

# Multiview Image-based Hand Geometry Refinement using Differentiable Monte Carlo Ray Tracing

Giorgos Karvounas<sup>12</sup>

gkarv@ics.forth.gr

Nikolaos Kyriazis<sup>1</sup>

kyriazis@ics.forth.gr

Iason Oikonomidis<sup>1</sup>

oikonom@ics.forth.gr

Aggeliki Tsoli<sup>1</sup>

aggeliki@ics.forth.gr

Antonis A. Argyros<sup>12</sup>

argyros@ics.forth.gr

<sup>1</sup> Computational Vision and

Robotics Laboratory

Institute of Computer Science

Foundation for Research and

Technology - Hellas (FORTH)

Heraklion, Crete

Greece

<sup>2</sup> Computer Science Department

University of Crete

Heraklion, Crete

Greece

---

## Abstract

The amount and quality of datasets and tools available in the research field of hand pose and shape estimation act as evidence to the significant progress that has been made. However, even the datasets of the highest quality, reported to date, have shortcomings in annotation. We propose a refinement approach, based on differentiable ray tracing, and demonstrate how a high-quality publicly available, multi-camera dataset of hands (InterHand2.6M) can become an even better dataset, with respect to annotation quality. Differentiable ray tracing has not been employed so far to relevant problems and is hereby shown to be superior to the approximative alternatives that have been employed in the past. To tackle the lack of reliable ground truth, as far as quantitative evaluation is concerned, we resort to realistic synthetic data, to show that the improvement we induce is indeed significant. The same becomes evident in real data through visual evaluation.

## 1 Introduction

Estimating the pose and shape of a human hand in 3D from RGB images is an important problem with numerous applications in human-computer interaction, augmented reality, robotics and more. State-of-the-art solutions have shown impressive results in real-world settings (see Section 2) with their remarkable performance heavily relying on advances in deep learning architectures [1, 2] and the availability of large datasets for training. Synthetic datasets have perfect ground truth information, but their lack of realism hinders generalization to the real domain. Real datasets model the real world directly, but come with noisy manual or automatically generated pose/shape annotations, that are nevertheless considered to be ground truth (see Figure 1). As a result, state-of-the-art 3D hand reconstructions from

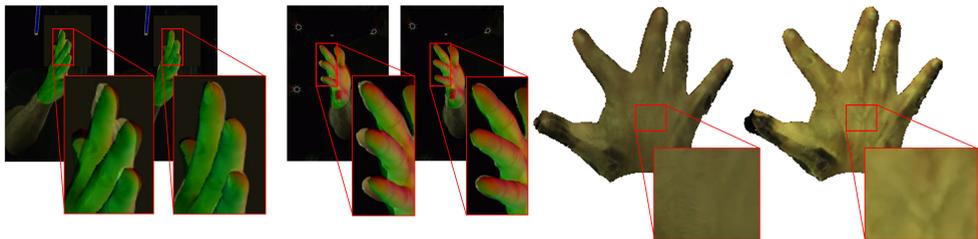


Figure 1: Real datasets of hands, *e.g.* multiview dataset InterHand2.6M [34], as well as results of state-of-the-art 3D hand pose/shape estimators, yield imperfect backprojections, which is evidence of error. We propose a method to improve the geometrical estimation of the hand, by exploiting color consistency, through a novel differentiable ray tracer, by Li *et al.* [24]. In the three exemplar comparisons taken from the same instance, the left sides correspond to input state and the right sides to output state, after our method has been applied. As is evident, a better fitting geometry increases color consistency, which is reflected in the sharpness and definition of the appearance map, on the right.

visual data exhibit mediocre image-to-model alignment. Increasing the amount of training data will improve the results only to the degree that the quality of the accompanied annotations is high. The ability to improve upon imperfect hand pose and shape estimates provided as image annotations or as the output of a neural network during training or testing can greatly benefit the performance of hand reconstruction methods.

In this work, we propose a method for refining an imperfect estimate of the geometry of a hand using calibrated multiview images as input. We follow a render-and-compare approach, and optimize the shape, pose and global orientation of the hand so that the discrepancy between the rendered and the observed hand, across multiple views, is minimized. The effectiveness of our method lies (a) in the density of the information exploited as constraints to the problem, *i.e.* all the pixels in all images that correspond to the hand, as well as (b) the appropriateness of sampling in rendering, taken into account in the forward and the backward step of the optimization.

No previous work addresses unsupervised hand geometry refinement. However, the constancy of the appearance across images has been exploited in the context of hand or human pose estimation in recent works [2, 33, 41]. Most rendering components invoked in the literature can be mapped to Pytorch3D [46], which implements Neural Rendering [20] and Soft Rasterization [27]. Contrary to what has been employed so far, we make use of differentiable ray tracing [24, 36], and show that it delivers improved results, in comparison. For an impression of the gap between rasterization and ray tracing, and sampling correctness the reader is referred to [13]. We employ the differentiable Monte Carlo ray tracer by Li *et al.* [24], which provides a solution to gradient discontinuities at the edges. The latter practically implies that the benefits of ray tracing can also be exploited in the backward pass, instead of just the forward pass.

We showcase our method on a subset of the recently proposed real and large-scale dataset InterHand2.6M [34]. InterHand2.6M contains imperfect ground truth annotations (see Figure 1) comprising 3D joint locations, derived in a semi-automatic way, and associated hand shape and pose descriptions, expressed via the MANO hand model, with a reported error of around 5mm. To circumvent the lack of ground truth data that are reliable enough for quanti-

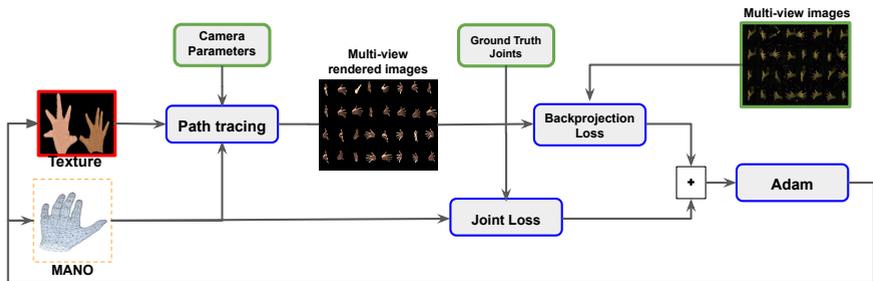


Figure 2: To perform multiview refinement of 3D hand geometry we compute the scalar difference between the backprojection of the currently estimated hand geometry and appearance and the observed input images. We update the estimates according to the gradient of the difference at the estimation point. Backprojection is performed using a ray tracer. 3D joint estimates, when available, can optionally be incorporated to aid search.

tative performance assessment, we resort to synthetic experimentation that closely resembles the setup in InterHand2.6M. We are able to successfully refine human hand estimates corresponding to a large variety of hand poses, shapes and appearances.

Our contributions can be summarized as follows.

- We propose an effective unsupervised method for hand pose and shape refinement from multiview images, given an initial estimate of the 3D hand geometry and, optionally, noisy target 3D joint locations. We complement the sparse noisy information about the 3D joint locations with rich image cues related to the appearance of the hand.
- We employ the proposed method to improve on a real dataset that is already on par with the state of the art with respect to annotation quality, while being richer in other aspects (number of samples, diversity in hand shape, pose, appearance, etc), yielding a dataset that is overall of even better quality.
- We employ, for the first time, a differentiable ray tracer to the problem of estimating the 3D pose and shape of hands, that is shown to be theoretically and practically superior to the rasterization methods that have been employed so far.
- We provide a suite of synthetic experimentation to connect the degree of color consistency, as a function of geometry and appearance, with the improvement in 3D estimation of the hand, for the described scenario.

## 2 Literature Overview

The task of 3D hand pose estimation from markerless visual input is a significant problem with numerous applications. It has been explored since the '90s [47] and is still actively researched today [14, 16, 18, 52, 55, 56, 58]. Numerous important applications can be built based on hand pose estimation systems, including Human-Computer Interaction such as Augmented and Virtual Reality, Performance Capture and Assessment, and Sign Language Recognition and Translation. The task is very challenging, due to several interdependent, complicating factors such as fast hand motions, self-occlusions, wide hand articulation range, and frequent interaction with the environment [9].

Given the difficulty of the problem, the early approaches predominantly relied on multi-

view input to alleviate ambiguities arising from occlusions, an approach that has recently gained again popularity [8, 15, 16, 39, 40, 47, 62, 62]. Stereo input was similarly explored [25, 41, 48]. With the wide availability of commodity depth sensors around 2010, the research also focused on monocular depth or RGBD input [11, 21, 30, 32, 37, 88, 44, 57, 58, 59, 63]. However, shortly thereafter, the success of deep learning led to the proliferation of robust systems that perform on monocular RGB input [9, 5, 10, 12, 14, 18, 23, 35, 42, 60, 55, 56, 60, 64, 65]. More recently, event cameras have also been investigated [52].

The early approaches typically estimated the parameters of a kinematics skeleton with predefined bone lengths, either in the form of joint angles, or as 3D joint positions directly [8, 47]. The accuracy achieved with the availability of depth information yielded the first practical real-time systems [88], and enabled on-the-fly adjustment of the bone lengths to the observed hand [30, 61]. The development and availability of a parametric hand pose and shape model [50] in turn fueled research towards estimating the surface of the observed hand [8, 12, 23, 33, 63]. Finally, some recent works are presenting the first attempts to model the appearance of the hand to some extent [1, 53, 45]. Their main focus is on the estimation of an appropriate hand model, while the affects of lighting and shadows are factored out.

Differentiable graphics renderers have lately gained traction as per their impact in many Computer Vision related tasks, too. Notably, such an approach was adopted specifically for hand pose estimation in the early work by De la Gorce *et al.* [8]. In that work, a graphics pipeline was custom fitted with the ability to compute derivatives with respect to parameters of the scene. More recently, several generic renderers with built-in automatic differentiation have been presented [20, 24, 27, 28, 29, 36] with [20] being the most widely adopted renderer. In this work, we propose the first method for inferring hand geometry from visual data that employs differentiable ray tracing [24]. This design choice allows for far more accurate modeling of the lighting in the scene than existing differentiable renderers, taking a big step towards handling the full complexity of real scenes. As we show in our experimental evaluation, the accuracy of estimating the hand geometry when we employ [24] instead of the state-of-the-art rasterization-based renderer [20] is significantly improved. Rendering via differentiable ray tracing can, additionally, benefit other domains that rely on analysis-by-synthesis approaches such as human pose estimation, 6D object pose estimation etc.

There is no previous work that focuses solely on unsupervised refinement of hand geometry. However, refinement of the hand geometry considering image evidence has been used within broader frameworks for reconstructing hands from visual data. For instance, a rendered hand is compared against the segmented silhouette of the observed hand [10, 9] or its texture, where the hand texture is modeled using a uniform color prior [7] or a texture model [45]. Although using the appearance of the hand for matching carries more information than using the hand silhouette, inaccurate inference of the appearance of the hand in [7, 45] led to no improvement on hand pose and shape estimation. We present an effective way for leveraging the appearance of the hand for pose and shape refinement via matching the appearance of the hand across multiple views. There exists limited previous work that has relied on photometric consistency among multiple views for hand pose estimation [63, 41] and human body estimation [1, 43]. The quality of the results is highly affected by the quality of the rendered hand or body geometries and the corresponding derivatives from the differentiable renderer. Note that dealing with human hands compared to the human body poses the extra challenge that the appearance of the fingers is self-similar.

### 3 Method

The input to our method is a calibrated multiframe,  $M$ , an initial estimate of the hand geometry parameters,  $h$ , and optionally a set of 3D annotations,  $J$ , for the joints of the hand.  $M$  consists of a set of RGB images  $\{O_i\}, i = 1, \dots, N$  along with the extrinsic and intrinsic calibration information  $\{C_i\}, i = 1, \dots, N$  of the corresponding cameras where  $N$  is the number of cameras. We represent the hand geometry using the MANO hand model [54], which is parameterized by  $h = \{\beta, \theta, \tau, \phi\}$ , where  $\beta \in R^{10}$  is the identity shape,  $\theta \in R^{45}$  the pose,  $\tau \in R^3$  the 3D translation of the root and  $\phi \in R^3$  the 3D rotation of the root. Any other differentiable representation of the hand that can be decoded into 3D geometry, e.g. [33], is employable, too. The optional 3D joint annotations,  $J$ , can be either a product of  $h$  or a loosely coupled estimate, as is the case in the InterHand2.6M dataset.

The output of our method is the finetuning,  $\delta h$ , of the hand geometry parameters,  $h$ , such that the new geometry derived from  $h$  and  $\delta h$ , when backprojected, fits the observations  $O_i$  more accurately than  $h$ . Similarly to previous work [44], we rely on the color consistency assumption. Contrary to [44], we provide an alternative formulation of the problem, striving for higher fidelity. More specifically, instead of requiring all rasterized pixels corresponding to projected geometry on each image to have the same color intensity in all reprojections on other images, we assume that each part of the geometry is consistently colored, through a common color map, *i.e.* an image encoded in a set of variables  $c$ , and require that all backprojections, up to per camera  $i$  color transforms  $T_i$ , match the observations. The main differences between the two approaches are that color consistency in our approach is a hard constraint, rather than a soft one, and that we do not compare rasterizations to color intensities interpolated in image space, which are prone to aliasing, but, rather, we require for each backprojected pixel, that constitutes an integration of samples, to match the observations in the corresponding pixels, which are also regarded as the integrals they really are. The integrals in reference regard the rendering equation [19] which we compute through [24]. This approach guarantees that the constraints drawn from observations are properly defined and their effects are properly backpropagated from image differentiation to geometrical changes, with sampling appropriateness being the highlight.

We, thus, define the following optimization problem:

$$\delta \hat{h}, \hat{c}, \hat{T} = \underset{\delta h, c, T}{\operatorname{argmin}} E(\delta h, c, T), \quad (1)$$

where

$$E(\delta h, c, T) = E_1(\delta h, c, T) + \lambda E_2(\delta h). \quad (2)$$

The data term,  $E_1$ , of Equation (2) amounts to the backprojection error between the renderings of the estimated geometry across multiple views and the corresponding observed multiview images (Section 3.1). The prior term,  $E_2$ , acts as a penalty for solutions that are far from the initial estimate of 3D joint locations (Section 3.2).

#### 3.1 Data term

We define the backprojection error as

$$E_1(\delta h, c, T) = \sum_{i=1, \dots, N} f_1(O_i, T_i \cdot R(G_v(h + \delta h), V(c), C_i)). \quad (3)$$

**Geometry:**  $G_v(x)$  computes the 3D vertices of the geometry represented by  $x$ . Without loss of generality, in the presented work we implement this function through the MANO model [24]. This means that  $x \in \mathbb{R}^{61}$ , encoding global position and orientation, as well as non-dimensionality-reduced articulation. Instead of our highly regularized approach, one might opt for a non-parametric representation, to also accommodate off-model deformations, without further changes in modeling.

**Color mapping:** Each surface element of the hand geometry is mapped to a pixel colour during rendering. This mapping takes the form of an image which is mapped to the geometry through a per pixel coordinate transform (see [25]). We employ the same mechanism, but instead of a single image we consider a set  $c = \{V_{low}, V_{high}\}$  of two images of different sizes:  $V_{low} \in \mathbb{R}^{8 \times 8 \times 3}$  and  $V_{high} \in \mathbb{R}^{512 \times 512 \times 3}$ . Thus,  $V(c) = V_{low} + V_{high}$ .  $V_{low}$  is scheduled early during optimization (see Section 3.3) and represents low frequency base color information.  $V_{high}$  is added to the optimization schedule in a 2nd phase, as a detail residual on top of  $V_{low}$ , which through addition also includes medium and high frequency information.  $V_{low}$  is scaled up in size so as to match  $V_{high}$  prior to addition, through nearest neighbor interpolation.

**Rendering:**  $R(g, V, C)$  is a function that renders an image, through ray tracing, out of the geometry  $g$ , using a camera specification  $C$ , and a color map  $V \in \mathbb{R}^{H \times W \times 3}$  of height  $H$  and width  $W$ , with known vertex associations to  $g$ . The renderer  $R$ , is implemented using the differentiable ray tracer by Li et al. [24]. Color consistency is hardly met in real conditions. To account for slightly different colorization across cameras we allow the rendering result for camera  $i$  to also vary with respect to a  $3 \times 3$  linear color transform  $T_i$ , which is the same for all pixels corresponding to camera  $i$ . Best results have been attained for such transforms that are close but not equal to the identity transform.

**Render-and-compare:** The comparison between an observed image  $o$  and a rendered hypothesis  $r$  amounts to the aggregation of the per pixel color differences ( $D_c$ ) and the differences between edge detections ( $D_e$ ) computed on  $o, r$ :

$$\begin{aligned} f_1(o, r) &= D_c(o, r) + D_e(o, r) \\ D_c(o, r) &= |o - r| \\ D_e(o, r) &= |Canny(o) * K - Canny(r) * K| \end{aligned} \quad (4)$$

where  $Canny(\cdot)$  denotes the application of the Canny edge detector [6], implemented through [49].  $K$  is a  $11 \times 11$  Gaussian kernel of  $\sigma = 5$ , that is convolved with the edge detections, prior to differentiation. The detected edges reflect the prominence of the hands against the background. The diffusive blurring counters the sharpness of the edge detection results and, thus, practically widens the basin of convergence.

### 3.2 Prior term

We define the joint error as

$$E_2(\delta h) = f_2(J, G_J(h + \delta h)), \quad f_2(x, y) = \begin{cases} 0.5(x - y)^2 / \gamma, & |x - y| < \gamma \\ |x - y| - 0.5\gamma, & \text{otherwise} \end{cases}, \quad (5)$$

where  $G_J(x)$  computes the 3D joint locations of the hand geometry  $x$  which are compared against the target 3D joint locations,  $J$ . The error function  $f_2$  is a coordinate-wise smooth

L1 loss. The purpose of this term is to favor solutions whose joint predictions are within  $\gamma$  millimeters from  $J$  in all axes. We use  $\gamma = 5$  in all our experiments. It is worth noting that the weight of this term  $\lambda$  can be zero for the majority of cases, but it is assigned a small value ( $\lambda = 10^{-3}$ ) to robustify optimization against poses with high ambiguity (e.g. closed fist) or cases where the initial starting point is significantly far from the sought optimum.

### 3.3 Optimization

The optimization was carried out using the Adam algorithm [22]. The simultaneous optimization of geometry and color is a very challenging problem, as the optimization dynamics yield a race between which variable set settles first, leading to unfavorable local optima. To remedy this, we introduce a warm-up phase of 50 iterations where only variables  $T$  (per-camera color transform) and  $V_{low}$  (low detail color map) are considered. This warm-up phase is succeeded by an optimization of all the variables ( $\delta h, c, T$ ), for a remainder of 150 iterations. This scheduling has been experimentally shown to be robust for the case of the InterHand2.6M dataset. The implementation was done in PyTorch. For the average case of 30 views in a multiframe, each optimization iteration requires 0.93s for the rasterization path and 13.77s for the ray tracing path with 16 number of samples per pixel in both the forward and the backward pass, as estimated on a system powered by a *Nvidia Titan V GPU*. For the 200 iterations we employed this yields a duration of 3.1min for the rasterization path and 45.9min for the ray tracing path to process a multiframe.

## 4 Results

### 4.1 Data and metrics

We evaluated our pose and shape refinement method on real data from the InterHand2.6M dataset [34]. To circumvent the lack of reliable ground truth, as required for quantitative assessment, we additionally performed experiments on synthetic data with perfect ground truth. These data closely simulated the InterHand2.6M dataset. In the next sections, we report the following error metrics:

1. **mean Euclidean per-vertex distance**  $\varepsilon_v$  between the ground truth and predicted hand geometry (in *mm*),
2. **backprojection error on image intensities**  $\varepsilon_b$  ( $D_c$  term in Equation (4)) between the rendered projection of the refined geometry and the observed image, expressed in intensity levels (range 0 – 255).

### 4.2 Quantitative analysis on synthetic data

#### 4.2.1 Geometry refinement from noisy input

We assessed the performance of our proposed method for varying levels of noise regarding the initial estimate of the hand geometry. We closely simulated the capture environment of the InterHand2.6M dataset, using the same number and location of virtual cameras (see Figure 3), and considered 5 gestures, each performed by 2 different subjects, resulting in a total of 10 multiframe or 305 independent frames (camera counts varied between 29 and 32). The gestures were chosen to depict varying degrees of self-occlusion. Noise was simulated

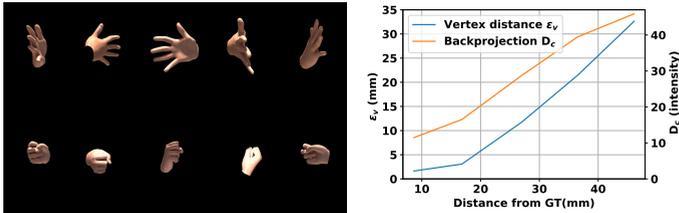


Figure 3: **(left)** Examples of synthetic data. **(right)** Performance of our method on hand geometry refinement using synthetic data. The geometry of the hand is initialized at various distances from ground truth as shown in the  $x$ -axis. The  $y$ -axes show how the mean per-vertex error ( $\epsilon_v$ ) and the backprojection error ( $\epsilon_b$ ) after refinement exhibit similar behavior. Please refer to Section 4.1 for the symbol definitions.

by adding a random Gaussian vector to the MANO pose and shape parameters of the selected gestures and hand shapes. We defined the noise level as the mean per-vertex distance between the initial and perturbed hand geometry. The hand textures were generated using HTML [45] approximating the hand textures observed in InterHand2.6M. Examples of our synthetic data are shown in Figure 3 on the left. On the right, we show the mean per-vertex error of the refined hand geometry and the corresponding backprojection error for all subjects and hand poses. The  $x$  axis denotes the noise level. Note that we also considered noise levels beyond the average levels estimated in InterHand2.6M (5mm). We observe a monotonic relation between the two types of error, *i.e.* the error on the estimated 3D geometry decreases as the backprojection error decreases. Similarly, we expect that decreasing the backprojection error on all views on real data will be accompanied by improvement in the actual 3D estimation.

## 4.2.2 Ray tracing vs rasterization

For ray tracing we use the differentiable Monte Carlo ray tracer by Li et al. [24]. As a rasterizer to compare with, we use the hard rasterization variant of Pytorch3D [46], which is equivalent to Neural Rendering [20], and has been commonly employed in the relevant literature, in its original form or in slight variations.

A few important notes are the following. The ray tracer by Li et al. [24], apart from implementing straightforward ray tracing and gradient computation, provides a continuous sampling solution for the edges, which in practice means that the basin of convergence for our backprojection problem is effectively widened. Other than that, the only hyperparameter of the employed ray tracer is the number of samples per pixel, which increases fidelity, *i.e.* the approximation of the true integrals involved in the rendering equation. Hard rasterization does not have provision for discontinuities and widening of the convergence basin, but soft rasterization [27] does, by means of per pixel aggregation of neighboring information. The way soft rasterization is employed in optimization is that the “softness”, controlled by a few parameters ( $\gamma, \sigma$ , blur amount, number of faces per pixel, etc.), is scheduled in a coarse-to-fine fashion, across iterations [26]. This means that at the last iterations, or when close to the sought solution, hard rasterization is effectively used. Since what is being discussed resides in this ballpark, we compare to hard rasterization.

To compare between the two rendering methodologies we define the same, common problem, of refining a pose given calibrated multicamera RGB input. For the comparison everything remains the same, apart from the differentiable renderer itself. We define initial

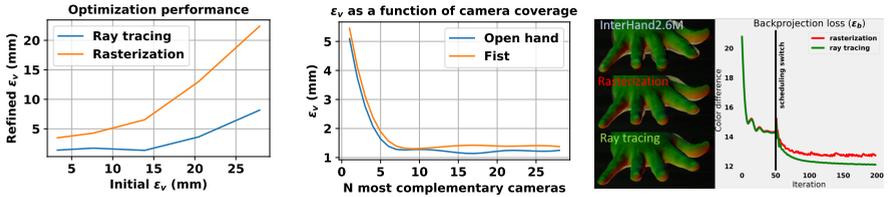


Figure 4: **(left)** Comparison between ray tracing and rasterization in optimization. **(center)**  $\epsilon_v$  as a function of optical coverage for two diverse scenarios. We observe that 7 cameras are enough to get most of the benefit. **(right)** The superiority of ray tracing [24] to rasterization [46] in this task, all else being equal, is also evident in real data.

search locations that are increasingly further from the ground truth and try, through optimization, to recover the ground truth itself. The results shown in Figure 4 indicate that ray tracing is significantly more effective in the task described, with the advantage of simplicity and the disadvantage of computational cost, compared to rasterization.

### 4.2.3 Camera coverage

We evaluated our method by varying the number of camera views used as input. We noticed that the proposed method does not require all 32 cameras in InterHand2.6M in order to incur a significant improvement on estimating the hand geometry. However, complementarity in coverage among the camera views was important, as expected. To demonstrate this we ordered the cameras in a new enumeration  $C_i, i = 1, \dots, N$ , such that

$$\text{comp}(C_{1,\dots,i}) \leq \text{comp}(C_{1,\dots,i+j}), \forall i, j > 0, \quad (6)$$

where  $\text{comp}(C_{1,\dots,i})$  computes the complementarity for cameras 1 to  $i$ , by aggregating estimates of how much of the surface of the hand becomes visible, on average. Put simply, we reordered the cameras such that every next camera incorporation maximizes the complementarity at every step. Given synthetic data as in Figure 4, we initialized our optimization from poses that had  $\epsilon_v = 5\text{mm}$ . Then, we performed our optimization in runs, where each run added the next camera, in the order described earlier. Among all results we picked two diverse instances to summarize the effect of camera coverage. The first instance regards an open hand and the other regards a fist. Both instances vary with respect to occlusions and ambiguity in observation. The results shown in Figure 4 indicate that, although more cameras are beneficial, most of the benefits can be gained by having 7 complementary cameras.

## 4.3 Qualitative analysis on real data

We evaluated our method on a subset of InterHand2.6M [42]. InterHand2.6M is a large-scale dataset with real multiview images depicting hands of various subjects performing sign language and everyday life gestures. The dataset provides annotations on 3D joint locations using a semi-automatic neural annotation method [31]. Given the 3D joint locations, pose and shape parameters of the MANO hand model were also extracted on the whole dataset with an average joint location error of 5mm [42]. However, when it comes to fine prediction of hand geometry, the provided annotations can mainly serve as references for the true location and geometry of the hand rather than ground truth. In this section, we provide qualitative evidence of our method improving the bundled annotations in InterHand2.6M. We ran

our geometry refinement method on a subset of InterHand2.6M consisting of 7 subjects (6 male, 1 female) and a diverse set of poses, ranging from a highly self-occluded hand to a completely open hand. An illustration of the predicted hand geometry for various subjects and poses is shown in Figure 5. The hand is colored based on the surface normals. It is worth noting that the comparative evaluation which validated our design choice of employing [24] (see Section 4.2.2) evidently generalizes to the case of real data, too (see Figure 4). More qualitative results are provided in the supplementary material accompanying this paper.

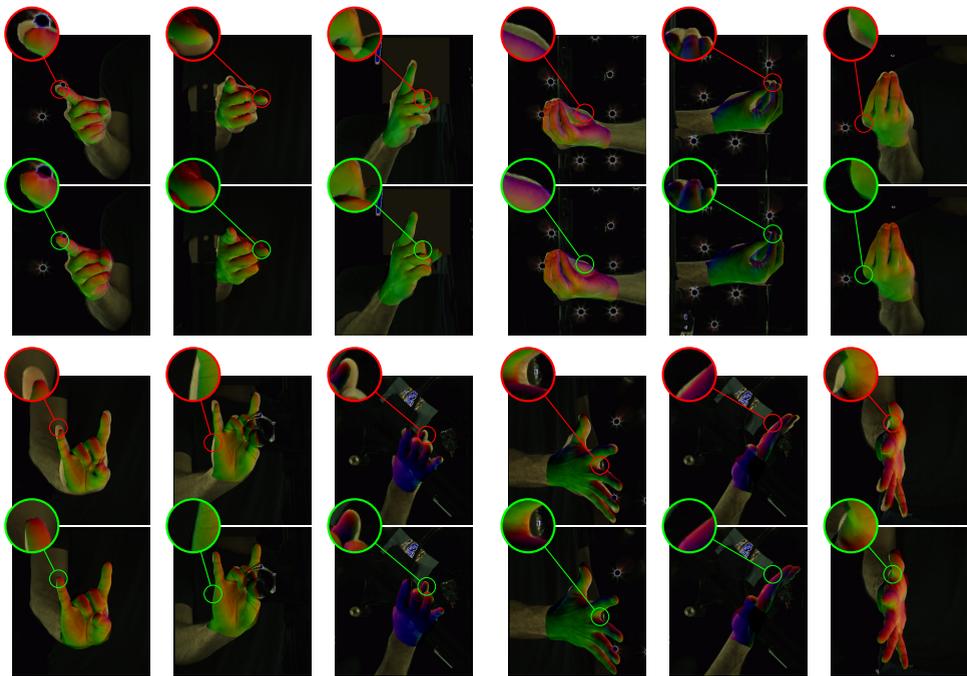


Figure 5: Example results of our method using real data from InterHand2.6M. Odd rows show the hand annotations in InterHand2.6M based on MANO. Even rows show the refined hand geometry with our method. The proposed approach can handle a variety of hand poses and is robust to moderate hand occlusions.

## 5 Discussion

We proposed a method for multiview refinement of the geometry of a hand, based on the assumption of color consistency among views and the employment of a powerful ray tracer. We presented a quantified analysis of the benefits of our method using a series of synthetic experiments and showed qualitatively that we improved upon the ground truth annotations of InterHand2.6M, a state-of-the-art dataset of real images with hands. What has been discussed so far has a straightforward impact on refining the output or the training procedure of supervised methods for hand geometry estimation. However, we find most exciting the fact that differentiable ray tracing, as employed here, for the first time, in the problem of hand pose/shape estimation, may open the door to self-supervised training and training in the wild. A differentiable ray tracer such as the one in [24] presents the best chance at properly disentangling image formation, which is key to generalization, and, in turn, self-supervision. The latter constitutes the basis of our future work in this direction.

## Acknowledgements

We gratefully acknowledge the support of NVIDIA Corporation with the donation of a Titan V GPU used for the execution of this research. The research project was supported by the Hellenic Foundation for Research and Innovation (H.F.R.I) under the "1<sup>st</sup> Call for H.F.R.I Research Projects to support Faculty members and Researchers and the procurement of high-cost research equipment" Project I.C.Humans, Number: 91.

## 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 1067–1076, 2019.
- [2] Federica Bogo, Javier Romero, Matthew Loper, and Michael J Black. Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3794–3801, 2014.
- [3] Adnane Boukhayma, Rodrigo de Bem, and Philip HS Torr. 3d hand shape and pose from images in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10843–10852, 2019.
- [4] 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.
- [5] Yujun Cai, Lihao Ge, Jianfei Cai, Nadia Magnenat-Thalmann, and Junsong Yuan. 3d hand pose estimation using synthetic data and weakly labeled rgb images. *IEEE transactions on pattern analysis and machine intelligence*, 2020.
- [6] John Canny. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, pages 679–698, 1986.
- [7] Yujin Chen, Zhigang Tu, Di Kang, Linchao Bao, Ying Zhang, Xuefei Zhe, Ruizhi Chen, and Junsong Yuan. Model-based 3d hand reconstruction via self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10451–10460, 2021.
- [8] Martin de La Gorce, Nikos Paragios, and David J Fleet. Model-based hand tracking with texture, shading and self-occlusions. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, 2008.
- [9] Ali Erol, George Bebis, Mircea Nicolescu, Richard D Boyle, and Xander Twombly. Vision-based hand pose estimation: A review. *Computer Vision and Image Understanding*, 108(1-2):52–73, 2007.
- [10] Yafei Gao, Yida Wang, Pietro Falco, Nassir Navab, and Federico Tombari. Variational object-aware 3-d hand pose from a single rgb image. *IEEE Robotics and Automation Letters*, 4(4):4239–4246, 2019.

- [11] Liuhao Ge, Hui Liang, Junsong Yuan, and Daniel Thalmann. Robust 3d hand pose estimation from single depth images using multi-view cnns. *IEEE Transactions on Image Processing*, 27(9):4422–4436, 2018.
- [12] Liuhao 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 10833–10842, 2019.
- [13] Andrew S Glassner. *An introduction to ray tracing*. Morgan Kaufmann, 1989.
- [14] Shreyas Hampali, Sayan Deb Sarkar, Mahdi Rad, and Vincent Lepetit. Handsformer: Keypoint transformer for monocular 3d pose estimation of hands and object in interaction. *arXiv preprint arXiv:2104.14639*, 2021.
- [15] Shangchen Han, Beibei Liu, Robert Wang, Yuting Ye, Christopher D Twigg, and Kenrick Kin. Online optical marker-based hand tracking with deep labels. *ACM Transactions on Graphics (TOG)*, 2018.
- [16] Shangchen Han, Beibei Liu, Randi Cabezas, Christopher D Twigg, Peizhao Zhang, Jeff Petkau, Tsz-Ho Yu, Chun-Jung Tai, Muzaffer Akbay, Zheng Wang, et al. Megatrack: monochrome egocentric articulated hand-tracking for virtual reality. *ACM Transactions on Graphics (TOG)*, 2020.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [18] Umar Iqbal, Pavlo Molchanov, Thomas Breuel Juergen Gall, and Jan Kautz. Hand pose estimation via latent 2.5d heatmap regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 118–134, 2018.
- [19] James T Kajiya. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, pages 143–150, 1986.
- [20] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [21] Cem Keskin, Furkan Kıraç, Yunus Emre Kara, and Lale Akarun. Real time hand pose estimation using depth sensors. In *Consumer depth cameras for computer vision*, pages 119–137. Springer, 2013.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [23] Dominik Kulon, Riza Alp Guler, Iasonas Kokkinos, Michael M Bronstein, and Stefanos Zafeiriou. Weakly-supervised mesh-convolutional hand reconstruction in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4990–5000, 2020.
- [24] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable monte carlo ray tracing through edge sampling. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 2018.

- [25] Yuncheng Li, Zehao Xue, Yingying Wang, Liuhao Ge, Zhou Ren, and Jonathan Rodriguez. End-to-end 3d hand pose estimation from stereo cameras. In *BMVC*, volume 1, page 2, 2019.
- [26] Lingjie Liu, Weipeng Xu, Marc Habermann, Michael Zollhoefer, Florian Bernard, Hyeonwoo Kim, Wenping Wang, and Christian Theobalt. Neural human video rendering by learning dynamic textures and rendering-to-video translation. *IEEE Transactions on Visualization and Computer Graphics*, 2020.
- [27] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. A general differentiable mesh renderer for image-based 3d reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [28] Matthew M Loper and Michael J Black. Opendr: An approximate differentiable renderer. In *European Conference on Computer Vision*, pages 154–169. Springer, 2014.
- [29] Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. Reparameterizing discontinuous integrands for differentiable rendering. *ACM Transactions on Graphics (TOG)*, 38(6):1–14, 2019.
- [30] Alexandros Makris and A Argyros. Model-based 3d hand tracking with on-line hand shape adaptation. In *Proc. BMVC*, pages 77–1, 2015.
- [31] Gyeongsik Moon and Kyoung Mu Lee. Neuralannot: Neural annotator for in-the-wild expressive 3d human pose and mesh training sets. *arXiv preprint arXiv:2011.11232*, 2020.
- [32] 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.
- [33] Gyeongsik Moon, Takaaki Shiratori, and Kyoung Mu Lee. Deephandmesh: A weakly-supervised deep encoder-decoder framework for high-fidelity hand mesh modeling. In *European Conference on Computer Vision (ECCV)*, 2020.
- [34] Gyeongsik Moon, Shoou-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.
- [35] 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.
- [36] Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. Mitsuba 2: A retargetable forward and inverse renderer. *ACM Transactions on Graphics (TOG)*, 38(6):1–17, 2019.
- [37] Markus Oberweger, Paul Wohlhart, and Vincent Lepetit. Training a feedback loop for hand pose estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 3316–3324, 2015.

- [38] Iason Oikonomidis, Nikolaos Kyriazis, and Antonis A Argyros. Efficient model-based 3d tracking of hand articulations using kinect. In *BMVC*, volume 1, page 3, 2011.
- [39] Iason Oikonomidis, Nikolaos Kyriazis, and Antonis A Argyros. Full dof tracking of a hand interacting with an object by modeling occlusions and physical constraints. In *2011 International Conference on Computer Vision*, pages 2088–2095. IEEE, 2011.
- [40] Iasonas Oikonomidis, Nikolaos Kyriazis, and Antonis A Argyros. Markerless and efficient 26-dof hand pose recovery. In *Asian Conference on Computer Vision*, pages 744–757. Springer, 2010.
- [41] Paschalis Panteleris and Antonis Argyros. Back to rgb: 3d tracking of hands and hand-object interactions based on short-baseline stereo. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 575–584, 2017.
- [42] 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.
- [43] Georgios Pavlakos, Nikos Kolotouros, and Kostas Daniilidis. Texturepose: Supervising human mesh estimation with texture consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 803–812, 2019.
- [44] Chen Qian, Xiao Sun, Yichen Wei, Xiaoou Tang, and Jian Sun. Realtime and robust hand tracking from depth. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1106–1113, 2014.
- [45] Neng Qian, Jiayi Wang, Franziska Mueller, Florian Bernard, Vladislav Golyanik, and Christian Theobalt. HTML: A Parametric Hand Texture Model for 3D Hand Reconstruction and Personalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2020.
- [46] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020.
- [47] James M Rehg and Takeo Kanade. Digiteyes: Vision-based hand tracking for human-computer interaction. In *Proceedings of 1994 IEEE Workshop on Motion of Non-rigid and Articulated Objects*, pages 16–22. IEEE, 1994.
- [48] Rilwan Remilekun Basaru, Greg Slabaugh, Eduardo Alonso, and Chris Child. Hand pose estimation using deep stereovision and markov-chain monte carlo. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 595–603, 2017.
- [49] Edgar Riba, Dmytro Mishkin, Daniel Ponsa, Ethan Rublee, and Gary Bradski. Kornia: an open source differentiable computer vision library for pytorch. In *Winter Conference on Applications of Computer Vision*, 2020. URL <https://arxiv.org/pdf/1910.02190.pdf>.
- [50] Javier Romero, Hedvig Kjellström, and Danica Kragic. Monocular real-time 3d articulated hand pose estimation. In *2009 9th IEEE-RAS International Conference on Humanoid Robots*, pages 87–92. IEEE, 2009.

- [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)*, 2017.
- [52] Viktor Rudnev, Vladislav Golyanik, Jiayi Wang, Hans-Peter Seidel, Franziska Mueller, Mohamed Elgharib, and Christian Theobalt. Eventhands: Real-time neural 3d hand reconstruction from an event stream. *arXiv preprint arXiv:2012.06475*, 2020.
- [53] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [54] Breannan Smith, Chenglei Wu, He Wen, Patrick Peluse, Yaser Sheikh, Jessica K Hodgins, and Takaaki Shiratori. Constraining dense hand surface tracking with elasticity. *ACM Transactions on Graphics (TOG)*, 2020.
- [55] Adrian Spurr, Aneesh Dahiya, Xucong Zhang, Xi Wang, and Otmar Hilliges. Self-supervised 3d hand pose estimation from monocular rgb via contrastive learning. *arXiv preprint arXiv:2106.05953*, 2021.
- [56] Adrian Spurr, Pavlo Molchanov, Umar Iqbal, Jan Kautz, and Otmar Hilliges. Adversarial motion modelling helps semi-supervised hand pose estimation. *arXiv preprint arXiv:2106.05954*, 2021.
- [57] Srinath Sridhar, Antti Oulasvirta, and Christian Theobalt. Interactive markerless articulated hand motion tracking using rgb and depth data. In *Proceedings of the IEEE international conference on computer vision*, pages 2456–2463, 2013.
- [58] Andrea Tagliasacchi, Matthias Schröder, Anastasia Tkach, Sofien Bouaziz, Mario Botsch, and Mark Pauly. Robust articulated-icp for real-time hand tracking. In *Computer Graphics Forum*, volume 34, pages 101–114. Wiley Online Library, 2015.
- [59] Danhang Tang, Hyung Jin Chang, Alykhan Tejani, and Tae-Kyun Kim. Latent regression forest: Structured estimation of 3d articulated hand posture. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3786–3793, 2014.
- [60] Peter Thompson and Aphrodite Galata. Hand tracking from monocular rgb with dense semantic labels. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 394–401. IEEE, 2020.
- [61] Anastasia Tkach, Andrea Tagliasacchi, Edoardo Remelli, Mark Pauly, and Andrew Fitzgibbon. Online generative model personalization for hand tracking. *ACM Transactions on Graphics (ToG)*, 36(6):1–11, 2017.
- [62] Dimitrios Tzionas, Luca Ballan, Abhilash Srikantha, Pablo Aponte, Marc Pollefeys, and Juergen Gall. Capturing hands in action using discriminative salient points and physics simulation. *International Journal of Computer Vision*, 118(2):172–193, 2016.
- [63] Jan Wöhlke, Shile Li, and Dongheui Lee. Model-based hand pose estimation for generalized hand shape with appearance normalization. *arXiv preprint arXiv:1807.00898*, 2018.

- [64] Christian Zimmermann and Thomas Brox. Learning to estimate 3d hand pose from single rgb images. Technical report, arXiv:1705.01389, 2017. URL <https://lmb.informatik.uni-freiburg.de/projects/hand3d/>. <https://arxiv.org/abs/1705.01389>.
- [65] Christian Zimmermann, Max Argus, and Thomas Brox. Contrastive representation learning for hand shape estimation. *arXiv preprint arXiv:2106.04324*, 2021.