

A Skeleton-Driven Neural Occupancy Representation for Articulated Hands

Conference Paper**Author(s):**

Karunratanakul, Korrawe; [Spurr, Adrian](#) ; Fan, Zicong; [Hilliges, Otmar](#) ; Tang, Siyu

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000518040>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/3DV53792.2021.00012>

A Skeleton-Driven Neural Occupancy Representation for Articulated Hands

Korraue Karunratanakul Adrian Spurr Zicong Fan Otmar Hilliges Siyu Tang
ETH Zürich, Switzerland

{korraue.karunratanakul, adrian.spurr, zicong.fan, otmar.hilliges, siyu.tang}@inf.ethz.ch

<https://korraue.github.io/halo>




Figure 1. We introduce a novel neural implicit surface representation of human hands (HALO) that is fully driven by keypoint-based skeleton articulation. Taking 3D keypoints as input, a fully differentiable implicit occupancy representation produces high-fidelity reconstruction of the hand surface (top row). We show that HALO facilitates the conditional generation of articulated hands that grasp 3D objects in a realistic and physically plausible manner (bottom row).

Abstract

We present *Hand Articulated Occupancy (HALO)*, a novel representation of articulated hands that bridges the advantages of 3D keypoints and neural implicit surfaces and can be used in end-to-end trainable architectures. Unlike existing statistical parametric hand models (e.g. MANO), HALO directly leverages the 3D joint skeleton as input and produces a neural occupancy volume representing the posed hand surface. The key benefits of HALO are (1) it is driven by 3D keypoints, which have benefits in terms of accuracy and are easier to learn for neural networks than the latent hand-model parameters; (2) it provides a differentiable volumetric occupancy representation of the posed hand; (3) it can be trained end-to-end, allowing the formulation of losses on the hand surface that benefit the learning of 3D keypoints. We demonstrate the applicability of HALO to the task of conditional generation of hands that grasp 3D objects. The differentiable nature of HALO is shown to improve the quality of the synthesized hands both in terms of physical plausibility and user preference.

1. Introduction

Humans grasp and manipulate objects with their hands. Modeling 3D poses and surfaces of human hands is important for numerous applications such as animation, games, augmented and virtual reality. Existing hand representations in the literature can be categorized into two paradigms: skeleton representations [19, 28, 40, 41, 71] and mesh-based representations [39, 51]. Even though 3D skeletons are defined in the Euclidean space and are easy to interface with deep neural networks, they lack surface information, therefore not suitable for reasoning about hand-object interaction. In contrast, mesh-based hand models provide surfaces and can thus explicitly reason about physical interactions, such as hand-object manipulation [25, 59]. However, their pose and shape parametrizations are often hard to directly interpret and more difficult to learn in an end-to-end fashion than 3D keypoints.

In this work, we aim to bridge the gap between 3D keypoints and dense surface models. To this end, we propose **Hand Articulated Occupancy (HALO)**, a novel hand representation that is driven by keypoint-based skeleton articulation.

ulation and provides a high-fidelity, resolution-independent neural implicit surface. Importantly, the proposed representation is fully differentiable and can be trained end to end such that volume-based losses can back-propagate gradients to the 3D keypoints. Specifically, HALO makes two key innovations: (1) a fully differentiable skeleton canonicalization algorithm and (2) a shape-aware neural articulated hand representation that is driven directly by a skeleton and naturally affords differentiable reasoning about surface and volumetric occupancy.

Differentiable skeleton canonicalization. The core advantage of the HALO model is that its surface representation is *fully differentiable* and driven by the *3D joint skeleton*. Although a naive iterative fitting procedure can fit a parametric hand model to a target skeleton, it prevents gradient flow from hand surface to its keypoints. In addition, due to the ambiguities of the twist angle around finger bones, the optimization procedure may lead to an unnatural hand surface. By leveraging biomechanical constraints that ensure plausible hand poses, we propose a novel differentiable layer that converts 3D keypoint locations to the corresponding bone rotations and translations with respect to a canonicalized pose space in a single forward pass. We call this operation a canonicalization layer. Importantly, this layer does not rely on iterative fitting and can therefore be effectively used in back-propagation-based learning. Furthermore, this layer enables the learning of implicit hand shape representation in the canonical space. This significantly improves the generalization capability of the learned representations across different hand poses and shapes, as shown in Sec. 5.

Skeleton-driven neural articulated hand. Recently, several neural occupancy networks for human body modeling, e.g. [15, 36, 53], have been proposed. Despite demonstrating the feasibility of parameterizing articulated deformations via neural implicit surfaces, all these methods require ground-truth bone transformations as input, making them unable to interface with keypoints. In addition, it is uncertain and has not been demonstrated how these models will function under noisy predictions without the ground truth transformations. Note that making implicit representations applicable to articulated hands, driven by skeleton, is non-trivial. The key challenge is how to infer realistic shapes of unseen hands from the skeletons with highly articulated hand poses. Our solution includes two steps. First, by leveraging biomechanical constraints, we ensure bijective mappings between 3D skeletons and hand surfaces using the canonicalization layer, which greatly simplifies the learning of the implicit surface for highly articulated hand skeletons; Second, in the canonical space, we learn both identity-dependent and pose-dependent deformations of the hand surface with a set of multi-layer perceptrons that generalize well across different hand shapes and poses. We systematically compare HALO with several baseline methods.

The results show that HALO significantly improves the accuracy, generality, and visual fidelity over the baselines.

Hand grasps synthesis. To demonstrate the utility of the keypoint-driven implicit surface representation, we employ HALO for the conditional generation of hands grabbing 3D objects. We propose a novel generation pipeline that synthesizes hand keypoints and yields neural occupancy of the hand via HALO. By exploiting the differentiable nature of HALO, we design a hand-object interpenetration loss to guide the training of the 3D keypoint generator in an end-to-end fashion. Our experiments show that this loss leads to convincing hand-object contact of the generated hands. Furthermore, compared to a MANO-based method (GrabNet [59]), HALO produces more physically plausible and visually convincing grasps even before refinement, suggesting that the volume-based losses are effective for learning the 3D keypoint generator.

Overall, the contributions include (1) HALO, the first neural implicit hand model that is driven by keypoint-based skeleton articulation, provides smooth neural implicit surfaces, and enables differentiable reasoning about volumetric occupancy; (2) A differentiable skeleton canonicalization layer that maps any skeleton to the canonical pose with a unique bio-mechanically valid transformations; (3) A realistic human grasp generation framework that leverages the efficient volumetric occupancy checks enabled by HALO.

2. Related Work

Hand pose and mesh estimation. Hand pose estimation is a long-standing question and several learning-based approaches have been introduced. These approaches generally involve predicting 3D keypoint locations [8, 13, 16, 18, 19, 28, 40, 41, 43, 54, 57, 58, 60, 64, 66, 71], regressing MANO [51] parameters [1, 2, 4, 25, 26, 68], or directly predicting the full dense surface of the hand [20, 30, 39, 61]. The methods that directly predict 3D key points usually achieve better performance, however, they do not yield dense surface which is crucial for hand interaction. Iteratively fitting a templated-based model such as MANO to the keypoints could recover the dense surface but also make the process non-differentiable [44, 46, 62]. Alternatively, the dense surface can also be estimated from 3D or 2D keypoints [11, 65, 70]. However, such estimation could result in a change of hand pose from the input keypoints. In contrast, our model produces a hand surface that faithfully respects the input pose and allows surface or volumetric losses to back-propagate directly to the keypoints.

Hand representation. The surface of 3D hands can be represented explicitly or implicitly. The commonly used template-based approaches such as MANO [51] induce a prior of poses and shapes over its learned parameter space for regularization. However, using the learned parameters

also increases the learning complexity as image features do not directly correspond to the learned parameters of a hand model. In [1, 4, 25, 26, 68], the MANO parameters are predicted using weak supervision such as hand masks [1, 68] or 2D annotations [4, 68]. Another way of representing 3D hand is to directly store the dense vertex locations of the MANO template [20, 30, 38]. While being more generalizable by avoiding constraints in the parameter space, these approaches require the corresponding dense 3D annotations, which might be difficult to acquire. Our work differs in the way that we can recover dense hand surfaces from 3D keypoints, eliminating the need for estimating model parameters or predicting 3D vertex locations.

Implicit representation. Several works represent object shapes by learning an implicit function using neural networks [10, 14, 21, 22, 35, 37, 47, 55], which allows for the modeling of arbitrary object topologies with dynamic resolution. Many approaches for learning such implicit functions from various input types were also proposed [32, 33, 45, 48, 52, 56]. These works focus on rigid objects and do not permit shape deformation. Recently, the interest is also on learning an articulated implicit function for human body [3, 15, 27, 36, 63]. NASA [15] represents human bodies using a set of implicit functions, but the model is limited to a specific body shape. LEAP [36] proposes to learn inverse linear blend skinning functions for multiple body shapes, however, it relies on ground truth bone transformation matrices instead of 3D joint locations. To the best of our knowledge, there are no implicit hand representations that can generalize well to various shapes. Grasping Field [29] learns an implicit function for hand and objects together to represent contact but treats every posed hand as a rigid object. As a result, the complexity of learning a wide range of poses increases significantly. In this work, we leverage biomechanical constraints of human hand to learn a novel hand model that only takes a skeleton as input and generalizes to different hand shapes and poses.

Hand-object interaction. There have been many studies into hand interaction with objects in various settings [5, 7, 9, 12, 17, 18, 23, 24, 29, 31, 42, 59]. Recently, the community has begun exploring the task of generating plausible hand grasps given an object with notable studies including [12], [29], and [59]. GanHand [12] generates grasps for each object in a given RGB image by first predicting a grasp type from grasp taxonomy [17] and its initial orientation, and then optimizing for better contact with the object. GrabNet [59] uses Basis Point Set [49] to represent 3D objects as input to generate MANO parameters. The predicted hand is then fed to a refinement model to improve the contact. Grasping Field [29] learns a signed distance field for both hand surface and object surface in one space, allowing the contact to be learned as regions where distances to both surfaces are zeros. However, the output surface cannot be




Figure 2. **Local coordinate systems and rotation matrices defined in the systems.** (a) We adopt the technique from [57] to construct local coordinate systems F_i for each segment of the kinematic chains of the hand. (b) Each F_i is constructed by measuring the flexion angle θ^f and the abduction angle θ^a relative to the parent bone. See Supp. Mat. for further details.

articulated and requires hand model fitting. Our work differs from others in the way that we use the proposed hand representation to model the contact while keeping the synthesis task as simple as generating 3D keypoints.

3. HALO: Hand Articulated Occupancy

The HALO model is a skeleton-driven neural occupancy function, formally defined as $\mathcal{O}_w(x|\mathbf{J}) \rightarrow \{0, 1\}$. Parameterized by neural network weights w , it maps a 3D point x to its occupancy value given the hand skeleton represented by a set of 3D keypoint locations \mathbf{J} . In this section, we first describe how to convert an arbitrary 3D joint skeleton to the reference canonical pose in a differentiable and consistent manner, then we introduce our simple yet effective neural occupancy network for hands.

Notations. Given a hand skeleton represented by 3D keypoints $\mathbf{J} : \mathbb{R}^{21 \times 3}$, we denote θ_i^f and θ_i^a to be the flexion/extension and abduction/adduction angles of bone i relative to its parent, respectively. For simplicity, we refer to them as flexion and abduction angle. The angle between a palmar bone i and its adjacent palmar bone $i + 1$ is denoted as θ_i^p . Lastly, θ_i^n is the palmar plane angle between plane n_i and n_{i+1} spanned by the palmar bone $i - 1$, i , $i + 1$ respectively. We denote the properties of the reference canonical hand with $^c(\cdot)$. We refer to Supp. Mat. for further details.

3.1. Canonicalization of 3D Hand Skeleton

Our goal is to learn a neural representation of the surface of human hands in the canonical space. Furthermore, we want to deform this shape based on the spatial configuration of the underlying skeleton, represented by 3D keypoints. To do so, we require a mechanism that allows us to convert the 3D keypoints into valid skeletons in the canonical pose in terms of joint angles. As the keypoints have no notion of the surface, naively converting them to axis-angles does not work due to the unconstrained twist of bones. While twist

does not affect keypoints, they affect the surface.

We take inspiration from Spurr et al. [57] which defines a consistent local coordinate system for each bone to measure the bone angles for semi-supervised learning. Our objective is to derive a differentiable mapping layer that i) provides means to convert predicted keypoints to the rest pose and back, and ii) ensures that the skeleton is free of implausible twists that would influence the surface.

Building on [57], we represent each finger bone by two rotation angles, flexion and abduction, relative to its parent bone (Fig. 2). Each bone cannot rotate about itself, thus, no twist. However, such formulation ignores the palm configuration, which is needed for defining the canonical pose. In this work, we propose a method to parameterize the pose of a palm in order to define a consistent canonical pose. We decouple the palmar bone configuration into 1) finger spreading and 2) palm arching. The spreading of fingers is captured via the angles between two adjacent palmar bones. The arching of the palm is defined by the angle between the two planes spanned by three adjacent palmar bones. The resulting palmar region then serves as a frame of reference for the remaining fingers. Please refer to Supp. Mat. Fig. B.1 for better visualization.

Converting 3D Keypoints to Bone Transformations. Formally, we seek the unique set of transformations $\{\mathbf{B}^{-1}\}$ that maps the skeleton $\mathbf{J} : \mathbb{R}^{21 \times 3}$ to the canonical pose ${}^c\mathbf{J}$. Given a skeleton, we obtain the set $\{\mathbf{B}^{-1}\}$ by sequentially performing the following operations: First, we rotate each finger to match the description of our canonical palm pose, which we define as a flat hand with fixed angles between palmar bones; Second, we compute joint angles and local coordinate systems following [57] (Fig. 2b), which we use to iteratively undo the angles along the kinematic chain (Fig. 2a) to acquire the canonical pose. By combining the transformations from both steps, then adjust for the conversion from keypoints to bone vectors, we could obtain a set of transformation matrices $\{\mathbf{B}^{-1}\}$ that maps the given skeleton \mathbf{J}_i to our canonical pose ${}^c\mathbf{J}_i$. Formally,

$${}^c\mathbf{J}_i = \mathbf{B}_i^{-1} \mathbf{J}_i \quad , \quad (1)$$

where $\mathbf{B}^{-1} = \mathbf{T}\mathbf{S}\mathbf{F}'\mathbf{R}^c\mathbf{F}\mathbf{P}\mathbf{K}$.

Here $\mathbf{K} : \mathbb{R}^{3 \times 21} \mapsto \mathbb{R}^{3 \times 20}$ is a function that maps the keypoints $\mathbf{J} \in \mathbb{R}^{21 \times 3}$ to bone vectors by translating them to the local origin and scaling to unit norm; \mathbf{P} normalizes the palmar bone and palmar planes angles; \mathbf{F} then maps the bones to their local coordinate frames; \mathbf{R}^c rotates each bone to have the same flexion and abduction angles as the canonical pose. Finally, \mathbf{F}' maps each coordinate frame back to the global coordinate system; \mathbf{S} reverts bones back to their original length and \mathbf{T} translates the bones to the tip of their parent bones.

This set of transformations is unique for each skeleton

pose and only allows biomechanically valid transformation. For details, we kindly refer the reader to our Supp.Mat.

3.2. Neural Occupancy Networks for Hands

Here, we describe how to leverage the unique mapping $\{\mathbf{B}^{-1}\}$ between the posed skeleton and the canonical skeleton to learn the neural hand representation that generalizes to different shapes with highly articulated poses. We draw inspirations from NASA [15] and explore similar neural network structure due to its simplicity and efficacy.

NASA [15]. NASA learns an implicit representation of a human body $\mathcal{O}_w(x|\theta) \rightarrow \{0, 1\}$, conditioned on the pose descriptor θ . Specifically, it defines the implicit surface for each body part separately. Let \mathbf{B}_b^{-1} be the transformation to the canonical pose for bone b , NASA can be denoted by:

$$\mathcal{O}_w(x|\theta) = \max_b \{\bar{\mathcal{O}}_w^b(\mathbf{B}_b^{-1}x, \Pi_w^b[\{\mathbf{B}^{-1}t_0\}])\}. \quad (2)$$

where the pose descriptor θ is defined by a collection of transformation matrices $\{\mathbf{B}_b\}_{b=1}^B$, and the probability of $\mathcal{O}_w(x|\theta)$ is derived from the maximum occupancy probability across B child occupancy functions $\bar{\mathcal{O}}_w^b(\cdot, \cdot)$, where each represents the body part of the bone b . For a query point x , each child function $\bar{\mathcal{O}}_w^b(\cdot, \cdot)$ maps x to its local coordinate system by the transformation matrix \mathbf{B}_b^{-1} , so that the local shape of each body part can be learned. The term $\Pi_w^b[\cdot]$ is used to provide global pose information to each child function. Essentially, by querying the occupancy value x using $\mathbf{B}_b^{-1}x$, the NASA model learns a template shape and the correction based on the global pose with $\Pi_w^b[\cdot]$. Note that, the bone transformation $\{\mathbf{B}_b\}_{b=1}^B$ is assumed to be given. For more details, we kindly refer the reader to [15].

Neural Occupancy Networks for Hands. A naive adaptation of the NASA model for human hand results in erroneous surface reconstructions as shown in Sec. 5.1 (Fig. 5). In order to represent hands with highly articulated poses and diverse shapes, we propose to learn the child occupancy functions by conditioning on a shape descriptor β , effectively learning $\mathcal{O}_w(x|\theta, \beta)$. We assume that the identify-dependent deformations of the hand are highly correlated to the bones, hence, we propose the use of a collection of bone lengths as the shape descriptor. In particular, we propose a simple yet effective bone length encoder $f^b(\mathbf{D})$ that takes $\mathbf{D} = [d_1, d_2, \dots, d_B] \in \mathbb{R}^B$ the bone length of individual bones as inputs. We emphasize that under the proposed formulation, we could learn the hand surface using only the key points \mathbf{J} , as the pose descriptor \mathbf{B}_b^{-1} is derived from \mathbf{J} by Eq. 1. Our final occupancy function is given by:

$$\mathcal{O}_w(x|\theta, \beta) = \max_b \{\bar{\mathcal{O}}_w^b(\mathbf{B}_b^{-1}x, \Pi_w^b[\{\mathbf{B}^{-1}t_0\}], f^b(\mathbf{D}))\} \quad (3)$$

where each implicit function $\bar{\mathcal{O}}_w^b$ learns the corresponding part shape based on the hand pose descriptor $\Pi_w^b[\{\mathbf{B}^{-1}t_0\}]$ and our bone length descriptor $f^b(\mathbf{D})$.

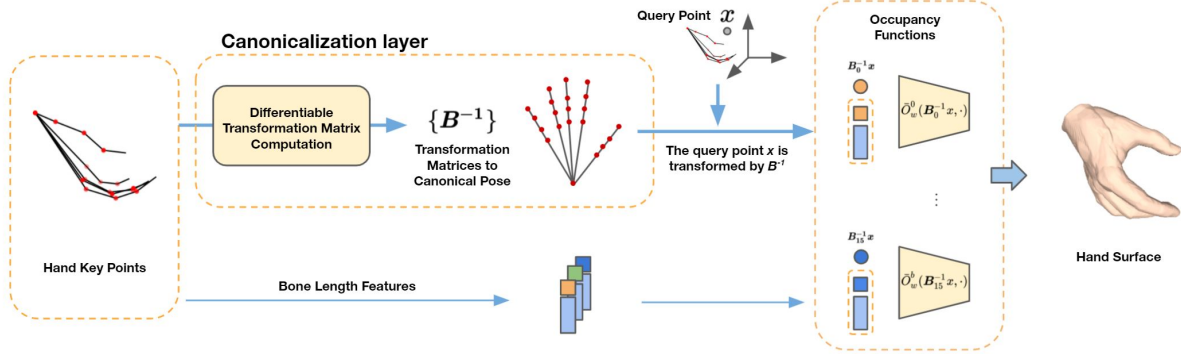


Figure 3. **Overview of HALO.** Given a hand skeleton, HALO derives bone transformations $\{B^{-1}\}$ that map the bones to the canonical pose using the canonicalization layer. The query point x is then transformed into the canonical space as $B_i^{-1}x$ for the occupancy check.




Figure 4. **HALO-VAE Architecture.** Given object point cloud, the VAE model synthesizes a grasping hand represented by a set of 3D keypoints. The input hand keypoints and the hand encoder are only used during the training. From the synthesized hand skeleton, HALO then produces the hand surface. The whole pipeline is end-to-end trainable, therefore, the VAE can leverage volume-based losses to improve the generation quality of the 3D keypoints.

Shape descriptor variations. We investigate two versions of bone length encoders $f^b(\mathbf{D})$: the local bone encoder $f_l^b(\mathbf{D}) = d_b$ and the global encoder $f_g^b(\mathbf{D}) = [d_b; MLP(\mathbf{D})]$, where the MLP for the global encoder has two linear layers. We follow a similar training strategy as [15], for more training details, please refer to the Supp. Mat.

3.3. Skeleton-driven Articulated Hand Model

To build a skeleton-driven articulated hand model, we combine the previously described canonicalization layer and the neural hand surface together. Specifically, HALO takes the input 3D keypoints to compute bone transformations $\{B^{-1}\}$ for the occupancy networks using the canonicalization layer. As the canonicalization layer is differentiable, the model can be trained end-to-end and allows volume-based losses from the surface to back-propagate to the keypoints. The overview of HALO is shown in Fig. 3. Note that the bone lengths \mathbf{D} can also be computed from the keypoints. During inference, only 3D keypoints are needed as input to reconstruct hand surface.

4. Human Grasps Generation

We show the utility of the HALO model in the challenging task of grasps generation. Given an object, we generate diverse grasps with natural and plausible hand-object interaction. Our grasp generation pipeline consists of two parts: a 3D keypoints generator based on a variational autoencoder (VAE) and the HALO model for obtaining the hand surface.

HALO-VAE Architecture. The architecture of the HALO-VAE model is illustrated in Fig. 4. During training, the object point cloud is first passed to the object encoder, which is a modification of PointNet [50] with residual connection [35], to obtain an object latent code. The object latent code is then concatenated to the 3D hand joint location, $\mathbf{J} \in \mathbb{R}^{21 \times 3}$, and passed to the VAE encoder. The decoder reconstructs the 3D hand joint positions conditioned on the hand and object latent representation. From the key points, the surface is obtained using HALO through the skeleton canonicalization layer.

The advantages of using HALO are two-fold. First, we decouple the complexity of learning the pose, represented by the skeleton, from that of learning the surface that corresponds to the pose; Second, the implicit model enables fast intersection tests between hand and object, which can be used to efficiently compute an interpenetration loss. Combined with the differentiable skeleton canonicalization layer, the interpenetration loss can be used to improve the keypoints generator in both end-to-end training and post-optimization refinement.

Our grasp generation pipeline is similar to [59] and [29], but with the following key differences. First, in [29], the output is a rigid implicit surface that cannot be articulated. To obtain an animatable hand for downstream tasks, additional MANO model fitting is required. Second, in [59], the grasps generator is trained to produce the MANO parameters which is not directly related to the Euclidean space where the hand and the object live in. The challenge of interfacing the MANO parameters with deep neural networks

is reflected in the GrabNet (CoarseNet) [59] results which will be discussed in the experiments section.

4.1. Learning and Losses

To train the VAE model, we use the following losses: the KL-divergence loss on the hand latent Z , L2 loss on the predicted key points, L1 bone lengths loss, and the bone angle losses. The bone angle losses are used to provide additional supervisions for learning the hand structure which consists of 1) flexion angles θ_i^f and abduction angles θ_i^a , 2) angles between adjacent palmar bones θ_i^p and 3) angles between palmar planes θ_i^n . The bone angles are the same as used in Sec. 3. The losses are defined as L1 angle difference between the prediction and the ground truth.

Interpenetration loss. In addition to the losses on the keypoints, we also use the interpenetration loss on the hand surface to avoid collision between hand and object. The key idea is to penalize every point inside the object that is also occupied by the hand. Concretely, for a set of points sampled inside the object P_o and the predicted key points \bar{J} , the interpenetration loss is defined as:

$$\mathcal{L}_{in}(\bar{J}) = \sum_{p \in P_o} \mathcal{O}_w(p | \Theta(\bar{J}), f_g(D_{\bar{J}})), \text{ where } \mathcal{O}_w(\cdot) > 0.5 \quad (4)$$

where $D_{\bar{J}}$ is the bone length vector for \bar{J} and $\Theta(\bar{J})$ maps the predicted key points to the HALO pose vector θ using the differentiable transformation matrices in Eq. 1.

4.2. Optimization-based Refinement

To demonstrate that the efficient intersection tests enabled by HALO can be used for optimization, we refine the sampled hands by changing the global translation t to avoid collision with the object. The refinement is run for 10 steps with the interpenetration loss term in Eq. 4. The optimization objective is:

$$\min_t \mathcal{L}_{in}(\bar{J} + t) \quad (5)$$

This simple optimization step aims at refining the contact after the initial prediction of HALO-VAE. It is analogous to the RefineNet in [59], but with an explicit objective to avoid collision instead of being a neural network denoiser.

5. Experiments

In this section, we assess our skeleton-driven hand model and the grasp synthesis pipeline. First, in Sec. 5.1, we validate the efficacy of HALO as a neural implicit hand model and compare it to the surface baseline [15] and keypoints-to-surface baselines [11, 70]. Second, we show in Sec. 5.2 that HALO can be used effectively in generative tasks which require surface-based reasoning in form of grasp synthesis. For more experiments, please see supplementary materials.

	IoU \uparrow	Cham. (mm) \downarrow	Norm. \uparrow
NASA [15]	0.896	1.057	0.955
NASA+surf.	0.883	1.177	0.944
NASA+surf.+local b.	0.913	0.884	0.950
HALO (ours)	0.932	0.719	0.959
HALO keypoints (ours)	0.930	0.740	0.959

Table 1. **Evaluation on IoU, Chamfer-distance (L1), and normal consistency score (Norm.) between NASA [15] and HALO.** All models are trained using ground truth bone transformations except for *HALO keypoints* where only the 3D key points are given. *NASA+surf.* indicates a NASA model with resampled surface points for the skinning loss \mathcal{L}_s and *local b.* indicates that a corresponding bone length is given to each occupancy function. The results show that HALO outperforms NASA [15] on IoU and Chamfer-L1 by large margins, suggesting the superior performance of HALO in terms of fidelity and generality.

Methods	IOU \uparrow	Cham. (mm) \downarrow	MPJPE (mm) \downarrow
Choi et al. [11]	0.43	4.651	14.1
Zhou et al. [70]	0.54	2.811	7.95
HALO keypoints (ours)	0.93	0.740	0

Table 2. **Comparison between the estimated, root-aligned surfaces when only 3D keypoints are given as input.**

5.1. Neural Hand Model

We first evaluate the performance of the proposed implicit surface representation and analyze the effect of the keypoint-to-transformation mapping layer.

Training data. To train our neural occupancy hand model, we utilize MANO [51] hand meshes. Following [15], for each mesh we sample points with two strategies: 1) uniformly sampling in the hand bounding box, 2) sampling on the surface with additional isotropic Gaussian noise. Only the uniformly sampled points are used for evaluation. The associated occupancy value of each query point is computed by casting a ray from the sampled point and counting the number of intersections along the ray. The ground truth bone transformation matrices are computed along the kinematic chain to transform the template MANO hand into the target pose. The skinning weights are taken from the skinning weights of MANO. We use the Youtube3D (YT3D) hands dataset [30] in all our experiments. The YT3D training set contains 50,175 hand meshes of hundreds of subjects performing a wide variety of tasks in 102 videos. The test set covers 1,525 meshes from 7 videos.

Evaluation metrics. For 3D surface reconstruction evaluation, we compute the mean Intersection over Union (IoU), Chamfer-L1 distance, and normal consistency score [35].

5.1.1 Comparison to implicit surface baseline

Here we investigate the generalization ability of the proposed HALO model to represent articulated hands with var-




Figure 5. **Qualitative results of HALO and HALO-VAE.** (Left) Comparison between the NASA and HALO. Two views are shown for the same pose for comparison. (Right) Hands sampled from HALO-VAE.

	HALO-VAE w/o \mathcal{L}_{in}	HALO-VAE (ours)	[59]-coarse	HALO-VAE + Optim	[59]-refine
Int. vol. (cm ³) ↓ / Cont. ra. (%) ↑					
Binoculars	9.19 / 1.00	6.35 / 1.00	8.97 / 0.95	4.47 / 1.00	3.24 / 1.00
Camera	3.99 / 1.00	4.47 / 1.00	3.31 / 0.75	1.66 / 1.00	1.46 / 1.00
Frying pan	0.25 / 0.85	0.22 / 0.65	0.94 / 0.80	0.05 / 0.35	1.10 / 0.85
Mug	3.38 / 1.00	6.41 / 1.00	5.72 / 1.00	3.48 / 1.00	4.48 / 1.00
Toothpaste	6.25 / 1.00	1.85 / 1.00	6.05 / 1.00	1.11 / 1.00	2.28 / 1.00
Wineglass	1.41 / 0.90	2.35 / 1.00	2.62 / 1.00	1.56 / 0.80	1.60 / 0.95
Average	4.08 / 0.96	3.61 / 0.94	4.60 / 0.92	2.06 / 0.85	2.36 / 0.97
Diversity					
Entropy ↑	2.88	2.88	2.84	2.88	2.83
Cluster size ↑	2.25	2.15	1.65	2.15	1.70

Table 3. **Evaluation on interpenetration volume (Int. vol.), contact ratio (Cont. ra.) and diversity.** HALO-VAE w/o \mathcal{L}_{in} denotes the HALO-VAE model *without* the interpenetration loss. Best numbers are in bold, except for the contact ratios when more than one model have the same best result.

ious **poses** and **shapes**. The results are summarized in Tab. 1 **Baseline.** We use the NASA model [15] as our baseline. The NASA model is designed to represent an implicit function of an articulated body. However, by changing the input dimension and the number of part-models to match the number of hand parts, it can also be used to represent an articulated hand. We train the baseline model using the bone transformation matrices taken from MANO and the sampled query points. For details on implementation and network architecture we refer to the supplementary.

Surface vertex re-sampling. In [15], the surface vertices v used for enforcing the part models in the skinning loss \mathcal{L}_s are the mesh vertices of SMPL [34]. Similarly, we use MANO surface vertices during training. However, we notice that the human-designed mesh often has many more vertices in the area around the joints which could cause the part models to bias toward the bone endpoints. Thus, we propose to re-sample the surface vertices uniformly on the mesh surface. This result is performance degradation but the bone connections are more natural with less artifact.

Local and global bone encoders. The bone lengths of a human hand greatly influence the hand shape. Therefore, for the local bone encoder, we add the bone length d_i to

the back-projected query point as input to the part model $[B_i^{-1}x; d_i]$. As shown in Tab. 1, the local bone encoder improves the reconstruction quality both in terms of IoU and Chamfer-L1 distance. We further extend the local bone encoder by considering all the bone lengths as input. A concatenated vector of bone lengths is first fed into a small feed-forward neural network to get the global bone feature $f_g(D)$, which is then concatenated with the query point x and the local bone length d_i as input to the part model.

Results. By combining the local and global bone encoders, HALO significantly improves the 3D surface reconstruction quality compared to NASA. As shown in Tab. 1, the IoU is increased from 0.896 to 0.932 and the Chamfer-L1 distance is decreased from 1.057mm to 0.719mm.

We provide a qualitative comparison between NASA, and HALO in Fig. 5, confirming the quantitative results. The proposed HALO representation generalizes well for highly articulated poses, whereas the NASA model produces severe artifacts at the connection between parts.

5.1.2 3D keypoints to hand surface

Tab. 1 also shows the result from HALO that only takes 3D keypoints as input. The keypoint model achieves compara-

% users rated as ‘more realistic’	
HALO-VAE w/o \mathcal{L}_{in}	HALO-VAE
32.77 %	67.23 %
HALO-VAE	GrabNet - coarse
55.75 %	44.25 %
HALO-VAE + Optim	GrabNet - refine
51.29 %	48.71 %

Table 4. **Binary choice user study.** The number show the percentage of users who rate the corresponding method as more realistic.

ble surface reconstruction performance as when the ground truth transformation are given, showing the effectiveness of our method. We show the qualitative results in Fig. 1 and 5.

In addition, to evaluate the keypoint-to-surface pipeline, we then compare HALO to the equivalent component in [70] and [11] which estimates hand surface from 3D keypoints. The evaluation is done on same the Youtube3D test set where the ground truth 3D keypoints are given as input. As [11] requires both 2D and 3D coordinates as inputs, the 2D keypoints is obtain by projecting the 3D keypoints perpendicular to the palm. For evaluation, we also report the 3D joint error between the predicted hand and the input joints. This metric measures if the input keypoints are faithfully respected by the models. By design, the HALO model does not change the keypoint locations from input to output, thus does not have this error. The comparison in Tab. 2 shows that [70] and [11] change the hand pose and shape in the prediction while HALO faithfully reconstructs the hand surface according to the given keypoints.

5.2. Grasp Synthesis

To assess the utility of HALO in downstream tasks we demonstrate our grasp generative model, HALO-VAE.

Dataset. We leverage the recently introduced GRAB dataset [6, 59] and compare our results to GrabNet [59]. We compare both to the initial (coarse) predictions of GrabNet [59] and the refined results which matches with our own two-stage generation process. The test set contains 6 unseen objects. For each object, we fix the object orientation and sample 20 hand proposals from each model.

Physics Metrics. Following [67, 69], we evaluate the physical plausibility (interpenetration volume and contact ratio) and diversity, and provide results from a perceptual study. To evaluate the interpenetration and contact, we measure the ratio of frames in which the hand is in contact with the object and average the interpenetration volume. The volume is calculated by voxelizing hand and object mesh with 1mm cubes and counting the number of intersecting cubes.

User study. We asked 75 participants in a forced-alternative-choice perceptual study to ‘select the grasp that is more realistic’. For each question, the user is shown 4 views per grasp and forced to select one. We compare all possible combinations on the same object. Each question is assigned to at least 2 participants, totaling 4,800 data points per pair of model comparison. To ensure that the grasps from HALO-VAE and GrabNet have the exact same texture, we fit MANO to our generated key points for rendering.

Diversity. Following [69], we compute the diversity of the sampled grasps by performing k-means with 20 clusters on all samples, then evaluate the entropy of the cluster assignment and the average cluster size. More diversity results in higher value for both metrics. We use the flatten key point

locations of the hands after aligning the root joint and the plane spanned by middle and index palmar bone as features.

Results. We first validate the efficacy of the interpenetration loss (Eq. 4). We compare the HALO-VAE models with and without the interpenetration loss. The results show that the interpenetration loss helps in: 1) reducing the collision between the objects and the generated hands (Tab. 3, col.1-2), and 2) largely improves the user preference of the corresponding model (Tab. 4, first row), demonstrating the efficacy of the proposed neural occupancy representation of articulated hand for reasoning about hand-object interaction.

Next, we compare HALO-VAE with GrabNet [59]. Both HALO-VAE and GrabNet-coarse are CVAE based generative models and end-to-end trainable, the key difference is that GrabNet-coarse generates MANO model parameters whereas HALO-VAE generates 3D keypoints. As shown in Tab. 3, HALO-VAE outperforms GrabNet-coarse by a large margin for interpenetration volume and sample diversity. Moreover, the HALO-VAE model without interpenetration also compares favorably to GrabNet-coarse, suggesting that the 3D keypoints-based representation is well suited to interface with deep neural networks.

Finally, we compare our optimization-based refinement with GrabNet-refine. To the best of our knowledge, the RefineNet is not trained end-to-end with GrabNet-coarse and used for three steps during the inference. As shown in Tab. 4, our refined grasps attain a higher user score, suggesting they are more realistic and natural compared to the grasps refined by GrabNet-refine.

6. Discussion and Conclusion

In this work, we introduce HALO, a novel surface representation for articulated hands that can generalize to different hand poses and shapes. We address the issue of the transformation matrix requirement for inferring the 3D occupancy hand by proposing a skeleton canonicalization algorithm that computes valid transformations from 3D keypoints. The experiments show that our proposed hand model outperforms the baseline and can represent a wide range of hand poses and shapes. Finally, we demonstrate the HALO can be used to train an end-to-end grasp generator conditioned on an object and produces hand grasps with natural and realistic interaction. We believe that HALO can be useful in future work attempting to reconstruct the surface of articulated hands directly from images via differentiable rendering and for several downstream tasks that need to perform surface-based computation such as collision detection and response.

Acknowledgement. We sincerely acknowledge Shaofei Wang and Marko Mihajlovic for the insightful discussions and helps with the baselines.

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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1067–1076, 2019. 2, 3
- [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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6121–6131, 2020. 2
- [3] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Loopreg: Self-supervised learning of implicit surface correspondences, pose and shape for 3d human mesh registration. *Advances in Neural Information Processing Systems*, 33, 2020. 3
- [4] Adnane Boukhayma, Rodrigo de Bem, and Philip HS Torr. 3d hand shape and pose from images in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10843–10852, 2019. 2, 3
- [5] Samarth Brahmhatt, Cusuh Ham, Charles C Kemp, and James Hays. Contactdb: Analyzing and predicting grasp contact via thermal imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8709–8719, 2019. 3
- [6] Samarth Brahmhatt, Cusuh Ham, Charles C. Kemp, and James Hays. ContactDB: Analyzing and predicting grasp contact via thermal imaging. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 8, 6
- [7] Samarth Brahmhatt, Chengcheng Tang, Christopher D. Twigg, Charles C. Kemp, and James Hays. ContactPose: A dataset of grasps with object contact and hand pose. In *The European Conference on Computer Vision (ECCV)*, August 2020. 3
- [8] Yujun Cai, Lihao Ge, Jianfei Cai, and Junsong Yuan. Weakly-supervised 3d hand pose estimation from monocular rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 666–682, 2018. 2
- [9] Yu-Wei Chao, Wei Yang, Yu Xiang, Pavlo Molchanov, Ankur Handa, Jonathan Tremblay, Yashraj S. Narang, Karl Van Wyk, Umar Iqbal, Stan Birchfield, Jan Kautz, and Dieter Fox. DexYCB: A benchmark for capturing hand grasping of objects. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [10] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [11] Hongsuk Choi, Gyeongsik Moon, and Kyoung Mu Lee. Pose2mesh: Graph convolutional network for 3d human pose and mesh recovery from a 2d human pose. In *European Conference on Computer Vision (ECCV)*, 2020. 2, 6, 8
- [12] Enric Corona, Albert Pumarola, Guillem Alenyà, Francesc Moreno-Noguer, and Grégory Rogez. Ganhand: Predicting human grasp affordances in multi-object scenes. In *CVPR*, 2020. 3
- [13] Martin de La Gorce, David J Fleet, and Nikos Paragios. Model-based 3d hand pose estimation from monocular video. *IEEE transactions on pattern analysis and machine intelligence*, 33(9):1793–1805, 2011. 2
- [14] Boyang Deng, Kyle Genova, Soroosh Yazdani, Sofien Bouaziz, Geoffrey Hinton, and Andrea Tagliasacchi. Cvxnet: Learnable convex decomposition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 31–44, 2020. 3
- [15] Boyang Deng, JP Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. Neural articulated shape approximation. *European Conference on Computer Vision (ECCV)*, 2020. 2, 3, 4, 5, 6, 7
- [16] Bardia Doosti, Shujon Naha, Majid Mirbagheri, and David J Crandall. Hope-net: A graph-based model for hand-object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6608–6617, 2020. 2
- [17] Thomas Feix, Javier Romero, Heinz-Bodo Schmiedmayer, Aaron M Dollar, and Danica Kragic. The grasp taxonomy of human grasp types. *IEEE Transactions on human-machine systems*, 46(1):66–77, 2015. 3
- [18] Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, and Tae-Kyun Kim. First-person hand action benchmark with rgb-d videos and 3d hand pose annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 409–419, 2018. 2, 3
- [19] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3593–3601, 2016. 1, 2
- [20] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 10833–10842, 2019. 2, 3
- [21] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Local deep implicit functions for 3d shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 3
- [22] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7154–7164, 2019. 3
- [23] Patrick Grady, Chengcheng Tang, Christopher D Twigg, Minh Vo, Samarth Brahmhatt, and Charles C Kemp. Contactopt: Optimizing contact to improve grasps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1471–1481, 2021. 3
- [24] Shreyas Hampali, Mahdi Rad, Markus Oberweger, and Vincent Lepetit. Honnotate: A method for 3d annotation of hand and object poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3196–3206, 2020. 3, 1
- [25] Yana Hasson, Bugra Tekin, Federica Bogo, Ivan Laptev, Marc Pollefeys, and Cordelia Schmid. Leveraging photo-

- metric consistency over time for sparsely supervised hand-object reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 571–580, 2020. 1, 2, 3
- [26] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11807–11816, 2019. 2, 3, 1
- [27] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3093–3102, 2020. 3
- [28] Umar Iqbal, Pavlo Molchanov, Thomas Breuel Juergen Gall, and Jan Kautz. Hand pose estimation via latent 2.5 d heatmap regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 118–134, 2018. 1, 2
- [29] Korrawe Karunratanakul, Jinlong Yang, Yan Zhang, Michael Black, Krikamol Muandet, and Siyu Tang. Grasping field: Learning implicit representations for human grasps. *arXiv preprint arXiv:2008.04451*, 2020. 3, 5, 1, 2
- [30] Dominik Kulon, Riza Alp Guler, Iasonas Kokkinos, Michael M. Bronstein, and Stefanos Zafeiriou. Weakly-supervised mesh-convolutional hand reconstruction in the wild. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2, 3, 6
- [31] Shaowei Liu, Hanwen Jiang, Jiarui Xu, Sifei Liu, and Xiaolong Wang. Semi-supervised 3d hand-object poses estimation with interactions in time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14687–14697, 2021. 3
- [32] Shichen Liu, Shunsuke Saito, Weikai Chen, and Hao Li. Learning to infer implicit surfaces without 3d supervision. In *Advances in Neural Information Processing Systems*, pages 8295–8306, 2019. 3
- [33] Shaohui Liu, Yinda Zhang, Songyou Peng, Boxin Shi, Marc Pollefeys, and Zhaopeng Cui. Dist: Rendering deep implicit signed distance function with differentiable sphere tracing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2019–2028, 2020. 3
- [34] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. 7
- [35] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 3, 5, 6
- [36] Marko Mihajlovic, Yan Zhang, Michael J Black, and Siyu Tang. Leap: Learning articulated occupancy of people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10461–10471, 2021. 2, 3
- [37] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421. Springer, 2020. 3
- [38] Gyeongsik Moon and Kyoung Mu Lee. I2l-meshnet: Image-to-lixel prediction network for accurate 3d human pose and mesh estimation from a single rgb image. In *European Conference on Computer Vision (ECCV)*, 2020. 3
- [39] Gyeongsik Moon, Takaaki Shiratori, and Kyoung Mu Lee. Deephandmesh: A weakly-supervised deep encoder-decoder framework for high-fidelity hand mesh modeling. *European Conference on Computer Vision (ECCV)*, 2020. 1, 2
- [40] 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 *Proceedings of the IEEE conference on computer vision and pattern Recognition*, pages 5079–5088, 2018. 1, 2
- [41] Gyeongsik Moon, Shouo-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 *European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 7
- [42] Arsalan Mousavian, Clemens Eppner, and Dieter Fox. 6-dof graspnet: Variational grasp generation for object manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2901–2910, 2019. 3
- [43] Franziska Mueller, Florian Bernard, Oleksandr Sotnychenko, Dushyant Mehta, Srinath Sridhar, Dan Casas, and Christian Theobalt. Generated hands for real-time 3d hand tracking from monocular rgb. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 49–59, 2018. 2
- [44] Franziska Mueller, Micah Davis, Florian Bernard, Oleksandr Sotnychenko, Miekeal 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 Transactions on Graphics (TOG)*, 38(4), 2019. 2
- [45] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020. 3
- [46] Paschalis Panteleris, Iason Oikonomidis, and Antonis Argyros. Using a single rgb frame for real time 3d hand pose estimation in the wild. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 436–445. IEEE, 2018. 2
- [47] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 3
- [48] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. *arXiv preprint arXiv:2003.04618*, 2020. 3
- [49] Sergey Prokudin, Christoph Lassner, and Javier Romero. Efficient learning on point clouds with basis point sets. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. 3, 1

- [50] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *CoRR*, abs/1612.00593, 2016. 5, 1
- [51] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6), 2017. 1, 2, 6, 5
- [52] 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 *Proceedings of the IEEE International Conference on Computer Vision*, pages 2304–2314, 2019. 3
- [53] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J. Black. SCANimate: Weakly supervised learning of skinned clothed avatar networks. In *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. 2
- [54] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1145–1153, 2017. 2
- [55] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33, 2020. 3
- [56] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, pages 1121–1132, 2019. 3
- [57] Adrian Spurr, Umar Iqbal, Pavlo Molchanov, Otmar Hilliges, and Jan Kautz. Weakly supervised 3d hand pose estimation via biomechanical constraints. In *European Conference on Computer Vision (ECCV)*, 2020. 2, 3, 4
- [58] Adrian Spurr, Jie Song, Seonwook Park, and Otmar Hilliges. Cross-modal deep variational hand pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 89–98, 2018. 2
- [59] 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)*, 2020. 1, 2, 3, 5, 6, 7, 8
- [60] Bugra Tekin, Federica Bogo, and Marc Pollefeys. H+ o: Unified egocentric recognition of 3d hand-object poses and interactions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4511–4520, 2019. 2
- [61] Chengde Wan, Thomas Probst, Luc Van Gool, and Angela Yao. Dual grid net: Hand mesh vertex regression from single depth maps. In *European Conference on Computer Vision*, pages 442–459. Springer, 2020. 2
- [62] 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 Transactions on Graphics (TOG)*, 39(6), 12 2020. 2
- [63] Shaofei Wang, Andreas Geiger, and Siyu Tang. Locally aware piecewise transformation fields for 3d human mesh registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7639–7648, 2021. 3
- [64] Ying Wu, John Lin, and Thomas S Huang. Analyzing and capturing articulated hand motion in image sequences. *IEEE transactions on pattern analysis and machine intelligence*, 27(12):1910–1922, 2005. 2
- [65] Lixin Yang, Jiasen Li, Wenqiang Xu, Yiqun Diao, and Cewu Lu. Bihand: Recovering hand mesh with multi-stage bisected hourglass networks. In *BMVC*, 2020. 2
- [66] Linlin Yang and Angela Yao. Disentangling latent hands for image synthesis and pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9877–9886, 2019. 2
- [67] Siwei Zhang, Yan Zhang, Qianli Ma, Michael J Black, and Siyu Tang. Place: Proximity learning of articulation and contact in 3d environments. In *8th international conference on 3D Vision (3DV 2020)(virtual)*, 2020. 8
- [68] Xiong Zhang, Qiang Li, Hong Mo, Wenbo Zhang, and Wen Zheng. End-to-end hand mesh recovery from a monocular rgb image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2354–2364, 2019. 2, 3
- [69] Yan Zhang, Mohamed Hassan, Heiko Neumann, Michael J Black, and Siyu Tang. Generating 3d people in scenes without people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6194–6204, 2020. 8
- [70] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2, 6, 8
- [71] Christian Zimmermann and Thomas Brox. Learning to estimate 3d hand pose from single rgb images. In *Proceedings of the IEEE international conference on computer vision*, pages 4903–4911, 2017. 1, 2