

Im2Hands: Learning Attentive Implicit Representation of Interacting Two-Hand Shapes

Jihyun Lee¹ Minhyuk Sung¹ Honggyu Choi¹ Tae-Kyun Kim^{1,2}
¹ KAIST ² Imperial College London




Figure 1. **Reconstructed two-hand shapes from an RGB image.** We present Im2Hands, the first neural implicit representation for two interacting hands. Compared to the existing mesh-based two-hand reconstruction method (IntagHand [23]), Im2Hands effectively captures the fine-grained geometry of two hands with higher shape-to-image coherency. The above results were produced from a single RGB image input, where ours utilized an off-the-shelf two-hand keypoint estimation method (DIGIT [11]) to condition the articulated occupancy.

Abstract

We present **Implicit Two Hands (Im2Hands)**, the first neural implicit representation of two interacting hands. Unlike existing methods on two-hand reconstruction that rely on a parametric hand model and/or low-resolution meshes, Im2Hands can produce fine-grained geometry of two hands with high hand-to-hand and hand-to-image coherency. To handle the shape complexity and interaction context between two hands, Im2Hands models the occupancy volume of two hands – conditioned on an RGB image and coarse 3D keypoints – by two novel attention-based modules responsible for (1) initial occupancy estimation and (2) context-aware occupancy refinement, respectively. Im2Hands first learns per-hand neural articulated occupancy in the canonical space designed for each hand using query-image attention. It then refines the initial two-hand occupancy in the posed space to enhance the coherency between the two hand shapes using query-anchor attention. In addition, we introduce an optional keypoint refinement module to enable robust two-hand shape estimation from predicted hand keypoints in a single-image reconstruction scenario. We experimentally demonstrate the effectiveness of Im2Hands on two-hand reconstruction in comparison to related methods, where ours achieves state-of-the-art results. Our code is publicly available at <https://github.com/jyunlee/Im2Hands>.

1. Introduction

Humans use hand-to-hand interaction in everyday activities, which makes modeling 3D shapes of two interacting hands important for various applications (e.g., human-computer interaction, robotics, and augmented or virtual reality). However, the domain of two-hand shape reconstruction remains relatively under-explored, while many existing studies have put efforts into *single-hand* reconstruction from RGB [1, 2, 4, 13, 21, 24, 44], depth [38], or sparse keypoints [18, 44]. These single hand-based methods are not effective when directly applied for interacting two-hand reconstruction, since it introduces additional challenges including inter-hand collisions and mutual occlusions.

Recently, few learning-based methods on two-hand shape reconstruction [23, 32, 42] have been proposed since the release of the large-scale interacting hand dataset (i.e., InterHand2.6M [27]). Two-Hand-Shape-Pose [42] and IHMR [32] reconstruct two-hands by estimating MANO [31] parameters, which are later mapped to triangular hand meshes using a pre-defined statistical model (i.e., MANO). IntagHand [23] directly regresses a fixed number of mesh vertex coordinates using a graph convolutional network (GCN). These methods mainly model the shape of two interacting hands based on a low-resolution mesh representation with a fixed topology of MANO (please refer to Figure 1).

In this paper, we present **Implicit Two Hands** (Im2Hands), the first neural implicit representation of two interacting hands. Unlike the existing mesh-based two-hand reconstruction methods, Im2Hands can capture the fine-grained geometry of two interacting hands by learning a *continuous* 3D occupancy field. Im2Hands (1) produces two-hand meshes with an arbitrary resolution, (2) does not require dense vertex correspondences or statistical model parameter annotations for training, and (3) learns output shapes with precise hand-to-hand and hand-to-image alignment. As two interacting hands are highly articulated objects, we take inspiration from recent neural *articulated* implicit functions [7, 8, 18, 25, 29, 34] that learn an implicit geometry leveraging the object canonical space computed from an input pose observation. Our two-hand articulated implicit function is also driven by input pose and shape observations, which are represented as sparse 3D keypoints and an RGB image, respectively.

To effectively handle the shape complexity and interaction context between two hands, Im2Hands consists of two novel attention-based modules responsible for initial hand occupancy estimation and context-aware two-hand occupancy refinement, respectively. The initial occupancy estimation network first predicts the articulated occupancy volume of each hand in the canonical space. Given a 3D query point, it (1) performs query canonicalization using the keypoint encoder of HALO [18] to effectively capture *pose-dependent hand deformations* and (2) extracts a hand shape feature using our novel query-image attention module to model *shape-dependent hand deformations*. As it is non-trivial to model two-hand interaction while learning in the canonical space defined for each hand, our context-aware occupancy refinement network modifies the initial two-hand occupancy in the original posed space to enhance hand-to-hand coherency. Given the initial two-hand shape represented as anchored point clouds, it uses query-anchor attention to learn a refined two-hand occupancy in a context-aware manner. Furthermore, we consider a practical scenario of two-hand reconstruction using Im2Hands from single images, where no ground truth keypoints are observed as inputs to our method. To this end, we introduce an *optional* input keypoint refinement network to enable more robust two-hand shape reconstruction by alleviating errors in the input 3D keypoints predicted from an off-the-shelf image-based two-hand keypoint estimation method (e.g., [11, 20, 23, 27, 42]).

Overall, our main contributions are summarized as follows:

- We introduce Im2Hands, the first neural implicit representation of two interacting hands. Im2Hands reconstructs resolution-independent geometry of two-hands with high hand-to-hand and hand-to-image coherency.

- To effectively learn an occupancy field of the complex two-hand geometries, we propose two novel attention-based modules that perform (1) initial occupancy estimation in the canonical space and (2) context-aware occupancy refinement in the original posed space, respectively. We additionally introduce an *optional* keypoint refinement module to enable more robust two-hand shape estimation using a single image input.
- We demonstrate the effectiveness of Im2Hands in comparison to the existing (1) two-hand mesh-based and (2) single-hand implicit function-based reconstruction methods, where Im2Hands achieves state-of-the-art results in interacting two-hand reconstruction.

2. Related Work

Single-Hand Reconstruction. Methods for single-hand reconstruction have been actively investigated in the past decades. Most existing deep learning-based approaches either reconstruct hand poses represented as 3D keypoints [5, 12, 17, 26, 35, 36, 40, 45], estimate MANO parameters [1, 2, 31], or directly regress mesh vertex coordinates [13, 21, 24, 38]. Inspired by the recent success of implicit representations in modeling human bodies [8, 25, 29, 34], few recent works [7, 18] also employ neural implicit functions for single-hand reconstruction. Compared to the existing methods based on mesh representations, these implicit function-based methods [7, 18] can produce fine-grained hand geometry in a resolution-independent manner.

Interacting Hand-Object Reconstruction. Similar to single-hand reconstruction, most of the existing methods [2, 6, 9, 15, 16] on hand-object reconstruction adopt MANO topology-based mesh representations to model a hand shape. Few recent works [18, 19] consider neural implicit representations to model hand-objects, but they often constrain their interaction based on contacts – which does not *necessarily* occur in hand-to-hand interactions [42]. In addition, these methods mainly consider a rigid object in the interaction, which makes them ineffective to be directly applied for articulated two-hand reconstruction.

Interacting Two-Hand Reconstruction. Compared to single-hand or hand-object reconstruction, two-hand reconstruction is more challenging, as more complex occlusions and deformations occur from the interaction between two articulated hands. For modeling two interacting hands, there have been methods recently proposed for two-hand *pose* reconstruction from RGB images [11, 20, 23, 27, 42], depth images [28, 37], or RGB-D video sequences [22, 30]. Yet, there are only a few methods that can directly reconstruct the *dense surface* of closely interacting two-hands [23, 28, 32, 42], which is more challenging and important in reasoning hand-to-hand interactions. Mueller *et al.* [28] propose an energy minimization framework to

fit MANO [31] parameters to the input depth image of two interacting hands. Two-Hand-Shape-Pose [42] introduce pose-aware attention and context-aware cascaded refinement modules to directly regress two-hand MANO parameters from an RGB image. IntagHand [23] propose an attention-based graph convolutional network (GCN) for two-hand vertex regression from an RGB image. These recent deep learning-based frameworks [23, 32, 39, 42] have shown that (1) an attention mechanism to model non-local interactions and (2) context-aware shape refinement steps are effective in two-hand reconstruction, which have inspired the design of our two-hand occupancy function. However, compared to these existing methods, our occupancy-based method can learn resolution-independent hand surface with better image-shape alignment, as our output space (i.e., occupancy field) itself is more directly aligned with the input image space.

Neural Articulated Implicit Representation. Many existing methods to model articulated objects (e.g., human bodies, hands) use neural *articulated* implicit representations [7, 8, 18, 25, 29, 34], which directly condition an implicit geometry on the object articulation. While many methods [8, 25, 29, 34] mainly demonstrate their results on human body modeling, there are only a few studies that explore the effectiveness of implicit representations for hands. LISA [7] proposes an articulated VolSDF [41] to model shape and appearance of single-hands, but it assumes that various input domains (i.e., multi-view RGB image sequences, 3D bone transformations and foreground masks) are available and requires model optimization for inference. HALO [18] proposes a single-hand articulated occupancy function driven by 3D keypoints via incorporating a novel canonicalization layer – eliminating the need for the ground truth 3D bone transformations used in most of the articulated implicit functions [7, 8, 25, 29, 34]. However, HALO models pose-independent shape variations only with hand bone lengths, which may be a strong assumption. We take inspiration from HALO to condition our two-hand occupancy on the pose represented as sparse 3D keypoints, but we further improve HALO by conditioning shape variations using an RGB image, which is a more direct form of shape observation than bone lengths.

3. Im2Hands: Implicit Two-Hand Function

Im2Hands is a neural occupancy representation of two interacting hands. Our two-hand occupancy network can be formally defined as:

$$\mathcal{O}(x | \alpha, \beta) \rightarrow [o_l, o_r], \quad (1)$$

where \mathcal{O} is a neural network with the learned weights. \mathcal{O} maps an input 3D query point $x \in \mathbb{R}^3$ to occupancy probabilities for each side of hand $o_l, o_r \in [0, 1]$ conditioned on

a shape observation α and a pose observation β , which are represented as an RGB image $I \in \mathbb{R}^{w \times h \times 3}$ and sparse 3D two-hand keypoints $J = \begin{bmatrix} J_l \\ J_r \end{bmatrix} \in \mathbb{R}^{42 \times 3}$, respectively.

One straightforward approach to design a neural implicit function for two hands would be directly applying the existing articulated occupancy function for *single hand* (i.e., HALO [18]) to both hands separately. This existing method can robustly model pose-dependent deformations via learning hand shapes in the canonical space, but it does not effectively capture other shape-dependent deformations (e.g., identity-dependent or soft tissue deformations). We thus design our initial hand occupancy estimation network (Section 3.1) that combines (1) HALO to robustly model *pose-dependent deformations* by learning in the hand canonical space and (2) our novel query-image attention module to capture *shape-dependent deformations* observed from an RGB image. However, it is difficult to handle two-hand interaction contexts while learning in the hand canonical spaces. Thus, we additionally propose a two-hand occupancy refinement network (Section 3.2) to perform context-aware shape refinement of interacting two hands in the original posed space. Furthermore, we consider two-hand reconstruction from single images, where the input two-hand keypoints can be obtained by applying an off-the-shelf two-hand pose estimation method (e.g., [11, 20, 23, 27, 42]). To this end, we additionally introduce an optional input keypoint refinement module (Section 3.3) to alleviate input keypoint errors to enable more robust shape estimation.

In what follows, we explain each of the proposed components of Im2Hands in more detail.

3.1. Initial Hand Occupancy Estimation

Given a 3D query point x , our initial hand occupancy network \mathcal{I} predicts initial occupancy probabilities for each hand $o_l^i, o_r^i \in [0, 1]$ conditioned on an RGB image I and two-hand keypoints J . Our network design is partly based on HALO [18], which models an implicit *single-hand* conditioned on 3D keypoints.

Background: HALO [18]. HALO aims to learn an occupancy field of a single hand driven by 3D keypoints. To learn hand shapes in the canonical space, HALO proposes a novel algorithm to transform the input 3D keypoints to canonicalization matrices $\{\mathbf{T}_b\}_{b=1}^B$, where \mathbf{T}_b is a transformation matrix to the canonical pose for hand bone b , and B is the number of hand bones. Using the obtained canonicalization matrices, hand occupancy at query x is modeled as:

$$\mathcal{H}(x | J) = \max_{b=1, \dots, B} \{\bar{\mathcal{H}}_b(\mathbf{T}_b x, f_b^\phi, f_b^\omega)\}, \quad (2)$$

where $\bar{\mathcal{H}}_b$ is an MLP-based part occupancy network to learn the shape of hand bone b . Each part occupancy network takes a canonicalized query point $\mathbf{T}_b x$ along with a bone

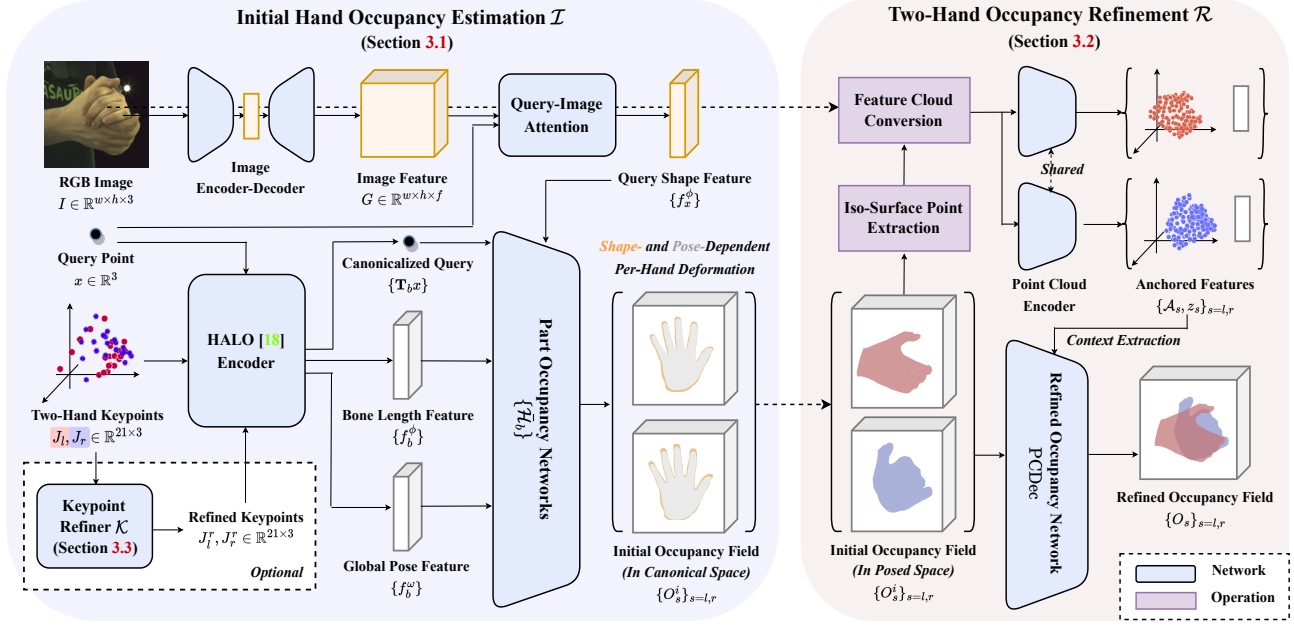


Figure 2. **Architecture overview.** Given an RGB image and coarse 3D two-hand keypoints, our method estimates two-hand occupancy volumes via (1) initial hand occupancy estimation and (2) two-hand occupancy refinement. In a nutshell, the initial hand occupancy network uses query-image attention and HALO [18] encoder to learn per-hand occupancy volumes in the canonical spaces. Then, the two-hand occupancy refinement network encodes the initial per-hand occupancy volumes into anchored features, which are then used to provide two-hand context information for refined occupancy estimation in the original posed space.

length shape feature f_b^ϕ and a pose feature f_b^ω for bone b . f_b^ϕ is extracted by an MLP encoder that takes a bone length vector $l \in \mathbb{R}^B$ computed from the input 3D keypoints. f_b^ω is obtained by another MLP encoder that takes global pose matrices $\{\mathbf{T}_b\}_{b=1}^B$ and a root translation vector $t \in \mathbb{R}^3$ as inputs. For more details about HALO, we kindly refer the reader to [18].

Our Initial Hand Occupancy Network. While HALO is effective in modeling pose-dependent deformations via query point canonicalization (i.e. $\mathbf{T}_b x$) and pose feature extraction (i.e., f_b^ω), its bone length shape feature (i.e., f_b^ϕ) cannot model shape-dependent deformations that cannot be described by bone lengths. Modeling such shape variations is especially important in interacting two-hand reconstruction, as soft tissue deformations commonly occur due to inter-hand contacts. Thus, we introduce an additional shape feature conditioned on an RGB image I , which provides a minimal observation for such shape-dependent deformations. We propose a per-query shape feature f_x^ϕ conditioned on the image I as follows:

$$f_x^\phi = \text{MSA}([\text{PosEnc}(x), \text{ImgEnc}(I)]), \quad (3)$$

where MSA denotes a multi-headed self-attention module that takes a query feature $\text{PosEnc}(x)$ and patch-wise image features $\text{ImgEnc}(I)$. To be more specific, PosEnc is an MLP that performs positional embedding to map the in-

put query coordinate x to a query feature vector $f_x \in \mathbb{R}^d$. ImgEnc is an image encoder of Vision Transformer [10], which maps the input image I into patch-wise image features $f_I \in \mathbb{R}^{p \times d}$, where p is the number of patches. Through this query-image attention, we can obtain a per-query shape feature f_x^ϕ contributed by the features of more relevant image patch regions to the input query x . In Section 4, we also show that using this query-image attention yields better results than directly using a local image feature located at the projected query position (e.g., PIFu [33]).

Finally, our initial per-hand occupancy is modeled by feeding our per-query shape feature f_x^ϕ along with the other inputs to the MLP-based part occupancy network $\bar{\mathcal{H}}_b$ in Equation 2:

$$\mathcal{I}(x | I, J) = \max_{b=1, \dots, B} \{\bar{\mathcal{H}}_b(\mathbf{T}_b x, f_b^\phi, f_x^\phi, f_b^\omega)\}. \quad (4)$$

Note that our occupancy is estimated for each hand separately in this stage. We use the shared network for both hands, but we distinguish each hand by feeding coarse 3D keypoints of each hand $J_L, J_R \in \mathbb{R}^{21 \times 3}$ separately as inputs to our network. As these keypoints are used to compute the canonicalization matrices and the pose feature, we can robustly learn our initial occupancy in the canonical space defined per-hand.

3.2. Two-Hand Occupancy Refinement

The previous step can robustly estimate per-hand occupancy volumes, but it does not effectively model the coherency between two hands due to the adaptation of the hand canonical spaces. As recent mesh-based two-hand reconstruction methods [23, 42] have shown that (1) two hand shapes are correlated with each other and thus (2) context-aware two-hand shape refinement is effective, we additionally propose a two-hand occupancy refinement network \mathcal{R} that refines the initial two-hand shape *in the original posed space*. It estimates the refined two-hand occupancies $o_l, o_r \in [0, 1]$ at the query point x conditioned on (1) the initial two-hand occupancy probabilities at x estimated by \mathcal{I} and (2) the input image I .

To condition our refined occupancy estimation on the interaction context, it is necessary to encode the initial geometry of two hands. To effectively learn features from the current two hand shapes, we first represent them as point clouds $\mathcal{P}_l, \mathcal{P}_r \in \mathbb{R}^{n \times 3}$, which are sets of iso-surface points of each side of hand geometry. These points can be easily obtained by collecting 512 farthest points among the query coordinates evaluated to be on surface by our initial hand occupancy network: $\{x \mid 0.5 - \epsilon \leq \mathcal{I}(x \mid I, J) \leq 0.5 + \epsilon\}$. We then encode each hand point cloud into a global latent vector and local latent vectors anchored in 3D space:

$$\{z_s, \mathcal{A}_s = \text{PCEnc}(\mathcal{P}_s, I)\}_{s=l,r}, \quad (5)$$

where $z_s \in \mathbb{R}^z$ is a global latent vector and $\mathcal{A}_s \in \mathbb{R}^{m \times (3+a)}$ is a set of m number of anchor points (i.e., a subset of the input point cloud selected by farthest point sampling) with a -dimensional per-point features. In our point cloud encoding module PCEnc , we first preprocess the input point cloud \mathcal{P}_s into a feature cloud $\mathcal{F}_s \in \mathbb{R}^{n \times (3+f)}$ to incorporate texture information observed from the input image I . To this end, we apply a simple encoder-decoder CNN to extract an image feature map $G \in \mathbb{R}^{w \times h \times f}$ and concatenate each point p in \mathcal{P}_s to an image feature vector located at the projected position of p in G . We then feed our feature cloud to the encoder of AIR-Net [14] that extracts global and locally anchored point cloud features using Point Transformer [43].

Using the obtained point cloud features, we extract a latent code that encodes interacting two-hand shapes and image context as follows:

$$z_c = \text{ContextEnc}(z_l, z_r, z_I). \quad (6)$$

ContextEnc is an MLP that extracts the context feature, and z_I is the global bottleneck feature from the image encoder-decoder previously used to extract G . Finally, we estimate our refined occupancy at the query point x as follows:

$$\{o_s = \text{PCDec}(x, o_s^i, \mathcal{A}_s, z_c)\}_{s=l,r}, \quad (7)$$

where PCDec is a point cloud decoder, for which we adopt a similar architecture to the decoder of AIR-Net [14]. The original AIR-Net decoder predicts an occupancy probability given a single point cloud via vector cross attention between the query point x and local anchor features \mathcal{A}_s and a global feature z_s (please refer to [14] for more details). Our decoder instead (1) uses a global latent vector z_c that encodes the context between the input image and the initial two hand shapes predicted by \mathcal{I} and (2) conditions the occupancy estimation on our initial occupancy probability by concatenating o_s^i at x to the input query coordinate. In summary, our refined occupancy estimation is effectively conditioned both on (1) local latent descriptors \mathcal{A}_s that encode information about the initial hand geometry, (2) a global latent descriptor z_c that encodes the global context of two-hand interaction and the input image, and (3) our initial occupancy estimation o_s^i .

3.3. Input Keypoint Refinement

In this section, we further consider image-based two-hand reconstruction using Im2Hands, where no ground truth hand keypoints are available. To enable robust shape reconstruction from keypoints *predicted* from an off-the-shelf image-based two-hand keypoint estimator (e.g., [11, 20, 23, 27, 42]), we introduce an optional keypoint refinement module \mathcal{K} that can alleviate noise in the input two-hand keypoints. Our keypoint refinement module is formulated as: $\mathcal{K}(J, I) \rightarrow J^r$, where $J, J^r \in \mathbb{R}^{42 \times 3}$ are initial and refined 3D two-hand keypoints, respectively. To be more specific, \mathcal{K} is designed as:

$$\mathcal{K}(J, I) = \text{MSA}([\text{GCN}(\text{KptEnc}(J)), \text{ImgEnc}(I)]). \quad (8)$$

In the above equation, we first extract input keypoint features by KptEnc that encodes the index and coordinate of each keypoint using a shared MLP. We then build a two-hand skeleton graph with initial node features set as features extracted from the previous KptEnc function. Next, we feed this initial two-hand skeleton graph to a GCN to extract keypoint features that further encode information of the initial two-hand structure. Finally, we apply multi-headed self-attention (MSA) between the keypoint features and image features – similarly to query-image attention in Equation 3 – to regress the refined two-hand keypoint positions. In Section 4, we experimentally demonstrate that \mathcal{K} significantly helps improving the quality of two-hand reconstruction from single images.

3.4. Loss Functions

We now explain loss functions to train our two-hand occupancy network.

Initial Hand Occupancy Network (\mathcal{I}). \mathcal{I} is trained by MSE loss that measures the deviation between the ground

truth and the predicted occupancy probabilities. For training query point generation, we combine (1) points uniformly sampled in the hand bounding box and (2) points sampled on the hand surface added with Gaussian noise following [18, 33].

Two-Hand Occupancy Refinement Network (\mathcal{R}). Similar to \mathcal{I} , \mathcal{R} is trained by MSE loss using the ground truth occupancy supervision. It is additionally trained with penetration loss, which penalizes the refined two-hand occupancy values that are estimated to be occupied in both hands at the same query position. Given a set of training samples \mathcal{X} , a set of sampled training query points \mathcal{P} , and \mathcal{R}_l and \mathcal{R}_r that output the refined occupancy probabilities for left and right hands, our penetration loss is defined as:

$$\mathcal{L}_{pen} = \frac{1}{|\mathcal{X}|} \sum_{(I,J) \in \mathcal{X}} \sum_{x \in \mathcal{P}} \mathcal{R}_l(x|I, J) \cdot \mathcal{R}_r(x|I, J), \quad (9)$$

where $\mathcal{R}_l(\cdot) > 0.5$ and $\mathcal{R}_r(\cdot) > 0.5$. We incorporate this loss function to avoid inter-penetration between two hands.

Input Keypoint Refinement Network (\mathcal{K}). \mathcal{K} is trained by MSE loss computed using the ground truth and the predicted two-hand keypoints. Due to the space limit, please refer to our supplementary section for training and network architecture details.

4. Experiments

4.1. Experimental Setups

4.1.1 Datasets, Metrics, and Keypoint Estimators

Datasets. We mainly use InterHand2.6M [27] dataset – the only interacting two-hand dataset with dense shape annotations – for both quantitative and qualitative evaluation. To maintain consistency with the previous work [23], we only use interacting hand (IH) samples annotated as *valid* hand type. The resulting dataset contains 366K training samples, 110K validation samples, and 261K test samples. For qualitative evaluation, we additionally demonstrate our results on RGB2Hands [39] and EgoHands [3] datasets, which contain RGB videos of interacting two hands without shape annotations.

Evaluation Metrics. For evaluating the quality of reconstructed two-hand shapes, we compute the mean Intersection over Union (IoU) and Chamfer L1-Distance (CD) between the predicted and the ground truth two-hand shapes. We also evaluate the accuracy of 3D hand keypoints after the proposed keypoint refinement using Mean Per Joint Position Error (MPJPE).

Two-Hand Keypoint Estimation Methods. To evaluate Im2Hands on *single-image* two-hand reconstruction, we leverage an off-the-shelf image-based two-hand keypoint estimation method. We consider DIGIT [11] and IntagHand [23], whose official implementation is publicly avail-

able. We would like to note that IntagHand is also proposed for two-hand *shape* estimation, thus we additionally consider it as a baseline for two-hand shape reconstruction. However, when using IntagHand for a *keypoint* estimation method in our image-based reconstruction experiments, we completely discard the predicted shape information and only use the estimated keypoints – considering it as a pure two-hand keypoint estimation method. We emphasize that our two-hand occupancy function works agnostic to the architecture of the previous two-hand keypoint pre-processor and can be used in a plug-and-play manner.

4.1.2 Baselines

As Im2Hands is the first neural implicit function for two interacting hands, we compare our method to (1) the existing *single-hand* reconstruction method using *implicit* representation [18] and (2) *two-hand* reconstruction methods using *mesh* representation [23, 42].

Implicit Single-Hand Reconstruction Methods. We consider HALO [18], which is a neural implicit single-hand representation driven by 3D keypoints. As HALO models a hand shape only using 3D keypoints, we modify HALO to additionally leverage an input RGB image by feeding pixel-aligned features [33] together with the other inputs to the part occupancy functions – to make more fair comparison to Im2Hands by using the same input data domains (see Section 4.2). Although LISA [7] is another implicit single-hand representation which takes both RGB image and 3D keypoints, we do not make direct comparison to LISA, as it requires more input data domains (e.g., foreground masks, multi-view RGB images) than Im2Hands.

Mesh-Based Two-Hand Reconstruction Methods. We compare our method to IntagHand [23] and Two-Hand-Shape-Pose [42], which are the state-of-the-art methods on interacting two-hand shape reconstruction. These methods reconstruct a fixed-topology mesh from an RGB image by using dense vertex correspondence [23] or statistical model parameter annotations [42] for training. Since they are designed for *single-image* two-hand reconstruction and do not take 3D hand keypoints as input, we also evaluate our model on image-based reconstruction by using 3D hand keypoints predicted from an image (see Section 4.3).

4.2. Reconstruction From Images and Keypoints

In this section, we examine the representation power of Im2Hands given an input RGB image paired with the ground truth 3D hand keypoints (e.g., obtained from a sensor). In Table 1, we quantitatively compare our reconstruction results to the baseline methods. In the fifth row, note that we also compare our method to a re-implemented version of HALO [18] that takes the same input data do-

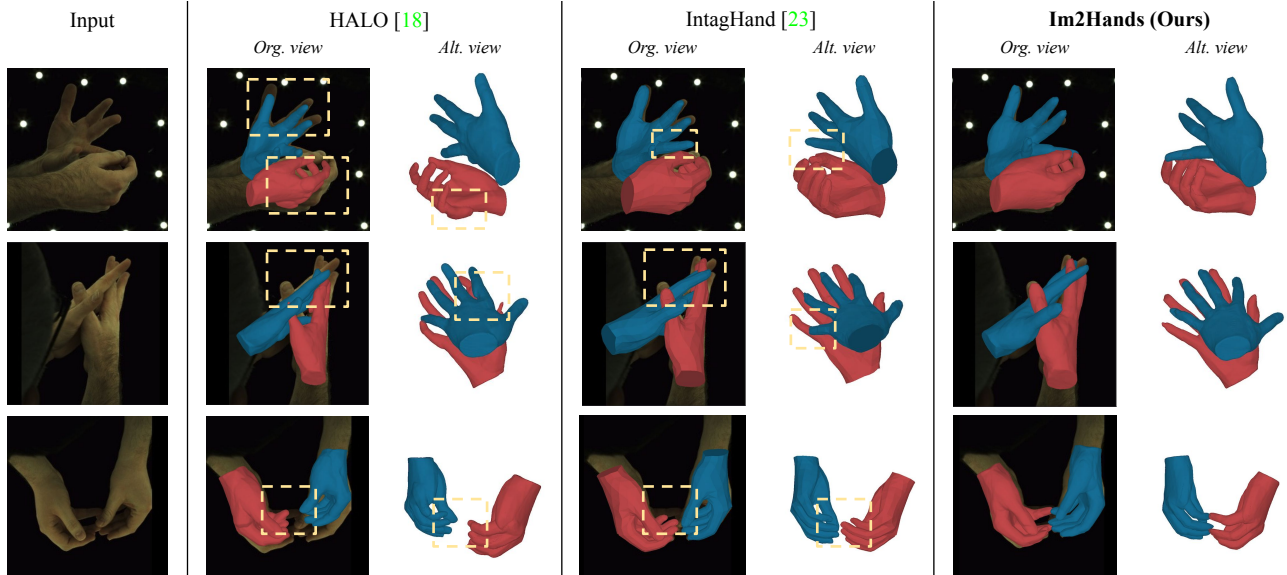


Figure 3. **Qualitative comparison of single-image interacting two-hand reconstruction on InterHand2.6M [27].** We compare our results to HALO [18] and IntagHand [23], where our method produces two-hand shapes with better hand-to-image and hand-to-hand alignment. Ours also properly captures finger contacts (*rows 1 and 3*) and avoids shape penetration (*row 2*). For our method and HALO, we use the hand keypoints predicted by DIGIT [11] to condition the articulated occupancy fields. Please see the supplementary section for more qualitative examples.

mains as ours (please refer to Section 4.1.2). In the table, Im2Hands outperforms all the other methods by effectively leveraging an RGB image and hand keypoints to model shape- and pose-dependent deformations, respectively.

Table 1. **Comparison of interacting two-hand reconstruction on InterHand2.6M [27].** For input domains of each method, **I**, **J**, and **L** denote RGB image, 3D hand keypoints, and hand bone lengths, respectively.

Method	Inputs	IoU (%) \uparrow	CD (mm) \downarrow
Two-Hand-Shape-Pose [42]	I, L	54.8	5.51
IntagHand [23]	I, L	67.0	3.88
HALO [18]	J	74.7	2.62
HALO (modified) [18]	I, J	75.8	2.51
Im2Hands (Ours)	I, J	77.8	2.30

4.3. Reconstruction From Single Images

In this section, we evaluate Im2Hands on single-image two-hand reconstruction. To this end, we estimate two-hand shapes using 3D hand keypoints predicted from an input image via off-the-shelf two-hand keypoint estimation methods [11, 23] and the proposed keypoint refinement module (\mathcal{K}). Note that the goal of this experiment is to evaluate each method in a setting where *no ground truth 3D hand keypoints* are available. Thus, we disabled the rescaling of the reconstructed hand joints and shapes performed by some of

the existing methods [23, 42] using the scale ratio calculated from a subset of the *ground truth* 3D keypoints at the test time. For those methods [23, 42] that require such rescaling, we use the mean scale of the hands in the training set of InterHand2.6M [27] to perform fair comparison.

In Table 2, we first evaluate the effectiveness of our keypoint refinement module on noisy two-hand keypoints predicted by DIGIT [11] and IntagHand [23]. Our method is successful in alleviating input keypoint errors. Especially, it is shown to be effective when the degree of the input keypoint error is high.

Table 2. **Effectiveness of our keypoint refinement module (\mathcal{K}) on InterHand2.6M [27].** We use two-hand keypoints predicted by DIGIT [11] and IntagHand [23] as inputs to our method.

Method	MPJPE (mm) \downarrow
DIGIT [11]	16.75
DIGIT [11] + \mathcal{K} (Ours)	10.70
IntagHand [23]	10.13
IntagHand [23] + \mathcal{K} (Ours)	9.68

We now report our two-hand shape reconstruction results from single images. In Table 3, our method achieves the state-of-the-art results in image-based interacting two-hand reconstruction on InterHand2.6M [27]. We would like to re-emphasize that the baselines originally designed for image-based two-hand reconstruction [23, 42] (1) can produce only fixed-resolution meshes and (2) require dense vertex corre-

spendences or statistical model parameter annotations for training, while our method outperforms them without such constraints. Also, note that most of the existing articulated implicit functions (e.g., [7, 8, 18]) are typically evaluated using *noiseless* keypoints or skeletons. In contrast, our method is experimentally demonstrated to achieve competitive results in a setting where *no ground truth keypoints* are available – by leveraging the proposed keypoint and two-hand shape refinement modules.

Table 3. **Comparison of single-image interacting two-hand reconstruction on InterHand2.6M [27].** We use two-hand keypoints predicted by DIGIT [11] and IntagHand [23] as inputs to our method and HALO [18].

Method	IoU (%) \uparrow	CD (mm) \downarrow
Two-Hand-Shape-Pose [42]	48.4	6.09
IntagHand [23]	59.0	4.69
DIGIT [11] + HALO [18]	45.1	7.64
IntagHand [23] + HALO [18]	53.8	5.38
DIGIT [11] + Im2Hands (Ours)	59.4	4.75
IntagHand [23] + Im2Hands (Ours)	62.1	4.35

Our qualitative comparison of image-based two-hand reconstruction is also shown in Figure 3. Note that HALO [18] produces non-smooth shapes due to the lack of shape or keypoint refinement steps given the noisy keypoint observation. IntagHand [23] does not properly model finger contacts (*rows 1 and 3*) or produces shape penetration (*row 2*). Our method generates more plausible two-hand shapes that also align well with the input image.

4.4. Generalizability Test

In Figure 4, we additionally show our qualitative results on RGB2Hands [39] and EgoHands [3] datasets. For two-hand keypoint estimation, we use DIGIT [11] and IntagHand [23] for RGB2Hands and EgoHands, respectively. Note that these reconstruction examples were directly generated by Im2Hands trained on InterHand2.6M [27] dataset only, demonstrating the generalization ability of Im2Hands. For more detailed setups for this experiment, please refer to our supplementary section.

4.5. Ablation Study

In this section, we perform an ablation study to investigate the effectiveness of the major modules (i.e., \mathcal{K} , \mathcal{I} , and \mathcal{R}) of Im2Hands using 3D hand keypoints predicted by DIGIT [11]. In Table 4, our full model is the most effective compared to the other model variants. Please refer to the supplementary section for ablation study results with respect to more detailed model variations (e.g. models after removing *each* of the proposed components inside \mathcal{I} and \mathcal{R}).

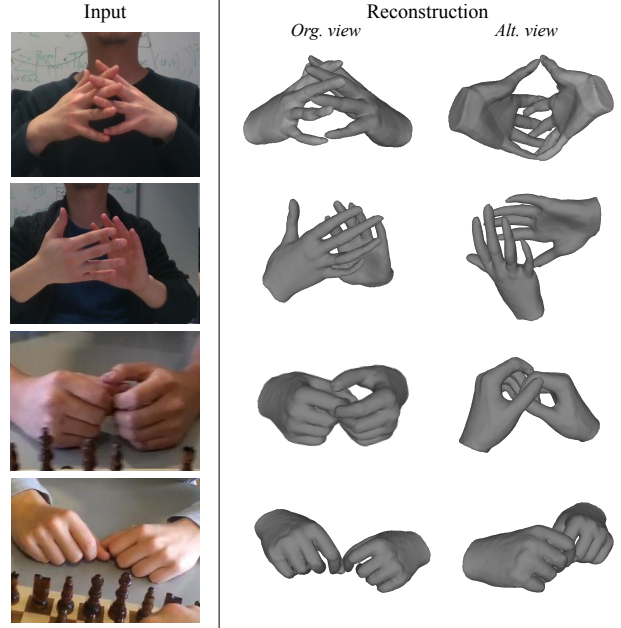


Figure 4. **Generalizability test on RGB2Hands [39] and EgoHands [3].** Top two rows show examples from RGB2Hands [39] and bottom two rows show examples from EgoHands [3].

Table 4. **Ablation study using the predicted 3D hand keypoints.** Experiments are performed on InterHand2.6M [27] dataset using the keypoints predicted by DIGIT [11].

Method	IoU (%) \uparrow	CD (mm) \downarrow
\mathcal{I}	45.6	7.61
$\mathcal{K} + \mathcal{I}$	55.4	5.36
Im2Hands ($\mathcal{K} + \mathcal{I} + \mathcal{R}$)	59.4	4.75

5. Conclusion and Future Work

We present Im2Hands, the first neural implicit representation of two interacting hands. To effectively model the interaction context between two articulated hands, we propose two modules for initial occupancy estimation in the hand canonical space and context-aware occupancy refinement in the original posed space, respectively. Furthermore, we introduce an optional input keypoint refinement module to enable robust shape reconstruction from single images. As a result, we achieve two-hand shape reconstructions with better hand-to-image and hand-to-hand alignment compared to the baseline methods.

Acknowledgements. This work is in part sponsored by NST grant (CRC 21011, MSIT) and KOCCA grant (R2022020028, MCST). Minhyuk Sung acknowledges the support of the NRF grant (RS-2023-00209723) and IITP grant (2022-0-00594) funded by the Korean government (MSIT), and grants from Adobe, ETRI, KT, and Samsung Electronics.

References

- [1] Seungryul Baek, Kwang In Kim, and Tae-Kyun Kim. Pushing the envelope for rgb-based dense 3d hand pose estimation via neural rendering. In *CVPR*, 2019. 1, 2
- [2] Seungryul Baek, Kwang In Kim, and Tae-Kyun Kim. Weakly-supervised domain adaptation via gan and mesh model for estimating 3d hand poses interacting objects. In *CVPR*, 2020. 1, 2
- [3] Sven Bambach, Stefan Lee, David J Crandall, and Chen Yu. Lending a hand: Detecting hands and recognizing activities in complex egocentric interactions. In *ICCV*, 2015. 6, 8
- [4] Adnane Boukhayma, Rodrigo de Bem, and Philip HS Torr. 3d hand shape and pose from images in the wild. In *CVPR*, 2019. 1
- [5] Yujun Cai, Lihao Ge, Jianfei Cai, and Junsong Yuan. Weakly-supervised 3d hand pose estimation from monocular rgb images. In *ECCV*, 2018. 2
- [6] Zhe Cao, Ilija Radosavovic, Angjoo Kanazawa, and Jitendra Malik. Reconstructing hand-object interactions in the wild. In *ICCV*, 2021. 2
- [7] Enric Corona, Tomas Hodan, Minh Vo, Francesc Moreno-Noguer, Chris Sweeney, Richard Newcombe, and Lingni Ma. Lisa: Learning implicit shape and appearance of hands. In *CVPR*, 2022. 2, 3, 6, 8
- [8] Boyang Deng, John P Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. Nasa neural articulated shape approximation. In *ECCV*, 2020. 2, 3, 8
- [9] Bardia Doosti, Shujon Naha, Majid Mirbagheri, and David J Crandall. Hope-net: A graph-based model for hand-object pose estimation. In *CVPR*, 2020. 2
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 4
- [11] Zicong Fan, Adrian Spurr, Muhammed Kocabas, Siyu Tang, Michael J Black, and Otmar Hilliges. Learning to disambiguate strongly interacting hands via probabilistic per-pixel part segmentation. In *3DV*, 2021. 1, 2, 3, 5, 6, 7, 8
- [12] Lihao Ge, Hui Liang, Junsong Yuan, and Daniel Thalmann. Robust 3d hand pose estimation in single depth images: from single-view cnn to multi-view cnns. In *CVPR*, 2016. 2
- [13] Lihao Ge, Zhou Ren, Yuncheng Li, Zehao Xue, Yingying Wang, Jianfei Cai, and Junsong Yuan. 3d hand shape and pose estimation from a single rgb image. In *CVPR*, 2019. 1, 2
- [14] Simon Giebenhain and Bastian Goldlücke. Air-nets: An attention-based framework for locally conditioned implicit representations. In *3DV*, 2021. 5
- [15] Yana Hasson, Bugra Tekin, Federica Bogo, Ivan Laptev, Marc Pollefeys, and Cordelia Schmid. Leveraging photometric consistency over time for sparsely supervised hand-object reconstruction. In *CVPR*, 2020. 2
- [16] 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 *CVPR*, 2019. 2
- [17] Umar Iqbal, Pavlo Molchanov, Thomas Breuel Juergen Gall, and Jan Kautz. Hand pose estimation via latent 2.5 d heatmap regression. In *ECCV*, 2018. 2
- [18] Korrawe Karunratanakul, Adrian Spurr, Zicong Fan, Otmar Hilliges, and Siyu Tang. A skeleton-driven neural occupancy representation for articulated hands. In *3DV*, 2021. 1, 2, 3, 4, 6, 7, 8
- [19] Korrawe Karunratanakul, Jinlong Yang, Yan Zhang, Michael J Black, Krikamol Muandet, and Siyu Tang. Grasping field: Learning implicit representations for human grasps. In *3DV*, 2020. 2
- [20] Dong Uk Kim, Kwang In Kim, and Seungryul Baek. End-to-end detection and pose estimation of two interacting hands. In *ICCV*, 2021. 2, 3, 5
- [21] Dominik Kulon, Riza Alp Guler, Iasonas Kokkinos, Michael M Bronstein, and Stefanos Zafeiriou. Weakly-supervised mesh-convolutional hand reconstruction in the wild. In *CVPR*, 2020. 1, 2
- [22] Nikolaos Kyriazis and Antonis Argyros. Scalable 3d tracking of multiple interacting objects. In *CVPR*, 2014. 2
- [23] Mengcheng Li, Liang An, Hongwen Zhang, Lianpeng Wu, Feng Chen, Tao Yu, and Yebin Liu. Interacting attention graph for single image two-hand reconstruction. In *CVPR*, 2022. 1, 2, 3, 5, 6, 7, 8
- [24] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *CVPR*, 2021. 1, 2
- [25] Marko Mihajlovic, Shunsuke Saito, Aayush Bansal, Michael Zollhoefer, and Siyu Tang. Coap: Compositional articulated occupancy of people. In *CVPR*, 2022. 2, 3
- [26] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. V2v-posenet: Voxel-to-voxel prediction network for accurate 3d hand and human pose estimation from a single depth map. In *CVPR*, 2018. 2
- [27] Gyeongsik Moon, Shou-I Yu, He Wen, Takaaki Shiratori, and Kyoung Mu Lee. Interhand2.6m: A dataset and baseline for 3d interacting hand pose estimation from a single rgb image. In *ECCV*, 2020. 1, 2, 3, 5, 6, 7, 8
- [28] Franziska Mueller, Micah Davis, Florian Bernard, Oleksandr Sotnychenko, Mickeal Verschoor, Miguel A Otaduy, Dan Casas, and Christian Theobalt. Real-time pose and shape reconstruction of two interacting hands with a single depth camera. *ACM TOG*, 2019. 2
- [29] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *ICCV*, 2021. 2, 3
- [30] Iasonas Oikonomidis, Nikolaos Kyriazis, and Antonis A Argyros. Tracking the articulated motion of two strongly interacting hands. In *CVPR*, 2012. 2
- [31] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: modeling and capturing hands and bodies together. *ACM TOG*, 2017. 1, 2, 3
- [32] Yu Rong, Jingbo Wang, Ziwei Liu, and Chen Change Loy. Monocular 3d reconstruction of interacting hands via collision-aware factorized refinements. In *3DV*, 2021. 1, 2, 3

- [33] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 4, 6
- [34] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. In *CVPR*, 2021. 2, 3
- [35] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multi-view bootstrapping. In *CVPR*, 2017. 2
- [36] Adrian Spurr, Umar Iqbal, Pavlo Molchanov, Otmar Hilliges, and Jan Kautz. Weakly supervised 3d hand pose estimation via biomechanical constraints. In *ECCV*, 2020. 2
- [37] Jonathan Taylor, Vladimir Tankovich, Danhang Tang, Cem Keskin, David Kim, Philip Davidson, Adarsh Kowdle, and Shahram Izadi. Articulated distance fields for ultra-fast tracking of hands interacting. *ACM TOG*, 2017. 2
- [38] Chengde Wan, Thomas Probst, Luc Van Gool, and Angela Yao. Dual grid net: Hand mesh vertex regression from single depth maps. In *ECCV*, 2020. 1, 2
- [39] Jiayi Wang, Franziska Mueller, Florian Bernard, Suzanne Sorli, Oleksandr Sotnychenko, Neng Qian, Miguel A Otaduy, Dan Casas, and Christian Theobalt. Rgb2hands: real-time tracking of 3d hand interactions from monocular rgb video. *ACM TOG*, 2020. 3, 6, 8
- [40] Ying Wu, John Lin, and Thomas S Huang. Analyzing and capturing articulated hand motion in image sequences. *IEEE TPAMI*, 2005. 2
- [41] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *NeurIPS*, 2021. 3
- [42] Baowen Zhang, Yangang Wang, Xiaoming Deng, Yinda Zhang, Ping Tan, Cuixia Ma, and Hongan Wang. Interacting two-hand 3d pose and shape reconstruction from single color image. In *ICCV*, 2021. 1, 2, 3, 5, 6, 7, 8
- [43] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *ICCV*, 2021. 5
- [44] Yuxiao Zhou, Marc Habermann, Weipeng Xu, Ikhsanul Habibie, Christian Theobalt, and Feng Xu. Monocular real-time hand shape and motion capture using multi-modal data. In *CVPR*, 2020. 1
- [45] Christian Zimmermann and Thomas Brox. Learning to estimate 3d hand pose from single rgb images. In *ICCV*, 2017. 2