negative classes appropriately weighted. While  $\mathcal{L}_S$  is applied directly on the output of  $\Theta_S$ , for the case of correspondences, the loss  $\mathcal{L}_C$  is applied on the feature similarity matrix  $FF^T$ , where  $F$  is the output of  $\Theta_C$ . We also experiment with the Euclidean distance as an alternative similarity metric to the dot prod-

uct used in  $FF^T$ , but confirm experimentally that it does not outperform it.

At inference, we propose, as novelty, to use both the estimated self-contact segmentation and the self-contact image support in order to limit erroneous correspondences in the predicted self-contact signature, as this has a large output space and is difficult to learn. First, we remove all correspondences involving regions not found in the predicted segmentation. Second, we remove all predicted correspondences between two regions whose estimated landmarks are not close to each other. This enforces consistency of signature with the image support, since two regions in correspondence should also have their image projection in proximity.

### Self-Contact Signatures for 3D Reconstruction

Self-contact signatures are also used to constrain 3d human reconstruction to be consistent. We showcase this using the optimization framework of (Zanfir, Marinoiu, and Sminchisescu 2018) augmented in (Fieraru et al. 2020) with interaction contact signature losses. The cost function, adapted for the reconstruction of a single person, and by using self-contact consistency, becomes:

$$L = L_S + L_{psr} + L_{col} + L_G \quad (4)$$

where  $L_S$  is the projection error with respect to estimated semantic body part labeling and 2d body pose,  $L_{psr}$  is a pose and shape regularization cost, and  $L_{col}$  is a self-collision penalty term.  $L_G = L_D + L_N$  is adapted to be a contact consistency cost for self-contact signature, where  $L_D$  minimizes the distance between pairs of regions in self-contact and  $L_N$  is a term aligning the orientation of region surfaces found in self-contact.

Please check our Sup. Mat. for further details on both the SCP network and the optimization framework.

### Proposed Datasets

**HumanSC3D.** As current 3d human pose datasets such as Human3.6m (Ionescu et al. 2014) or 3DPW (von Marcard et al. 2018) contain relatively few frames in self-contact, to evaluate our proposed methodology, we collect a new dataset of people in more challenging self-contact poses.

HumanSC3D contains 3d motion capture data of 6 human participants (3 men and 3 women between 20 and 30 years old with various fitness levels and body shapes), and videos captured by 4 synchronized RGB cameras. The subjects are shown a series of images with ordinary poses and self-contact and asked to reproduce only the type of contact they see (not the pose), such as: touching one’s chin, crossing one’s legs or arms, etc. In addition, they are instructed to continuously change their body position and orientation relative to the cameras for increased variability. For each scenario, we record a short clip that captures the transition from an A-pose to the desired self-contact and back to the A-pose. In total, each subject performs 172 motions, out of which 116 are standing, 20 sitting on the floor and 36 sitting on or standing next to a chair, summing up to a total of 1,246,487 ground truth 3d poses and associated RGB images.

We also capture each subject’s body shape using a 3d body scanner. By fitting our body model to the body scans

for shape and to the 3d marker positions for pose, we also obtain (pseudo) ground truth reconstructions (Xu et al. 2020).

For each of the 172 self-contact scenarios of any given subject we extract a middle frame (where the person is in self-contact) and manually label the body regions in contact and their correspondence. The annotation is performed by clicking on the surface of a 3d human body model with  $\sim 10k$  vertices. In addition, the annotators are also asked to roughly indicate the spatial support of the self-contact in the image, by clicking in the original image at the position of each self-contact. In this way, we obtain 4,128 images with people and associated self-contact information (multiple correspondences between two facets of the mesh and pixels of the image). Although in a controlled environment, we choose to manually annotate self-contact for higher quality control. Alternative approaches, such as multi-view marker-based reconstruction, can still fail under complex self-contact, especially for the hand regions where markers are sparsely placed or occluded.



Figure 3: Examples of images from our FlickrSC3D (top) and HumanSC3D (bottom) dataset. People in self-contact (left). People not in self-contact (center). Uncertain whether the person is in self-contact or not (right).

**FlickrSC3D.** To further extend our experiments to natural settings, we gather 3,969 images under the CC-BY license from Flickr, containing persons in self-contact. To obtain this data, we first crawl images with people by choosing a wide variety of tags (from daily activities to dance or sports) and run a person detector on the selection. Then, we pick images with persons in self-contact by manually classifying each person’s bounding box in one of the 3 classes: ”contact”, ”no contact” and ”uncertain contact” (see fig. 3). To ensure pose variability among images with persons in self-contact we additionally run a 3d pose estimator (Zanfir et al. 2018) and greedily select images that have a large 3d pose distance compared to the ones already selected. For the final pool of images we annotate the self-contact signature and image-space support of the signature, in a similar way to the HumanSC3D dataset. Statistics regarding the self-contact on the in-the-wild FlickrSC3D dataset can be visualized in fig. 4.

**Annotator consistency.** For a small subset of images from both datasets, we ask two raters to annotate the self-contact



Num. Reg.	Segmentation IoU		Signature IoU	
	HumanSC3D	FlickrSC3D	HumanSC3D	FlickrSC3D
75	0.469	0.528	0.315	0.422
37	0.560	0.564	0.512	0.475
17	0.703	0.664	0.590	0.579
9	0.787	0.768	0.685	0.692

Table 1: Annotator consistency, function of reg. granularity.

signature. We check the annotator consistency at different levels of granularity of the body regions (from a fine-grained split into 75 regions to a coarser one, of only 9 body regions). We measure the intersection over union (IoU), first at a body region segmentation level, and then also by taking into account the set of correspondences between regions. Results are shown in table 1. As in many other tasks, human annotation is not perfect, but it can be noticed that at smaller levels of granularity consistency increases and is of practical use (this certainly improves the quality of 3d reconstructions, as shown qualitatively and quantitatively in the following section).

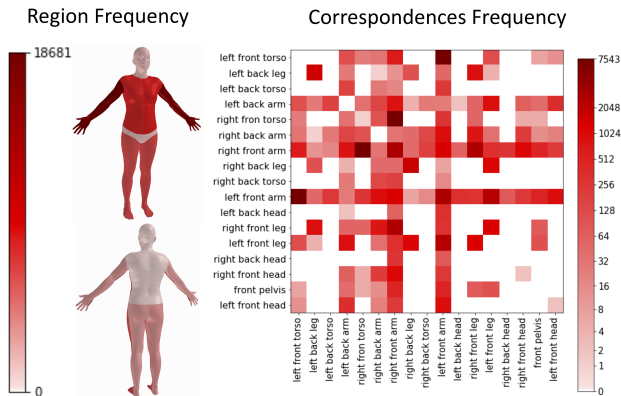


Figure 4: Body region frequency of self-contact (75 regions) (left). Note the left-right symmetry and the high frequency for the arms, hands, legs and torso regions. Self-contact correspondence counts (17 regions) (right).

## Experiments

**Self-Contact Image Support, Segmentation and Signature.** To assess the performance of our *SCP* network, we validate and test it on the FlickrSC3D dataset, which we split in the usual train (80%), test (10%) and validation (10%) subsets. We train using the training set, validate our meta-parameters on the validation set and show the quantitative analysis on the test set.

Since we are the first to propose explicitly learning the self-contact of the human body, there is no available method to compare against. The closest work in the literature is the ISP network (Fieraru et al. 2020) which predicts the contact signature between two humans in contact, which we adapt to learn a self-contact signature. We achieve it by removing one of the two computational pathways of ISP (the



Figure 5: GT image self-contact support (left). Intermediate contact image support of regions predicted to be in contact (within the segmentation prediction) (center). Final estimated contact image support of regions found to be in correspondence and having nearby spatial support (right).

graph convolutional layers for one person and its respective specialization tasks for segmentation and signature learning) and then train and validate on the self-contact dataset FlickrSC3D. We train all methods on the finest granularity available ( $N_R = 75$ ), but also report results on multiple coarser granularities ( $N_R = 37, 17$  and 9), following the region splitting used in ISP for comparison.

On both datasets, annotators have the freedom to choose whichever correspondences between facets of the human model and the image they prefer (as long as they are valid and the region-level self-contact segmentation is complete). This can lead to different multiple clicks in the image support of the same region. We set the ground truth image support of the respective region containing multiple coordinates as their average. Since the signature annotation is not necessarily complete (either because some regions are flagged *masked*, when it is unclear in the image if they are involved in self-contact or not, or because just a subset of correspondences is annotated), we neither penalize them in training nor in evaluation.

Table 2 shows quantitative results in terms of the intersection over union metric  $\text{IoU}_{N_R}$ , computed for different  $N_R$ . We show results for our method *SCP* and ablate different components. The strength of the proposed method is best seen at the finest level of granularity  $N_R = 75$ , where it more than doubles the signature prediction performance, at 0.301 vs. 0.133 for ISP. The scenario where no additional supervision is used (for the self-contact image support) and the landmark discovery is unconstrained (*SCP w/o  $\mathcal{L}_{sep}$  and  $\mathcal{L}_K$* ) also outperforms ISP, showing that the effectiveness of our method does not only stem from the newly proposed self-contact image support annotations, but also from the more robust and specialized architecture. The effectiveness of the separability constraint  $\mathcal{L}_{sep}$  is also shown for the self-contact signature task. Moreover, the experiment where the self-contact signature is not constrained to be consistent with the image support (0.244  $\text{IoU}_{75}$ ) shows the crucial, positive effect, of spotting and eliminating spurious correspondences by using the estimated image support. Fig. 5 visualizes the image support estimates after applying different constraints.

**Hand-Face Self-Contact** To assess the possible application to analyzing disease transmission, we evaluate the hand-to-face self-contact detection. On our collected HumanSC3D, the hand-face correspondence is present in 34% of images. On the problem of hand-to-face detection, *SCP* trained for general self-contact prediction obtains 46% recall and 75% precision. When retraining *SCP* with losses penalizing

Method	IoU <sub>75</sub>		IoU <sub>37</sub>		IoU <sub>17</sub>		IoU <sub>9</sub>	
	Segm.	Signature	Segm.	Signature	Segm.	Signature	Segm.	Signature
<i>SCP</i>	<b>0.469</b>	<b>0.301</b>	0.507	<b>0.339</b>	0.591	<b>0.442</b>	0.693	<b>0.550</b>
<i>SCP</i> w/o $\mathcal{L}_{sep}$	0.465	0.289	<b>0.510</b>	0.335	0.603	0.428	0.692	0.536
<i>SCP</i> w/o $\mathcal{L}_{sep}$ , w/o $\mathcal{L}_K$	0.460	0.236	0.502	0.283	<b>0.605</b>	0.395	<b>0.708</b>	0.514
<i>SCP</i> w/o imposing signature consistency with image support	0.469	0.244	0.507	0.288	0.591	0.395	0.693	0.501
ISP (Fieraru et al. 2020) (adapted for self-contact)	0.462	0.133	0.503	0.186	0.583	0.305	0.688	0.460
Human Consistency	0.528	0.422	0.564	0.475	0.664	0.579	0.768	0.692

Table 2: Results of our self-contact segmentation and signature estimation on FlickrSC3D, evaluated for different region granularities on the human 3d surface (from 75, down to 9 regions). Human consistency is the same as in table 1. We ablate the proposed losses and compare with the ISP baseline.




Figure 6: Pose error for a subset of HumanSC3D, each annotated by 2 human raters. We reconstruct the humans by enforcing contact consistency using the correspondences set by each annotator (*green* and *blue* line) and also without enforcing contact consistency (*red* line, image IDs are ordered following this error). Enforcing contact consistency leads to smaller reconstruction error in 71.5% cases (whether using correspondences from Ann. 1 or from Ann. 2), with 74% agreement over the effect of enforcing contact consistency (either annotations improving or deteriorating the reconstruction simultaneously). The higher the initial pose error, the higher the improvement when enforcing self-contact consistency.

only the hand-to-face self-contact, the detection improves, to 53% recall and 76% precision. These are both major improvements from the 10% recall and 66% precision obtained by the adapted ISP baseline. Still, the current results show that self-contact prediction is not a trivial problem, (see our Sup. Mat. video for failure cases), and can benefit from future research, potentially based on some of the methodology and data we provide.

**Self-Contact Signature for 3D Reconstruction** To quantitatively assess the impact of self-contact consistency constraints in the quality of reconstructions, we test our method on the HumanSC3D dataset, where we report the MPJPE (mean per-joint position error) to evaluate the inferred pose, the translation error, the contact distance error (the minimum Euclidean distance between each pair of facets from two regions annotated to be in self-contact correspondence), and the per-vertex Euclidean distance error, measured against the (pseudo) ground truth meshes. Table 3 shows improvement across all the metrics when annotated self-contact signature is used to further constrain the reconstruction, showing that our annotations are valuable. Fig. 6 plots the pose error for correspondences from two different annotators. While self-contact constraints do not always yield better reconstructions, on average, they do.

Fig. 7 shows reconstruction results for images in FlickrSC3D, both with and without self-contact consistency constraints. By adding the penalty on annotated regions in correspondence, we recover accurate and visually plausible 3d reconstructions of challenging human poses.

## Conclusions

We have presented the task of human self-contact estimation and the design of the *SCP* methodology to detect body surface regions in self-contact, the correspondences between them, and their spatial support. By integrating this methodology with 3d explicit self-contact losses, we have shown that 3d visual reconstructions of human self-touch events are possible with superior quantitative and perceptual results over non-contact baselines. The models we built had their component effectiveness evaluated based on a large dataset collected in the wild, containing 25,297 image-surface correspondence annotations, as well as a motion capture dataset containing 5,058 contact events and 1,246,487 ground truth 3d poses. This represents a considerable amount of logistic, collection and annotation work, involving human subjects, and will be made available to the research community.<sup>1</sup> Finally, we have demonstrated an application to detecting face-touch and showed how self-contact signatures can enable more expressive 3d reconstruction, thus opening path for subtle 3d behavioral reasoning in the future.

## Acknowledgments

This work was supported in part by the ERC Consolidator grant SEED, CNCS-UEFISCDI (PN-III-P4-ID-PCCF2016-0180) and SSF.

<sup>1</sup><http://vision.imar.ro/sc3d>

Optim. loss	W/o chair - standing				W/o chair - sitting				W/ chair				Overall			
	P	T	V	C	P	T	V	C	P	T	V	C	P	T	V	C
$L$	93.8	408.4	76.6	12.9	116.1	424.1	93.1	26.6	107.2	426.2	84.6	23.7	98.2	414.2	80.2	16.4
$L$ w/o $L_G$	106.0	419.3	121.0	210.2	145.2	436.0	147.4	182.7	131.6	431.9	122.7	189.3	114.4	423.4	124.4	203.4

Table 3: 3D human pose (P), translation (T), vertex (V) estimation errors, as well as mean 3d contact distance (C), expressed in mm, for the HumanSC3D dataset. Using the full optimization function, with the geometric alignment term on annotated self-contact signatures, decreases the pose, translation and vertex estimation errors as well as the 3d distance between surfaces annotated as being in contact.




Figure 7: 3D pose and shape reconstructions using our annotated self-contact data. Original image (left). Reconstruction without considering the self-contact and the associated loss (center). Reconstruction that uses the self-contact annotations and the corresponding loss (right).

## Broader Impact

While research is still experimental, in the long run our models can be integrated into systems performing large-scale psycho-social studies of human behavior. Such models can also be used with personal assistants that can operate under high privacy standards and can be accountable to humans. Assistants can rely on self-contact signatures to reason about a person's internal state including emotional response, and could provide feedback to that person over a period of time, for increase awareness or for positively changing habits. Our work can also be potentially relevant to detect and correct the unconscious behaviour of touching one's mouth or face with the hand (Kwok, Gralton, and McLaws 2015), in order to avoid infections with various pathogens, if proper hygiene is not maintained. In this regard, the models can be potentially used in hospitals in order to monitor hygiene of both patients and medical personnel. The work can also be potentially applied in the monitoring and treatment of people with a history of self-harm (Hawton et al. 2015). During data collection, we aimed to reduce bias, by having a diverse and representative collection of humans in self-contact, within our limited subject and annotation budget.

## References

- Benzine, A.; Luvison, B.; Pham, Q. C.; and Achard, C. 2019. Deep, Robust and Single Shot 3D Multi-Person Human Pose Estimation from Monocular Images. In *ICIP*.
- Bogo, F.; Kanazawa, A.; Lassner, C.; Gehler, P.; Romero, J.; and Black, M. J. 2016. Keep it SMPL: Automatic Estimation of 3D Human Pose and Shape from a Single Image. In *ECCV*.
- Fieraru, M.; Zanfir, M.; Oneata, E.; Popa, A.-I.; Olaru, V.; and Sminchisescu, C. 2020. Three-dimensional Reconstruction of Human Interactions. In *CVPR*.
- Hassan, M.; Choutas, V.; Tzionas, D.; and Black, M. J. 2019. Resolving 3D Human Pose Ambiguities with 3D Scene Constraints. In *ICCV*. URL <https://prox.is.tue.mpg.de>.
- Hasson, Y.; Varol, G.; Tzionas, D.; Kalevatykh, I.; Black, M. J.; Laptev, I.; and Schmid, C. 2019. Learning Joint Reconstruction of Hands and Manipulated Objects. In *CVPR*.
- Hawton, K.; Haw, C.; Casey, D.; Bale, L.; Brand, F.; and Rutherford, D. 2015. Self-harm in Oxford, England: epidemiological and clinical trends, 1996–2010. *Social psychiatry and psychiatric epidemiology* 50(5): 695–704.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2014. Human3.6M: Large Scale Datasets and Predictive Methods for 3D Human Sensing in Natural Environments. *PAMI*.
- Kanazawa, A.; Zhang, J. Y.; Felsen, P.; and Malik, J. 2019. Learning 3D Human Dynamics from Video. In *CVPR*.
- Kocabas, M.; Athanasiou, N.; and Black, M. J. 2020. VIBE: Video Inference for Human Body Pose and Shape Estimation. In *CVPR*.
- Kolotouros, N.; Pavlakos, G.; Black, M. J.; and Daniilidis, K. 2019. Learning to Reconstruct 3D Human Pose and Shape via Model-fitting in the Loop. In *ICCV*.
- Kwok, Y. L. A.; Gralton, J.; and McLaws, M.-L. 2015. Face touching: A frequent habit that has implications for hand hygiene. *American journal of infection control* 43(2): 112–114.
- Li, J.; Wang, C.; Zhu, H.; Mao, Y.; Fang, H.-S.; and Lu, C. 2019a. Crowdpose: Efficient crowded scenes pose estimation and a new benchmark. In *CVPR*.
- Li, Z.; Sedlar, J.; Carpentier, J.; Laptev, I.; Mansard, N.; and Sivic, J. 2019b. Estimating 3d motion and forces of person-object interactions from monocular video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8640–8649.
- Liu, Y.; Gall, J.; Stoll, C.; Dai, Q.; Seidel, H.-P.; and Theobalt, C. 2013. Markerless Motion Capture of Multiple Characters Using Multi-view Image Segmentation. *PAMI*.
- Liu, Y.; Stoll, C.; Gall, J.; Seidel, H.-P.; and Theobalt, C. 2011. Markerless motion capture of interacting characters using multi-view image segmentation. In *CVPR*.
- Manoj, C.; Magesh, S.; Sankaran, A. S.; and Manikandan, M. S. 2011. Novel approach for detecting applause in continuous meeting speech. In *2011 3rd International Conference on Electronics Computer Technology*, volume 3, 182–186.
- Marinoiu, E.; Zanfir, M.; Olaru, V.; and Sminchisescu, C. 2018. 3D Human Sensing, Action and Emotion Recognition in Robot Assisted Therapy of Children with Autism. In *CVPR*.
- Mehta, D.; Sotnychenko, O.; Mueller, F.; Xu, W.; Sridhar, S.; Pons-Moll, G.; and Theobalt, C. 2018. Single-Shot Multi-person 3D Pose Estimation from Monocular RGB. In *3DV*.
- Mueller, F.; Davis, M.; Bernard, F.; Sotnychenko, O.; Verschoor, M.; Otaduy, M. A.; Casas, D.; and Theobalt, C. 2019. Real-time Pose and Shape Reconstruction of Two Interacting Hands With a Single Depth Camera. *ACM Transactions on Graphics (TOG)* 38(4).
- Mueller, S. M.; Martin, S.; and Grunwald, M. 2019. Self-touch: Contact durations and point of touch of spontaneous facial self-touches differ depending on cognitive and emotional load. *PLoS one* 14(3): e0213677.
- Pavlakos, G.; Choutas, V.; Ghorbani, N.; Bolkart, T.; Osman, A. A.; Tzionas, D.; and Black, M. J. 2019. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 10975–10985.
- Rudovic, O.; Lee, J.; Mascarell-Maricic, L.; Schuller, B. W.; and Picard, R. W. 2017. Measuring engagement in robot-assisted autism therapy: A cross-cultural study. *Frontiers in Robotics and AI*.
- Su, K.; Yu, D.; Xu, Z.; Geng, X.; and Wang, C. 2019. Multi-Person Pose Estimation with Enhanced Channel-wise and Spatial Information. In *CVPR*.



Taylor, J.; Tankovich, V.; Tang, D.; Keskin, C.; Kim, D.; Davidson, P.; Kowdle, A.; and Izadi, S. 2017. Articulated Distance Fields for Ultra-fast Tracking of Hands Interacting. *ACM Trans. Graph.* 36(6): 244:1–244:12. ISSN 0730-0301. doi:10.1145/3130800.3130853. URL <http://doi.acm.org/10.1145/3130800.3130853>.

Tzionas, D.; Ballan, L.; Srikantha, A.; Aponte, P.; Pollefeys, M.; and Gall, J. 2016. Capturing Hands in Action Using Discriminative Salient Points and Physics Simulation. *IJCV*.

Tzionas, D.; and Gall, J. 2013. A Comparison of Directional Distances for Hand Pose Estimation. In *German Conference on Pattern Recognition (GCPR)*, volume 8142 of *Lecture Notes in Computer Science*, 131–141. Springer. URL [http://dx.doi.org/10.1007/978-3-642-40602-7\\_14](http://dx.doi.org/10.1007/978-3-642-40602-7_14).

von Marcard, T.; Henschel, R.; Black, M.; Rosenhahn, B.; and Pons-Moll, G. 2018. Recovering Accurate 3D Human Pose in The Wild Using IMUs and a Moving Camera. In *ECCV*.

Xu, H.; Bazavan, E. G.; Zanfır, A.; Freeman, W. T.; Sukthankar, R.; and Sminchisescu, C. 2020. GHUM & GHUML: Generative 3D Human Shape and Articulated Pose Models. In *CVPR*.

Zanfır, A.; Marinoiu, E.; and Sminchisescu, C. 2018. Monocular 3D Pose and Shape Estimation of Multiple People in Natural Scenes - The Importance of Multiple Scene Constraints. In *CVPR*.

Zanfır, A.; Marinoiu, E.; Zanfır, M.; Popa, A.-I.; and Sminchisescu, C. 2018. Deep Network for the Integrated 3D Sensing of Multiple People in Natural Images. In *NIPS*.

Zhang, Y.; Guo, Y.; Jin, Y.; Luo, Y.; He, Z.; and Lee, H. 2018. Unsupervised discovery of object landmarks as structural representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2694–2703.

Zhang, Y.; Hassan, M.; Neumann, H.; Black, M. J.; and Tang, S. 2020. Generating 3D People in Scenes without People. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6194–6204.

Zou, Y.; Yang, J.; Ceylan, D.; Zhang, J.; Perazzi, F.; and Huang, J.-B. 2020. Reducing Footskate in Human Motion Reconstruction with Ground Contact Constraints. In *The IEEE Winter Conference on Applications of Computer Vision*, 459–468.