

Mask2Hand: Learning to Predict the 3D Hand Pose and Shape from Shadow

Li-Jen Chang*, Yu-Cheng Liao*, Chia-Hui Lin*, Shih-Fang Yang-Mao[†] and Hwann-Tzong Chen*

* National Tsing Hua University, Taiwan

E-mail: htchen@cs.nthu.edu.tw

[†] Electronic and Optoelectronic System Research Laboratories, Industrial Technology Research Institute, Hsinchu, Taiwan.

E-mail: YangMao@itri.org.tw

Abstract—We present a self-trainable method, Mask2Hand, which learns to solve the challenging task of predicting 3D hand pose and shape from a 2D binary mask of hand silhouette/shadow without additional manually-annotated data. Given the intrinsic camera parameters and the parametric hand model in the camera space, we adopt the differentiable rendering technique to project 3D estimations onto the 2D binary silhouette space. By applying a tailored combination of losses between the rendered silhouette and the input binary mask, we are able to integrate the self-guidance mechanism into our end-to-end optimization process for constraining global mesh registration and hand pose estimation. The experiments show that our method, which takes a single binary mask as the input, can achieve comparable prediction accuracy on both unaligned and aligned settings as state-of-the-art methods that require RGB or depth inputs.

I. INTRODUCTION

More than hundreds of years ago the historic art form of hand shadow puppetry had already been documented in many ancient countries. Since then, making hand shadow puppets has been a low-cost household entertainment activity that parents can play with their children. As far as human perception is concerned, it is interesting that humans can recognize the animals and characters portrayed by the hands and fingers from merely their projected shadows. It is also remarkable that people are able to develop and create different ways of making hand shadow puppets, which may involve the generative capability of imagining some target shape and adjusting the hands and fingers to match the imaginary shape with the shadow. Inspired by the discriminative and generative nature in the visual process of doing hand shadow puppets, we aim to explore the possibility of learning to estimate the 3D hand pose and shape from an input hand silhouette, through a similar trial-and-error manner of learning to do hand shadow puppets.

Our main goal in this endeavor is to train a renderer-equipped deep network that takes a binary mask of a hand's silhouette as the input and predicts the corresponding 3D hand shape that produces a silhouette very similar to the input. (Fig. 1 shows several example results obtained using our model.) To build such a model, we employ MANO [1] as the 3D hand-shape renderer in our deep network. We propose a new network, called Mask2Hand, which consists of an encoder, a MANO layer, a differentiable render, and the refinement module. Fig. 2 shows an overview of the proposed Mask2Hand

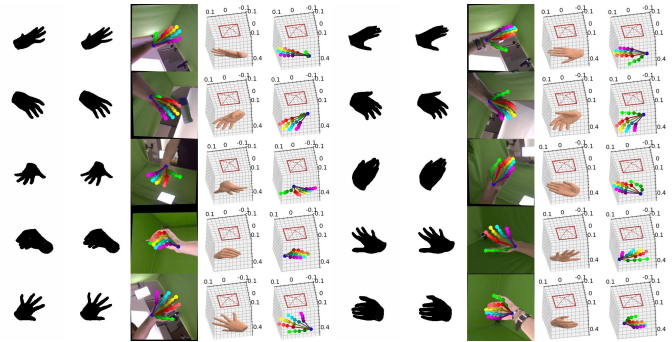


Fig. 1. The proposed method, Mask2Hand, learns to solve the challenging task of predicting the 3D hand pose and shape from a 2D binary mask of hand silhouette. Each row in the left or right half of the figure illustrates an example result of performing our method on real data. The first column shows the input of 2D binary mask. The second column shows the rendered mask using the predicted hand shape. The third column shows the predicted 2D skeleton depicted in the image space. Note that the color images are simply for visualization; our method does not use any RGB data. The last two columns show the predicted mesh and 3D skeleton.

network. Our network takes a binary mask as the input and uses the encoder to generate the required parameters for the MANO layer to render the hand mesh. The encoder learns to predict the pose-related principal components of MANO and the 3D global transformation of the hand from only a single input binary mask. Based on the principal components and global transformation, the MANO layer reconstructs the 3D hand mesh and the hand joint positions. We then use a differentiable renderer to generate the 2D projected silhouette from the hand mesh. Finally, the refinement module compares the projected silhouette with the input binary mask to refine the entire network. Three groups of loss functions are tailored for training the aforementioned modules, including the pose loss, silhouette loss, and mesh loss. The loss functions and the operations involved in the modules are all differentiable—we are able to train the entire network from end to end. Our code can be found at GitHub by searching Mask2Hand.

As this work aims at addressing a rather new and challenging task of predicting the 3D hand pose from a single 2D binary mask, it is quite encouraging that the proposed Mask2Hand network can achieve comparable prediction accuracy on 3D hand pose and shape estimation as state-of-the-art methods, although we only need a single binary input

image while theirs require an RGB image or a depth map. Note that our model can be trained under supervision with manually annotated joint positions or in a fully self-supervised manner without human annotations. Furthermore, since our method learns to predict the global transformation of the hand, its evaluation performance under the unaligned setting of hand pose estimation does not degrade as much as other methods do. Such a property is particularly favorable as, in real applications, it is not allowed to align the prediction with the unknown ground truth.

In addition to the motivation of hand shadow puppets, there are some other advantages and real-world applications of working on binary silhouettes. First, in contrast to RGB input, binary images can be obtained with relatively inexpensive sensors, such as low-resolution range sensors and low-resolution infrared sensors. They are especially useful when the lighting condition is poor and not much color information is available. Second, a line of research [2]–[4] in the hardware domain explores the usage of ambient light as a ubiquitous sensing medium. Light sensors, such as photodiodes or solar cells, measure and track the shadow blockage that a user’s hands cause with low cost and ultra-low power consumption, which have the potential to be widely used in IoT devices for shadow-based hand pose estimation systems. Finally, binary input is more robust to adversarial attacks of imperceptible pixel perturbation. From the literature on adversarial machine learning [5], we know that invisible image manipulation can seriously spoil the functionality of DNNs. However, this common type of attack is not directly applicable to the binary setting since small perturbation in pixel values is impossible with binary images. Also, pixel inversion in binary silhouettes can be easily identified and addressed.

II. RELATED WORK

One of the research endeavors inspiring us is *Shadow Theatre* [6], which aims to reproduce computationally the shadow puppet theatre—an ancient shadow-oriented form of performance art—just like how we are motivated by hand shadow puppetry. Won and Lee [6] develop an algorithm to incorporate several features and heuristic rules that are related to shadow generation using human bodies. The proposed method optimizes the poses of human bodies to match the given 2D target shapes. The underlying nonlinear high-dimensional optimization is solved using the proposed heuristic strategies so that a plausible solution can be found with reasonable computing resources. Nevertheless, the computation still requires tens to hundreds of minutes to complete on GPU.

A. 3D Hand Pose and Shape Estimation

Various methods have been proposed to address the problem of 3D hand pose and shape estimation. Generally, the techniques that are used to solve 3D human pose estimation can also be applied to 3D hand pose estimation. Based on the input format, we consider the following four categories of methods on 3D hand pose and shape estimation: RGB-based, depth-based, event-camera-based, and silhouette-based

methods. Among them, silhouette-based methods are most closely related to this work, where we seek to generate the 3D hand pose and mesh from merely a single binary-valued image of hand silhouette.

Despite the discrepancy in their input type, many of the methods take 2D joint detection as the basis. Such an approach is promising, and many existing deep-learning-based techniques can be employed. The following is a brief review of recent methods in the four categories.

a) RGB-based Methods.: Lin *et al.* propose a method called Mesh Transformer (METRO) [7] that explores non-local relationships between vertices and joints to reconstruct a 3D human pose from a single RGB image without using parametric mesh models. This method can also be extended to hand mesh reconstruction. Chen *et al.* [8] propose Camera-space Mesh Recovery (CMR) that generates 3D meshes from a single RGB image by extracting 2D cues including joint landmarks and silhouette. Chen *et al.* [9] present a self-supervised 3D hand reconstruction network called S²HAND that estimates 3D poses from 2D detected keypoints. Liu *et al.* [10] propose a semi-supervised 3D hand-object pose estimation method leveraging spatial-temporal consistency in videos to obtain pseudo-labels for self-training. Li *et al.* [11] propose a hand pose estimation method by using multi-task learning to categorize joints into groups so that different features can be learned to recover the 3D joint locations in a group-wise manner.

b) Depth-based Methods.: Xiong *et al.* [12] propose the Anchor-to-Joint Regression Network (A2J) that exploits informative anchor points to enhance the generalization ability of 3D pose estimation. Fang *et al.* [13] present a dense-prediction-based 3D hand pose estimation method called JGR-P2O that uses a pixel-to-offset prediction network and a joint graph reasoning module to enable end-to-end training and improve computational efficiency.

c) Silhouette-based Methods.: Silhouette-based methods have not yet been widely applied to hand/body pose estimation. The task to be solved is more challenging than those in other settings of RGB-based and depth-based methods. The difficulty comes from the lack of information in the input data, where the shading, color, and depth of the 3D objects are absent in the binary-valued images. Lee *et al.* present the Silhouette-Net [14] method that generates 3D hand poses based on binary-valued silhouettes and can achieve a performance similar to or better than depth-based methods if multiview silhouettes are provided. However, their method requires the guidance of depth maps during the training stage. In contrast, our method does not require additional depth information during training, and it only needs a single-view 2D binary mask of the hand as the input. Despite the challenging setting, estimating the 3D hand pose and mesh from limited binary information in the input is our major goal in this work.

B. 3D Hand Models

Several hand mesh models have been proposed, such as MANO [1], Sphere-Meshes [15], and Convex Parts [16]. With

the help of these models, hand shape generation and pose estimation can be processed via parametric methods. Typically, a hand model like MANO is designed to parameterize a triangle mesh into pose, shape, and rotation parameters. MANO has been widely used by hand pose estimation methods [17]–[26] to train their deep networks; the deep networks learn to predict the hand shape in the target image by regressing to the principal components of the MANO model. As shown in Fig. 2, our method also uses the MANO model to construct the target hand shape. By learning to optimize the parameters of MANO meshes, our method achieves the goal of generating the 3D hand mesh from a binary-valued silhouette image.

III. OUR APPROACH

To predict the hand pose and shape from a binary mask, we propose a two-phase end-to-end trainable network, *Mask2Hand*, that aims to reconstruct 3D hand meshes in the camera space. Fig. 2 illustrates the overall pipeline of our method. In what follows, we detail each module of the pipeline and describe how we formulate the loss functions.

A. Model Architecture

The core modules of our model include an encoder for estimating the MANO parameters, a MANO layer that reconstructs the 3D hand mesh, a differentiable renderer for optimizing the hand pose, and a RefineNet improving the accuracy.

a) Encoder.: Given a hand silhouette image presented as a binary mask, the encoder first uses a ResNet18 [27] backbone to extract features from the input. Then, these features are fed into four different estimators to encode 3D information into the PCA space of MANO [1] and predict the extrinsic camera parameters. More specifically, the output of the encoder consists of the first six or forty-five hand-pose-related principal components $\theta \in \mathbb{R}^6$ or \mathbb{R}^{45} of MANO, shape parameters $\beta \in \mathbb{R}^{10}$, 3D rotation $R \in \mathbb{R}^3$ in axis-angle representation, and 3D translation $T \in \mathbb{R}^3$.

b) MANO Layer.: The 3D hand shape is represented by a triangle mesh that contains a vertex set $V \in \mathbb{R}^{778 \times 3}$ and fixed faces F indicating the connection between them. In our network, the hand mesh and joints $J \in \mathbb{R}^{21 \times 3}$ are reconstructed using MANO [1]. Being a low-dimensional parametric hand model, MANO can generate the required hand mesh from parameters including the hand-pose-related principal components θ , shape $\beta \in \mathbb{R}^{10}$, rotation R , and translation T .

c) Differentiable Renderer.: Given intrinsic camera parameters and hand meshes in the camera space, we use the differentiable rendering technique [28] to project 3D estimations into the 2D image space. By applying various losses between the rendered mask of the hand silhouette and the input binary image, we can integrate self-supervision into our end-to-end optimization process for constraining global mesh registration. Here we use the PyTorch3D [29] implementation of the differentiable renderer. The loss functions will be detailed later in the subsequent sections.

d) RefineNet.: In the previous stage, we only consider several pose-related principal components of MANO for parametrizing hands in a low-dimensional space. The lower degree of freedom may have an expressiveness problem in that the finer details of hand gestures cannot be ideally reconstructed. Therefore, we incorporate a small convolutional neural network, called the RefineNet module¹, into our model to deal with this issue. The network takes the concatenation of the input image and the projected silhouette as input, and it produces residual coordinate values that serve as the point-wise offsets to the previously constructed mesh. The final output of our model is the summation of the originally reconstructed hand mesh and the offset predicted by the RefineNet. The refined joints are then estimated by MANO’s joint regressor given the refined mesh as its input.

B. Loss Functions

The complete objective function that we design for training our *Mask2Hand* model comprises three main loss terms:

$$\mathcal{L} = \mathcal{L}_{\text{pose}} + \mathcal{L}_{\text{silhouette}} + \mathcal{L}_{\text{mesh}}, \quad (1)$$

where we have the pose loss $\mathcal{L}_{\text{pose}}$ on the predicted joints, the silhouette loss $\mathcal{L}_{\text{silhouette}}$ on the 2D projected silhouette of the hand, and the mesh loss $\mathcal{L}_{\text{mesh}}$ on the reconstructed 3D mesh.

1) Pose Loss.: The pose loss in Eq. (1) further consists of two loss functions that use the unaligned and aligned joint positions to evaluate the preliminary and refined joint predictions by the squared L2-norm error:

$$\mathcal{L}_{\text{pose}} = \lambda_J \mathcal{L}_J + \lambda_{\text{align}J} \mathcal{L}_{\text{align}J}, \quad (2)$$

where we have \mathcal{L}_J for the unaligned joint evaluation and $\mathcal{L}_{\text{align}J}$ for the aligned joint evaluation, with λ_J and $\lambda_{\text{align}J}$ as the weight factors. The 3D positions of the K joints are represented as a K by 3 matrix, *i.e.*, $J \in \mathbb{R}^{K \times 3}$. The loss function \mathcal{L}_J of the unaligned joint evaluation is computed as

$$\mathcal{L}_J = \frac{1}{K} \sum_{i=1}^K \|J_i - \hat{J}_i\|_2^2 + \frac{1}{K} \sum_{i=1}^K \|J_i - \hat{J}_i^{\text{ref}}\|_2^2, \quad (3)$$

J_i is the ground-truth 3D position of joint i , \hat{J}_i is the preliminarily reconstructed position of joint i , and \hat{J}_i^{ref} is the final prediction after refinement.

For the aligned joint loss, the predicted joints are aligned with the ground truth using orthogonal Procrustes Analysis (PA) [31], and then the L2 loss is applied to the aligned and the ground-truth joints:

$$\mathcal{L}_{\text{align}J} = \frac{1}{K} \sum_{i=1}^K \|J_i - \tilde{J}_i\|_2^2 + \frac{1}{K} \sum_{i=1}^K \|J_i - \tilde{J}_i^{\text{ref}}\|_2^2, \quad (4)$$

where $\tilde{J} = \text{PA}(J, \hat{J})$ and $\tilde{J}^{\text{ref}} = \text{PA}(J, \hat{J}^{\text{ref}})$.

The procedure of alignment $\text{PA}(J, \hat{J})$ can be described as follows.

¹Our RefineNet module is different from the dense prediction network proposed in [30] despite the same name.

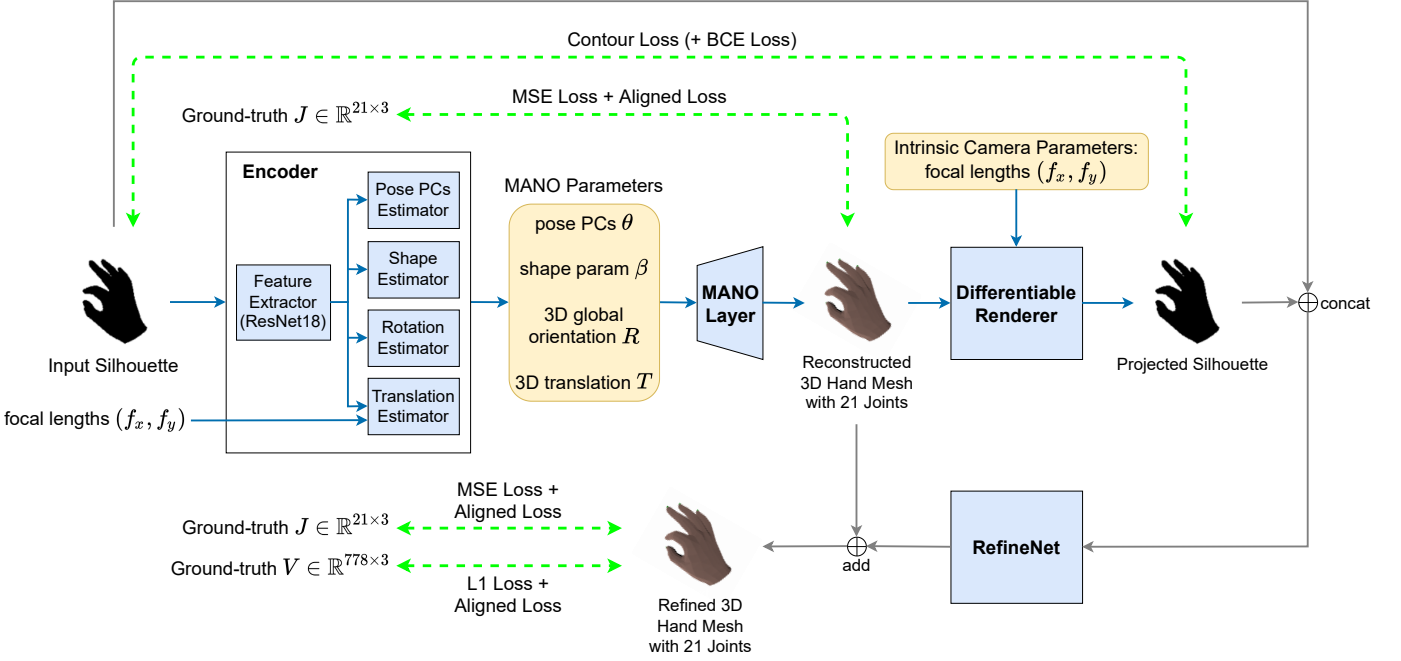


Fig. 2. The Mask2Hand network comprises the Encoder, MANO Layer, Differentiable Renderer, and RefineNet. The Encoder takes a binary mask as the input and generates the required parameters for the MANO Layer; the parameters include the shape, the six or forty-five pose-related principal components of MANO, the 3D global orientation, and translation of the hand. Based on these parameters, the MANO Layer reconstructs the 3D hand mesh (represented by 778 vertices) with the 21 positions of the hand joints. The Differentiable Renderer generates the 2D projected silhouette of the hand mesh, and the RefineNet can compare the projection with the input binary mask to refine the entire model. We design different loss functions for different modules in our model. The green dashed lines indicate what losses are applied to the associated modules. Please refer to the text for the detailed formulations of the loss functions.

- 1) Translate the joints so that their mean lies at the origin:

$$J \leftarrow J - \frac{1}{K} \sum_{i=1}^K J_i, \quad \hat{J} \leftarrow \hat{J} - \frac{1}{K} \sum_{i=1}^K \hat{J}_i.$$

- 2) Divide the joints set by its Frobenius norm to remove the uniform scaling component:

$$J \leftarrow J / \|J\|_F, \quad \hat{J} \leftarrow \hat{J} / \|\hat{J}\|_F.$$

- 3) Solve the orthogonal Procrustes problem: given matrices J and \hat{J} , find an orthogonal matrix $Q \in \mathbb{R}^{3 \times 3}$ that maps J most closely to \hat{J} , i.e., $\min_Q \|JQ - \hat{J}\|_F$ subject to $Q^\top Q = I_{3 \times 3}$. The solution consists of $Q = UV^\top$ and $s = \sum_i \sigma_i$, where $U\Sigma V^\top$ is the singular value decomposition of $J^\top \hat{J}$ and σ_i 's are the singular values.
- 4) $\tilde{J} = (\hat{J}Q^\top \times s) \times \|J\|_F + (\sum_{i=1}^K J_i) / K$, where the Procrustes output is scaled by the norm and shifted by the mean that we derive from J in the first two steps.

2) *Silhouette Loss.*: For the module of differentiable rendering, the ideal result is that the 2D projected hand silhouette matches the input binary mask exactly. Hence, our silhouette loss is based on the binary cross-entropy (BCE) and the contour loss:

$$\mathcal{L}_{\text{silhouette}} = \lambda_{\text{bce}} \text{BCE}(S, B) + \lambda_{\text{contour}} \text{ContourLoss}(S, B), \quad (5)$$

where S is the rendered silhouette and B is the input binary mask. Our contour loss is a differentiable estimation of Chamfer distance between two contours, calculated as a pixel-wise multiplication between the contour of rendered silhouette and the pre-computed distance field of the ground-truth silhouette's contour. The detailed steps for computing the Chamfer distance are summarized as follows:

- 1) Compute the distance transform [32] for the contour ∂B of the ground-truth binary mask B in the dataset to get the distance field $\Psi_{\partial B}$.
- 2) Apply differentiable binarization to the rendered silhouette image S . In our case, we convert a pixel value $x \in [0, 1]$ in S to an approximately binary value as $\frac{1}{1 + e^{-100(x-0.5)}} \simeq \{0, 1\}$.
- 3) Zero out the pixel values that are greater than the threshold 0.5 to deal with the noise near the wrist caused by PyTorch3D differentiable rendering. The result of Steps 2 & 3 is a clean and nearly binarized silhouette \bar{S} .
- 4) Apply Laplacian operator to \bar{S} to get $\Delta \bar{S}$. The contour $\partial \bar{S}$ is then obtained by $\partial \bar{S} = \tanh(\max\{\Delta \bar{S}, 0\})$, which maps the values on the contour to near 1 and all the rest to 0.
- 5) Do element-wise product between the resulting contour $\partial \bar{S}$ and the distance field $\Psi_{\partial B}$ to retrieve the corresponding distance values and then take the sum to get $\text{ContourLoss}(S, B)$.

All of the aforementioned computations are differentiable. Therefore, we can readily include the silhouette loss in our method, and end-to-end train the network.

3) *Mesh Loss.*: We employ the mesh loss to measure the error in predicting the 3D vertices of the hand mesh. Like the pose loss, the mesh loss is also evaluated on both the unaligned and aligned predictions:

$$\mathcal{L}_{\text{mesh}} = \lambda_V \mathcal{L}_V + \lambda_{\text{align}V} \mathcal{L}_{\text{align}V}. \quad (6)$$

The unaligned vertex loss \mathcal{L}_V consists of the L1 losses for the predictions from the preliminary and refined meshes:

$$\mathcal{L}_V = \frac{1}{M} \sum_{i=1}^M \|V_i - \hat{V}_i\|_1 + \frac{1}{M} \sum_{i=1}^M \|V_i - \hat{V}_i^{\text{ref}}\|_1, \quad (7)$$

where $V \in \mathbb{R}^{M \times 3}$ is the ground-truth 3D vertices of the hand mesh, \hat{V} is the preliminarily reconstructed mesh, and \hat{V}^{ref} is the final prediction of vertices after refinement.

For the aligned vertex loss $\mathcal{L}_{\text{align}V}$, the predicted vertices are aligned with the ground truth using orthogonal Procrustes analysis as described in the aligned pose loss Eq. (4), and then the L1 loss is applied to the aligned and the ground-truth mesh as

$$\mathcal{L}_{\text{align}V} = \frac{1}{M} \sum_{i=1}^M \|V_i - \tilde{V}_i\|_1 + \frac{1}{M} \sum_{i=1}^M \|V_i - \tilde{V}_i^{\text{ref}}\|_1, \quad (8)$$

where $\tilde{V} = \text{PA}(V, \hat{V})$ and $\tilde{V}^{\text{ref}} = \text{PA}(V, \hat{V}^{\text{ref}})$.

IV. EXPERIMENTS

Since the problem formulation of the task that we aim to address is different from those commonly adopted in prior work, we have not found state-of-the-art methods that take exactly the same input setting and could therefore be considered for direct comparison with our approach. Note that the work of Lee *et al.* [14] is only posted on arXiv as a preprint without available code. Therefore, we do not include their method in the following comparisons. Moreover, their method requires depth maps during training, which is more restrictive than ours. In our experiments, we first compare our method with the state-of-the-art depth-based model A2J [12] on a synthetic dataset and show that transferring our task for an existing model to solve is probably not as straightforward as it may seem. Further, we use the real data in FreiHAND to compare our method with the three state-of-the-art RGB-based methods: I2L-MeshNet [22], MeshTransformer [7], and CMR (ResNet18) [8]. The results show that our method performs particularly well under the unaligned evaluation setting of 3D hand pose estimation, which is more practical for real applications.

A. Datasets

We conduct experiments on the following two datasets.

a) *Synthetic Dataset.*: This dataset contains 20,000 training and 2,000 test binary images synthesized using MANO [1], [33]. For each sample in the dataset, its hand-pose principal components $\theta \in \mathbb{R}^6$ are randomly sampled from the uniform distribution between $[-2.0, 2.0]$, and each rotation angle of its 3D global orientation R is randomly sampled from the uniform distribution $[-\pi, \pi]$. We use MANO to generate the corresponding hand shapes from the sampled parameters. The ground-truth 3D coordinates of the 21 hand joints can thus be automatically derived from the MANO rendered hand shapes.

b) *FreiHAND* [34].: It is a common real-world RGB dataset with 32,560 training samples and 3,960 test samples. We binarize the segmentation masks available in the training set and use them as the input to our model. However, since its original evaluation set does not provide segmentation masks, we resort to splitting its original training set into a random partition of 26,000 training, 3,280 validation, and 3,280 test images to achieve a reasonably fair evaluation of our method.

For quantitative evaluation, we report the standard metrics used by the FreiHAND dataset on the predictions of 3D joints and meshes:

- **MPJPE**: the *mean per joint position error*, which measures the Euclidean distance (in cm) between the ground-truth joints and the predicted joints.
- **MPVPE**: the *mean per-vertex position error*, which measures the Euclidean distance (in cm) between the ground-truth vertices and the predicted vertices of a mesh.
- **AUC of PCK**: the *area under the curve* of the *percentage of correct keypoints*, which is plotted using 100 equally-spaced thresholds between 0 cm to 5 cm.
- **AUC of PCV**: the *area under the curve* of the *percentage of correct vertices*.
- **PA-MPJPE** or **PA-MPVPE**: It first applies Procrustes alignment between the ground truth and the prediction, and then calculates MPJPE or MPVPE. This metric aims to measure reconstruction error that neglects the effect of global rotation, translation, and scaling.
- **mIoU** and **Dice coefficient**: In addition to the preceding metrics used by the FreiHAND dataset for the evaluation of joint and mesh predictions, we also use the *mean intersection over union* and the *Dice coefficient* to evaluate the similarity between the input binary mask and the mask projected from the predicted hand mesh.

TABLE I
COMPARISON WITH THE DEPTH-BASED MODEL A2J [12] ON THE
SYNTHETIC DATASET. (PA MEANS PROCRUSTES ALIGNED)

Method	Input	MPJPE ↓	AUC _J ↑	PA-MPJPE	PA-AUC _J
A2J [12]	Mean depth + 5-view aug.	7.73 cm	0.35	1.93 cm	0.62
Ours	Binary (0, 255) + 5-view aug.	0.87 cm	0.83	0.53 cm	0.89

B. Comparison with State-of-the-Art Depth-based Methods

To demonstrate that state-of-the-art depth-based models cannot be easily transferred to the task with binary input, we

TABLE II
FREIHAND (Not Aligned). NOTE THAT THE mIoU AND THE DICE COEFFICIENT ARE ONLY RELEVANT TO THE BINARY INPUT CASE.

Method	Input	MPJPE (cm)↓	PCK↑	MPVPE (cm)↓	PCV↑	mIoU↑	Dice↑
Mesh-Transformer [7]	RGB	68.59	0.00	68.58	0.00	—	—
	Binary	68.59	0.00	68.59	0.00	0.20	0.12
CMR (ResNet18) [8]	RGB	3.49	0.42	3.49	0.42	—	—
	Binary	4.31	0.35	4.31	0.35	0.90	0.90
Ours (6 PCs)	Binary	3.56	0.40	3.57	0.40	0.88	0.87
Ours (45 PCs)		3.56	0.41	3.57	0.41	0.88	0.88

TABLE III
FREIHAND (Procrustes Aligned). THE ADDITIONAL ENTRIES OF PA-MPJPE AND PA-MPVPE OF I2L-MESHNET WITH RGB INPUT ARE COPIED FROM THE ORIGINAL PAPER [22] AND POSTED HERE FOR THE EASE OF REFERENCE, WHICH CAN BE CONSIDERED AS A STRONG BASELINE.

Method	Input	PA-MPJPE (cm)↓	AUC of PCK↑	PA-MPVPE (cm)↓	AUC of PCV↑
I2L-MeshNet [22]	RGB	0.74	—	0.76	—
Mesh-Transformer [7]	RGB	0.54	0.89	0.58	0.88
	Binary	0.60	0.88	0.64	0.87
CMR (ResNet18) [8]	RGB	0.67	0.87	0.68	0.87
	Binary	0.77	0.85	0.77	0.85
Ours (6 PCs)	Binary	0.73	0.85	0.76	0.85
Ours (45 PCs)		0.68	0.86	0.69	0.86

compare our model with A2J [12] on the synthetic dataset. The input image of our model is binarized to $\{0, 255\}$ while that of A2J is binarized to 0 (as the background) and the mean depth of joints to treat the depth-based model fairly. This way, in spite of the fact that the input image is binary-valued, we still provide sufficient statistics of depth information to A2J. As shown in Table I, for the unaligned setting, our model achieves MPJPE of 0.87 cm, and A2J reaches MPJPE of 7.73 cm. The results indicate that A2J fails to learn to localize 3D hand joints accurately when the input is binary, even if the reference mean depth is given. In contrast, our method has the promising ability to recover the positions of the joints in the camera space. In terms of MPJPE after Procrustes alignment, our approach also outperforms A2J by a large margin, as presented in Table I. Some other depth-based models are not quantitatively compared with ours for the following reasons. V2V-PoseNet [35] takes a 3D voxelized depth map as its input, which is expected to fail for silhouette-based input since the voxelization process cannot work for binary images. Hand PointNet [36] converts a depth map into a 3D point cloud and then uses it as the model’s input. Likewise, the conversion procedure cannot function properly for silhouettes, making it inadequate for the task of binary input.

C. Comparison with State-of-the-Art RGB-based Methods

For the comparison with state-of-the-art RGB-based methods, we select the leading models from the FreiHAND competition leaderboard at CodaLab ², including CMR [8] and MeshTransformer [7]. To make a fair comparison with our method, we use the officially released code to re-train their

²<https://competitions.codalab.org/competitions/21238>

TABLE IV
ABLATION STUDY. (MEANING OF * IN THE FIRST COLUMN: THE ALIGNED LOSS IS ONLY APPLIED TO THE REFINED JOINTS AND VERTICES. (A FOR ALIGNED LOSS, B FOR BCE LOSS, C FOR CONTOUR LOSS, R FOR REFINE NET, M FOR MESH LOSS.)

A	B	C	R	M	MPJPE (cm)↓	MPVPE (cm)↓	PA-MPJPE (cm)↓	PA-MPVPE (cm)↓	mIoU↑
✓	✓	✓	✓	✓	3.55	3.56	0.78	0.79	0.87
✓	✓	✓	✓	✓	3.67	3.84	0.86	1.18	0.76
✓	✓	✓	✓	✓	3.73	3.73	0.94	0.94	0.83
✓	✓	✓	✓	✓	3.60	3.61	0.80	0.81	0.87
✓	✓	✓	✓	✓	3.83	3.91	1.36	1.56	0.86
*	✓	✓	✓	✓	3.66	3.67	0.79	0.80	0.86
✓	✓	✓	✓	✓	3.56	3.57	0.78	0.79	0.87
✓	✓	✓	✓	✓	3.72	3.73	0.90	0.93	0.83
✓	✓	✓	✓	✓	3.66	3.67	1.04	1.06	0.88

models on the FreiHAND dataset randomly partitioned by us, as mentioned in Section IV-A, under the settings of both binary and RGB input. The former is for performance observation, while the latter is primarily for convincing proof that these models are not trained imperfectly. As presented in Table II, CMR achieves MPJPE and MPVPE of 4.31 cm with silhouette-based input, while MeshTransformer fails to localize 3D hand joints and mesh in the camera space. Our model reaches MPJPE of 3.56 cm and MPVPE of 3.57 cm, which outperforms the previous two methods by a large margin and is even comparable to the result of CMR with RGB input. The superior performance of our model in the unaligned case indicates that it has a better ability to reconstruct 3D hand mesh in the camera space, which is more practical in real applications. In terms of mIoU and Dice coefficient, our model achieves similar results as CMR.

When the predicted joints and vertices are aligned with the ground truth by Procrustes analysis, our approach reaches PA-MPJPE of 0.68 cm and PA-MPVPE of 0.69 cm. These experimental results are significantly better than CMR’s performance under the setting of binary input and are comparable to CMR’s results with RGB input. Furthermore, unlike MeshTransformer, which totally fails in the unaligned case, our method performs stably well in both cases. Such a balanced behavior is preferable in our task since localization and detail recovery are vitally important for projecting a silhouette similar to the input one.

D. Implementation Details

Our encoder’s backbone is based on ResNet18 [27]. Py-Torch3D’s implementation [29] is employed for our differentiable renderer. We use the Adam optimizer to train our network with a batch size of 32. The initial learning rate is set to 10^{-4} , and we use the ReduceOnPlateau scheduler to adjust the learning rate. The whole network is trained for 150 epochs, and the checkpoint with the lowest validation loss is selected for our final usage. For the FreiHAND dataset, we apply data augmentation of random rotation between $[-\pi, \pi)$ and scaling between $[0.9, 1.1)$ in the 2D image space. For the synthetic dataset, we randomly sample only one of the five views for each training data at every epoch to perform data

augmentation. We set $\lambda_J = 2 \times 10^{-3}$, $\lambda_{\text{align}J} = 2 \times 10^{-2}$, $\lambda_{\text{contour}} = 10^{-4}$, $\lambda_V = 0.1$, and $\lambda_{\text{align}V} = 1$ to balance each loss term. λ_{bce} is set to 0 for our final model and 0.5 for the ablation study.

E. Ablation Study

We evaluate different configurations of our method to analyze the effectiveness of each component and examine the effects of the loss functions.

1) *Effects of Shape Parameters.*: In our ablation study, we use six pose PCs and set β to the mean shape, *i.e.*, $\beta = \mathbf{0}$. By comparing the first row of Table IV and the penultimate row of Tables II & III, we can observe that regressing shape parameters in the first stage of our model helps reduce the burden of refinement and achieve better performance in the aligned case.

2) *Effectiveness of Each Component.*: As shown in Table IV, we explore the evaluation results on the FreiHAND dataset with the exclusion of different components. Our base model consists of Encoder, MANO Layer, and Differentiable Renderer. After incorporating the RefineNet module into our network with pose loss only, we observe that both MPJPE and PA-MPJPE decrease while MPVPE and PA-MPVPE undesirably increase. This result is in line with our expectation since the pose loss cannot impose strong enough constraints on the deformation of the preliminarily constructed hand mesh. Based on the observation, we further apply the mesh loss to the output of the RefineNet, which helps to strengthen the functionality of refinement and improve all evaluation metrics significantly.

To reveal the effects of other losses, we exclude them one at a time from our entire network. First, we remove the contour loss between the input binary image and the projected silhouette. The evaluation result shows that the model's performance on both hand pose and shape estimation degrades, especially the unaligned ones, demonstrating the effectiveness of the contour loss for achieving more accurate global mesh registration. Second, we eliminate the aligned losses of both joints and vertices. Our experimental result shows that MPJPE, MPVPE, PA-MPJPE, and PA-MPVPE increase by 0.28, 0.35, 0.58, and 0.77, respectively. Such a significant deterioration in performance indicates that the aligned losses not only play crucial roles in the reconstruction of fine-grained gestures but also help to improve mesh recovery in the camera space. Finally, we remove all losses inserted before RefineNet. The result shown in the sixth row of Table IV suggests that adding the pose loss and contour loss before the refinement stage can help our model improve the ability to localize the 3D hand mesh in the camera space.

3) *Effects of Adding the BCE Loss.*: We also do some experiments to see whether adding the binary cross-entropy loss between the input binary mask and the projected silhouette has any positive effect. As listed in Table IV, directly incorporating the BCE loss into our entire architecture contributes nothing to the final performance. To understand the root cause of this counterintuitive phenomenon, we take away the RefineNet module and the aligned loss one at a time with the presence

of BCE loss. The experimental results show that the BCE loss can enhance our model's performance by a large margin when the aligned losses are absent. This implies that the BCE loss does not have much effect when the model learns the shape estimation task to a certain degree guided by the stronger training objectives, *i.e.*, the aligned losses of joints and vertices.

V. CONCLUSION

This paper introduces the Mask2Hand network that learns to predict the 3D hand pose and shape from a 2D binary mask without relying on any RGB or depth information. Despite the more challenging input setting of our task in comparison with prior work on hand pose and shape estimation, we show that the proposed method can achieve comparable performance as state-of-the-art RGB-based and depth-based methods. With the parametric hand model and the differentiable rendering technique, we integrate the self-supervised mechanism into our end-to-end training process without the need for human annotations. We propose several loss functions to model different aspects of 3D hand pose and shape estimation. Since our method explicitly learns to predict the global transformation of the hand, its better performance in the unaligned setting is thus promising for real applications.

ACKNOWLEDGMENTS

This work was supported in part by Industrial Technology Research Institute, Taiwan. We are grateful to National Center for High-performance Computing for providing computational resources and facilities.

REFERENCES

- [1] J. Romero, D. Tzionas, and M. J. Black, "Embodied hands: Modeling and capturing hands and bodies together," *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 2017.
- [2] T. Li, X. Xiong, Y. Xie, G. Hito, X. Yang, and X. Zhou, "Reconstructing hand poses using visible light," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 1, no. 3, 71:1–71:20, 2017.
- [3] D. Ma, G. Lan, M. Hassan, *et al.*, "Solargest: Ubiquitous and battery-free gesture recognition using solar cells," in *The 25th Annual International Conference on Mobile Computing and Networking*, 2019, 12:1–12:15.
- [4] A. Varshney, A. Soleiman, L. Mottola, and T. Voigt, "Battery-free visible light sensing," in *Proceedings of the 4th ACM Workshop on Visible Light Communication Systems*, 2017, pp. 3–8.
- [5] N. Akhtar, A. Mian, N. Kardan, and M. Shah, "Advances in adversarial attacks and defenses in computer vision: A survey," *IEEE Access*, vol. 9, pp. 155 161–155 196, 2021.
- [6] J. Won and J. Lee, "Shadow theatre: Discovering human motion from a sequence of silhouettes," *ACM Trans. Graph.*, vol. 35, no. 4, 147:1–147:12, 2016.

- [7] K. Lin, L. Wang, and Z. Liu, “End-to-end human pose and mesh reconstruction with transformers,” in *CVPR*, 2021, pp. 1954–1963.
- [8] X. Chen, Y. Liu, C. Ma, *et al.*, “Camera-space hand mesh recovery via semantic aggregation and adaptive 2d-1d registration,” in *CVPR*, 2021, pp. 13 274–13 283.
- [9] Y. Chen, Z. Tu, D. Kang, *et al.*, “Model-based 3d hand reconstruction via self-supervised learning,” in *CVPR*, 2021, pp. 10 451–10 460.
- [10] S. Liu, H. Jiang, J. Xu, S. Liu, and X. Wang, “Semi-supervised 3d hand-object poses estimation with interactions in time,” in *CVPR*, 2021, pp. 14 687–14 697.
- [11] M. Li, Y. Gao, and N. Sang, “Exploiting learnable joint groups for hand pose estimation,” in *Thirty-Fifth AAAI Conference on Artificial Intelligence*, 2021, pp. 1921–1929.
- [12] F. Xiong, B. Zhang, Y. Xiao, *et al.*, “A2J: anchor-to-joint regression network for 3d articulated pose estimation from a single depth image,” in *IEEE International Conference on Computer Vision (ICCV)*, 2019, pp. 793–802.
- [13] L. Fang, X. Liu, L. Liu, H. Xu, and W. Kang, “JGR-P2O: joint graph reasoning based pixel-to-offset prediction network for 3d hand pose estimation from a single depth image,” in *European Conference on Computer Vision (ECCV)*, 2020.
- [14] K. Lee, S. Liu, H. Chen, and K. Ito, “Silhouette-net: 3d hand pose estimation from silhouettes,” *CoRR*, vol. abs/1912.12436, 2019.
- [15] A. Tkach, M. Pauly, and A. Tagliasacchi, “Spheremeshes for real-time hand modeling and tracking,” *ACM Trans. Graph.*, vol. 35, no. 6, pp. 222:1–222:11, 2016.
- [16] D. Tzionas, L. Ballan, A. Srikantha, P. Aponte, M. Pollefeys, and J. Gall, “Capturing hands in action using discriminative salient points and physics simulation,” *Int. J. Comput. Vis.*, vol. 118, no. 2, pp. 172–193, 2016.
- [17] A. Boukhayma, R. A. de Bem, and P. H. S. Torr, “3D hand shape and pose from images in the wild,” in *CVPR*, 2019, pp. 10 843–10 852.
- [18] S. Baek, K. I. Kim, and T. Kim, “Pushing the envelope for rgb-based dense 3d hand pose estimation via neural rendering,” in *CVPR*, 2019, pp. 1067–1076.
- [19] V. Choutas, G. Pavlakos, T. Bolkart, D. Tzionas, and M. J. Black, “Monocular expressive body regression through body-driven attention,” in *European Conference on Computer Vision (ECCV)*, 2020.
- [20] E. Corona, A. Pumarola, G. Alenyà, F. Moreno-Noguer, and G. Rogez, “GanHand: Predicting human grasp affordances in multi-object scenes,” in *CVPR*, 2020, pp. 5030–5040.
- [21] Y. Hasson, B. Tekin, F. Bogo, I. Laptev, M. Pollefeys, and C. Schmid, “Leveraging photometric consistency over time for sparsely supervised hand-object reconstruction,” in *CVPR*, 2020, pp. 568–577.
- [22] G. Moon and K. M. 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.
- [23] M. Seeber, R. Poranne, M. Pollefeys, and M. R. Oswald, “Realistichands: A hybrid model for 3d hand reconstruction,” in *International Conference on 3D Vision (3DV)*, 2021, pp. 22–31.
- [24] J. Wang, F. Mueller, F. Bernard, *et al.*, “Rgb2hands: Real-time tracking of 3d hand interactions from monocular RGB video,” *ACM Trans. Graph.*, vol. 39, no. 6, pp. 218:1–218:16, 2020.
- [25] X. Zhang, Q. Li, H. Mo, W. Zhang, and W. Zheng, “End-to-end hand mesh recovery from a monocular RGB image,” in *IEEE International Conference on Computer Vision (ICCV)*, 2019, pp. 2354–2364.
- [26] Y. Zhou, M. Habermann, W. Xu, I. Habibie, C. Theobalt, and F. Xu, “Monocular real-time hand shape and motion capture using multi-modal data,” in *CVPR*, 2020, pp. 5345–5354.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *CVPR*, 2016, pp. 770–778.
- [28] H. Kato, D. Beker, M. Morariu, *et al.*, “Differentiable rendering: A survey,” *CoRR*, vol. abs/2006.12057, 2020.
- [29] N. Ravi, J. Reizenstein, D. Novotný, *et al.*, “Accelerating 3d deep learning with pytorch3d,” *CoRR*, vol. abs/2007.08501, 2020.
- [30] G. Lin, F. Liu, A. Milan, C. Shen, and I. D. Reid, “Refinenet: Multi-path refinement networks for dense prediction,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 5, pp. 1228–1242, 2020.
- [31] P. Schönemann, “A generalized solution of the orthogonal procrustes problem,” *Psychometrika*, vol. 31, no. 1, pp. 1–10, 1966.
- [32] P. F. Felzenszwalb and D. P. Huttenlocher, “Distance transforms of sampled functions,” *Theory Comput.*, vol. 8, no. 1, pp. 415–428, 2012.
- [33] O. Taheri, N. Ghorbani, M. J. Black, and D. Tzionas, “GRAB: A dataset of whole-body human grasping of objects,” in *European Conference on Computer Vision (ECCV)*, 2020.
- [34] C. Zimmermann, D. Ceylan, J. Yang, B. C. Russell, M. J. Argus, and T. Brox, “Freihand: A dataset for markerless capture of hand pose and shape from single RGB images,” in *IEEE International Conference on Computer Vision (ICCV)*, 2019, pp. 813–822.
- [35] G. Moon, J. Y. Chang, and K. M. Lee, “V2v-posenet: Voxel-to-voxel prediction network for accurate 3d hand and human pose estimation from a single depth map,” in *CVPR*, 2018, pp. 5079–5088.
- [36] L. Ge, Y. Cai, J. Weng, and J. Yuan, “Hand pointnet: 3d hand pose estimation using point sets,” in *CVPR*, 2018, pp. 8417–8426.