

Geometric Correspondence Fields: Learned Differentiable Rendering for 3D Pose Refinement in the Wild

Alexander Grabner^{1,2}, Yaming Wang², Peizhao Zhang², Peihong Guo²,
Tong Xiao², Peter Vajda², Peter M. Roth¹, and Vincent Lepetit¹

¹ Graz University of Technology, Austria ² Facebook Inc., USA
{alexander.grabner, pmroth, lepetit}@icg.tugraz.at
{wym, stzpz, phg, xiaot, vajdap}@fb.com

Abstract. We present a novel 3D pose refinement approach based on differentiable rendering for objects of arbitrary categories in the wild. In contrast to previous methods, we make two main contributions: First, instead of comparing real-world images and synthetic renderings in the RGB or mask space, we compare them in a feature space optimized for 3D pose refinement. Second, we introduce a novel differentiable renderer that learns to approximate the rasterization backward pass from data instead of relying on a hand-crafted algorithm. For this purpose, we predict deep cross-domain correspondences between RGB images and 3D model renderings in the form of what we call geometric correspondence fields. These correspondence fields serve as pixel-level gradients which are analytically propagated backward through the rendering pipeline to perform a gradient-based optimization directly on the 3D pose. In this way, we precisely align 3D models to objects in RGB images which results in significantly improved 3D pose estimates. We evaluate our approach on the challenging Pix3D dataset and achieve up to 55% relative improvement compared to state-of-the-art refinement methods in multiple metrics.

1 Introduction

Recently, there have been significant advances in single image 3D object pose estimation thanks to deep learning [9, 38, 49]. However, the accuracy achieved by today’s feed-forward networks is not sufficient for many applications like augmented reality or robotics [10, 51]. As shown in Figure 1, feed-forward networks can robustly estimate the coarse high-level 3D rotation and 3D translation of objects in the wild (*left*) but fail to predict fine-grained 3D poses (*right*) [56].

To improve the accuracy of predicted 3D poses, refinement methods aim at aligning 3D models to objects in RGB images. In this context, many methods train a feed-forward network that directly predicts 3D pose updates given the input image and a 3D model rendering under the current 3D pose estimate [29, 45, 60]. In contrast, more recent methods use differentiable rendering [33] to explicitly optimize an objective function conditioned on the input image and



Fig. 1: Given an initial 3D pose predicted by a feed-forward network (*left*), we predict deep cross-domain correspondences between real-world RGB images and synthetic 3D model renderings in the form of geometric correspondence fields (*middle*) that enable us to refine the 3D pose in a differentiable rendering framework (*right*).

renderer inputs like the 3D pose [23, 41]. These methods yield more accurate 3D pose updates because they exploit prior knowledge about the rendering pipeline.

However, existing approaches based on differentiable rendering have significant shortcomings because they rely on comparisons in the RGB or mask space. First, methods which compare real-world images and synthetic renderings in the RGB space require photo-realistic 3D model renderings [33]. Generating such renderings is difficult because objects in the real world are subject to complex scene lighting, unknown reflection properties, and cluttered backgrounds. Moreover, many 3D models only provide geometry but no textures or materials which makes photo-realistic rendering impossible [46]. Second, methods which rely on comparisons in the mask space need to predict accurate masks from real-world RGB images [23, 41]. Generating such masks is difficult even using state-of-the-art approaches like Mask R-CNN [13]. Additionally, masks discard valuable object shape information which makes 2D-3D alignment ambiguous. As a consequence, the methods described above are not robust in practice. Finally, computing gradients for the non-differentiable rasterization operation in rendering is still an open research problem and existing approaches rely on hand-crafted approximations for this task [16, 23, 33].

To overcome these limitations, we compare RGB images and 3D model renderings in a feature space optimized for 3D pose refinement and learn to approximate the rasterization backward pass in differentiable rendering from data. In particular, we introduce a novel network architecture that jointly performs both tasks. Our network maps real-world images and synthetic renderings to a common feature space and predicts deep cross-domain correspondences in the form of *geometric correspondence fields* (see Figure 1, *middle*). Geometric correspondence fields hold 2D displacement vectors between corresponding 2D object points in RGB images and 3D model renderings similar to optical flow [7]. These predicted 2D displacement vectors serve as pixel-level gradients that enable us to approximate the rasterization backward pass and compute accurate gradients for renderer inputs like the 3D pose, the 3D model, or the camera intrinsics that minimize an ideal geometric reprojection loss.

Our approach has three main advantages: First, we can leverage depth, normal, and object coordinate [2] renderings which provide 3D pose information more explicitly than RGB and mask renderings [29]. Second, we avoid task-irrelevant appearance variations in the RGB space and 3D pose ambiguities in the mask space [11]. Third, we learn to approximate the rasterization backward pass from data instead of relying on a hand-crafted algorithm [16, 23, 33].

To demonstrate the benefits of our novel 3D pose refinement approach, we evaluate it on the challenging Pix3D [46] dataset. We present quantitative as well as qualitative results and significantly outperform state-of-the-art refinement methods in multiple metrics by up 55% relative. Finally, we combine our refinement approach with feed-forward 3D pose estimation [10] and 3D model retrieval [11] methods to predict fine-grained 3D poses for objects in the wild without providing initial 3D poses or ground truth 3D models at runtime **given only a single RGB image**. To summarize, our main contributions are:

- We introduce the first refinement method based on differentiable rendering that does not compare real-world images and synthetic renderings in the RGB or mask space but in a feature space optimized for the task at hand.
- We present a novel differentiable renderer that learns to approximate the rasterization backward pass instead of relying on a hand-crafted algorithm.

2 Related Work

In the following, we discuss prior work on differentiable rendering, 3D pose estimation, and 3D pose refinement.

2.1 Differentiable Rendering

Differentiable rendering [33] is a powerful concept that provides inverse graphics capabilities by computing gradients for 3D scene parameters from 2D image observations. This novel technique recently gained popularity for 3D reconstruction [26, 39], scene lighting estimation [1, 28], and texture prediction [22, 57].

However, rendering is a non-differentiable process due to the rasterization operation [23]. Thus, differentiable rendering approaches either try to mimic rasterization with differentiable operations [32, 41] or use conventional rasterization and approximate its backward pass [8, 16, 52].

In this work, we also approximate the rasterization backward pass but, in contrast to existing methods, do not rely on hand-crafted approximations. Instead, we train a network that performs the approximation. This idea is not only applicable for 3D pose estimation but also for other tasks like 3D reconstruction, human pose estimation, or the prediction of camera intrinsics in the future.

2.2 3D Pose Estimation

Modern 3D pose estimation approaches build on deep feed-forward networks and can be divided into two groups: Direct and correspondence-based methods.

Direct methods predict 3D pose parameters as raw network outputs. They use classification [48, 49], regression [36, 53], or hybrid variants of both [34, 56] to estimate 3D rotation and 3D translation [27, 37, 38] in an end-to-end manner. Recent approaches additionally integrate these techniques into detection pipelines to deal with multiple objects in a single image [24, 26, 51, 55].

In contrast, correspondence-based methods predict keypoint locations and recover 3D poses from 2D-3D correspondences using PnP algorithms [43, 45] or trained shape models [42]. In this context, different methods predict sparse object-specific keypoints [42–44], sparse virtual control points [9, 45, 47], or dense unsupervised 2D-3D correspondences [2, 3, 10, 21, 50].

In this work, we use the correspondence-based feed-forward approach presented in [10] to predict initial 3D poses for refinement.

2.3 3D Pose Refinement

3D pose refinement methods are based on the assumption that the projection of an object’s 3D model aligns with the object’s appearance in the image given the correct 3D pose. Thus, they compare renderings under the current 3D pose to the input image to get feedback on the prediction.

A simple approach to refine 3D poses is to generate many small perturbations and evaluate their accuracy using a scoring function [31, 58]. However, this is computationally expensive and the design of the scoring function is unclear. Therefore, other approaches try to predict iterative 3D pose updates with deep networks instead [29, 35, 45, 60]. In practice, though, the performance of these methods is limited because they cannot generalize to 3D poses or 3D models that have not been seen during training [45].

Recent approaches based on differentiable rendering overcome these limitations [23, 33, 41]. Compared to the methods described above, they analytically propagate error signals backward through the rendering pipeline to compute more accurate 3D pose updates. In this way, they exploit knowledge about the 3D scene geometry and the projection pipeline for the optimization.

In contrast to existing differentiable rendering approaches that rely on comparisons in the RGB [33] or mask [23, 41] space, we compare RGB images and 3D model renderings in a feature space that is optimized for 3D pose refinement.

3 Learned 3D Pose Refinement

Given a single RGB image, a 3D model, and an initial 3D pose, we compute iterative updates to refine the 3D pose, as shown in Figure 2. For this purpose, we first introduce the objective function that we optimize at runtime (see Sec. 3.1). We then explain how we compare the input RGB image to renderings under the current 3D pose in a feature space optimized for refinement (see Sec. 3.2), predict pixel-level gradients that minimize an ideal geometric reprojection loss in the form of geometric correspondence fields (see Sec. 3.3), and propagate gradients backward through the rendering pipeline to perform a gradient-based optimization directly on the 3D pose (see Sec. 3.4).

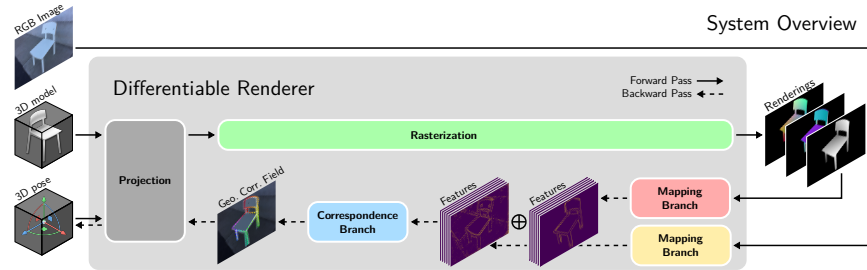


Fig. 2: Overview of our system. In the forward pass (\longrightarrow), we generate 3D model renderings under the current 3D pose. In the backward pass (\longleftarrow), we map the RGB image and our renderings to a common feature space and predict a geometric correspondence field that enables us to approximate the rasterization backward pass and compute gradients for the 3D pose that minimize an ideal geometric reprojection loss.

3.1 Runtime Object Function

Our approach to refine the 3D pose of an object is based on the numeric optimization of an objective function at runtime. In particular, we seek to minimize an ideal geometric reprojection loss

$$e(\mathcal{P}) = \frac{1}{2} \sum_i \|\text{proj}(\mathbf{M}_i, \mathcal{P}_{\text{gt}}) - \text{proj}(\mathbf{M}_i, \mathcal{P})\|_2^2 \quad (1)$$

for all provided 3D model vertices \mathbf{M}_i . In this case, $\text{proj}(\cdot)$ performs the projection from 3D space to the 2D image plane, \mathcal{P} denotes the 3D pose parameters, and \mathcal{P}_{gt} is the ground truth 3D pose. Hence, it is clear that $\text{argmin } e(\mathcal{P}) = \mathcal{P}_{\text{gt}}$.

To efficiently minimize $e(\mathcal{P})$ using a gradient-based optimization starting from an initial 3D pose, we compute gradients for the 3D pose using the Jacobian of $e(\mathcal{P})$ with respect to \mathcal{P} . Applying the chain rule yields the expression

$$\left(\frac{\partial e(\mathcal{P})}{\partial \mathcal{P}} \right) (\mathcal{P}_{\text{curr}}) = \sum_i \left[\frac{\partial \text{proj}(\mathbf{M}_i, \mathcal{P})}{\partial \mathcal{P}} \right]^T \left[\text{proj}(\mathbf{M}_i, \mathcal{P}_{\text{gt}}) - \text{proj}(\mathbf{M}_i, \mathcal{P}_{\text{curr}}) \right], \quad (2)$$

where $\mathcal{P}_{\text{curr}}$ is the current 3D pose estimate and the point where the Jacobian is evaluated. In this case, the term $\left[\frac{\partial \text{proj}(\mathbf{M}_i, \mathcal{P})}{\partial \mathcal{P}} \right]^T$ can be computed analytically because it is simply a sequence of differentiable operations. In contrast, the term $\left[\text{proj}(\mathbf{M}_i, \mathcal{P}_{\text{gt}}) - \text{proj}(\mathbf{M}_i, \mathcal{P}_{\text{curr}}) \right]$ cannot be computed analytically because the 3D model vertices projected under the ground truth 3D pose, *i.e.*, $\text{proj}(\mathbf{M}_i, \mathcal{P}_{\text{gt}})$, are unknown at runtime and can only be observed indirectly via the input image.

However, for visible vertices, this term can be calculated given a geometric correspondence field (see Sec. 3.4). Thus, we introduce a novel network architecture that learns to predict geometric correspondence fields given an RGB image and 3D model renderings under the current 3D pose estimate in the following. Moreover, we embed this network in a differentiable rendering framework to approximate the rasterization backward pass and compute gradients for renderer inputs like the 3D pose of an object in an end-to-end manner (see Figure 2).

3.2 Refinement Feature Space

The first step in our approach is to render the provided 3D model under the current 3D pose using the forward pass of our differentiable renderer (see Figure 2). In particular, we generate depth, normal, and object coordinate [2] renderings. These representations provide 3D pose and 3D shape information more explicitly than RGB or mask renderings which makes them particularly useful for 3D pose refinement [11]. By concatenating the different renderings along the channel dimension, we leverage complementary information from different representations in the backward pass rather than relying on a single type of rendering [33].

Next, we begin the backward pass of our differentiable renderer by mapping the input RGB image and our multi-representation renderings to a common feature space. For this task, we use two different network branches that bridge the domain gap between the real and rendered images (see Figure 2). Our mapping branches use a custom architecture based on task-specific design choices:

First, we want to predict local cross-domain correspondences under the assumption that the initial 3D pose is close to the ground truth 3D pose. Thus, we do not require features with global context but features with local discriminability. For this reason, we use small fully convolutional networks which are fast, memory-efficient, and learn low-level features that generalize well across different objects. Because the low-level structures appearing across different objects are similar, we do not require a different network for each object [45] but address objects of all categories with a single class-agnostic network for each domain.

Second, we want to predict correspondences with maximum spatial accuracy. Thus, we do not use pooling or downsampling but maintain the spatial resolution throughout the network. In this configuration, consecutive convolutions provide sufficient receptive field to learn advanced shape features which are superior to simple edge intensities [24], while higher layers benefit from full spatial parameter sharing during training which increases generalization. As a consequence, the effective minibatch size during training is higher than the number of images per minibatch because only a subset of all image pixels contributes to each computed feature. In addition, the resulting high spatial resolution feature space provides an optimal foundation for computing spatially accurate correspondences.

For the implementation of our mapping branches, we use fully convolutional networks consisting of an initial 7×7 Conv-BN-ReLU block, followed by three residual blocks [14, 15], and a 1×1 Conv-BN-ReLU block for dimensionality reduction. This architecture enforces local discriminability and high spatial resolution and maps RGB images and multi-representation renderings to $W \times H \times 64$ feature maps, where W and H are the spatial dimensions of the input image.

3.3 Geometric Correspondence Fields

After mapping RGB images and 3D model renderings to a common feature space, we compare their feature maps and predict cross-domain correspondences. For this purpose, we concatenate their feature maps and use another fully convolutional network branch to predict correspondences, as shown in Figure 2.

In particular, we regress per-pixel correspondence vectors in the form of geometric correspondence fields (see Figure 1, *middle*). Geometric correspondence fields hold 2D displacement vectors between corresponding 2D object points in real-world RGB images and synthetic 3D model renderings similar to optical flow [7]. These displacement vectors represent the projected relative 2D motion of individual 2D object points that is required to minimize the reprojection error and refine the 3D pose. A geometric correspondence field has the same spatial resolution as the respective input RGB image and two channels, *i.e.*, $W \times H \times 2$.

If an object’s 3D model and 3D pose are known, we can render the ground truth geometric correspondence field for an arbitrary 3D pose. For this task, we first compute the 2D displacement $\nabla \mathbf{m}_i = \text{proj}(\mathbf{M}_i, \mathcal{P}_{\text{gt}}) - \text{proj}(\mathbf{M}_i, \mathcal{P}_{\text{curr}})$ between the projection under the ground truth 3D pose \mathcal{P}_{gt} and the current 3D pose $\mathcal{P}_{\text{curr}}$ for each 3D model vertex \mathbf{M}_i . We then generate a ground truth geometric correspondence field $G(\mathcal{P}_{\text{curr}}, \mathcal{P}_{\text{gt}})$ by rendering the 3D model using a shader that interpolates the per-vertex 2D displacements $\nabla \mathbf{m}_i$ across the projected triangle surfaces using barycentric coordinates.

In our scenario, predicting correspondences using a network has two advantages compared to traditional correspondence matching [12]. First, predicting correspondences with convolutional kernels is significantly faster than exhaustive feature matching during both training and testing [59]. This is especially important in the case of dense correspondences. Second, training explicit correspondences can easily result in degenerated feature spaces and requires tedious regularization and hard negative sample mining [5].

For the implementation of our correspondence branch, we use three consecutive 7×7 Conv-BN-ReLU blocks followed by a final 7×7 convolution which reduces the channel dimensionality to two. For this network, a large receptive field is crucial to cover correspondences with high spatial displacement.

However, in many cases local correspondence prediction is ambiguous. For example, many objects are untextured and have homogeneous surfaces, *e.g.*, the backrest and the seating surface of the chair in Figure 1, which cause unreliable correspondence predictions. To address this problem, we additionally employ a geometric attention module which restricts the correspondence prediction to visible object regions with significant geometric discontinuities, as outlined in white underneath the 2D displacement vectors in Figure 1. We identify these regions by finding local variations in our renderings.

In particular, we detect rendering-specific intensity changes larger than a certain threshold within a local 5×5 window to construct a geometric attention mask w^{att} . For each pixel of w^{att} , we compute the geometric attention weight

$$w_{x,y}^{\text{att}} = \max_{u,v \in W} \left(\delta^R \left(R(\mathcal{P}_{\text{curr}})_{x,y}, R(\mathcal{P}_{\text{curr}})_{x-u,y-v} \right) \right) > t^R. \quad (3)$$

In this case, $R(\mathcal{P}_{\text{curr}})$ is a concatenation of depth, normal, and object coordinate renderings under the current 3D pose $\mathcal{P}_{\text{curr}}$, (x, y) is a pixel location, and (u, v) are pixel offsets within the window W . The comparison function $\delta^R(\cdot)$ and the threshold t^R are different for each type of rendering. For depth renderings, we

compute the absolute difference between normalized depth values and use a threshold of 0.1. For normal renderings, we compute the angle between normals and use a threshold of 15° . For object coordinate renderings, we compute the Euclidean distance between 3D points and use a threshold of 0.1. If any of these thresholds applies, the corresponding pixel (x, y) in our geometric attention mask w^{att} becomes active. Because we already generated these renderings before, our geometric attention mechanism requires almost no additional computations and is available during both training and testing.

Training. During training of our system, we optimize the learnable part of our differentiable renderer, *i.e.*, a joint network $f(\cdot)$ consisting of our two mapping branches and our correspondence branch with parameters θ (see Figure 2). Formally, we minimize the error between predicted $f(\cdot)$ and ground truth $G(\cdot)$ geometric correspondence fields as

$$\min_{\theta} \sum_{x,y} w_{x,y}^{att} \|f(I, R(\mathcal{P}_{curr}); \theta)_{x,y} - G(\mathcal{P}_{curr}, \mathcal{P}_{gt})_{x,y}\|_2^2. \quad (4)$$

In this case, w^{att} is a geometric attention mask, I is an RGB image, $R(\mathcal{P}_{curr})$ is a concatenation of depth, normal, and object coordinate renderings generated under a random 3D pose \mathcal{P}_{curr} produced by perturbing the ground truth 3D pose \mathcal{P}_{gt} , $G(\mathcal{P}_{curr}, \mathcal{P}_{gt})$ is the ground truth geometric correspondence field, and (x, y) is a pixel location. In particular, we first generate a random 3D pose \mathcal{P}_{curr} around the ground truth 3D pose \mathcal{P}_{gt} for each training sample in each iteration. For this purpose, we sample 3D pose perturbations from normal distributions and apply them to \mathcal{P}_{gt} to generate \mathcal{P}_{curr} . For 3D rotations, we use absolute perturbations with $\sigma = 5^\circ$. For 3D translations, we use relative perturbations with $\sigma = 0.1$. We then render the ground truth geometric correspondence field $G(\mathcal{P}_{curr}, \mathcal{P}_{gt})$ between the perturbed 3D pose \mathcal{P}_{curr} and the ground truth 3D pose \mathcal{P}_{gt} as described above, generate concatenated depth, normal, and object coordinate renderings $R(\mathcal{P}_{curr})$ under the perturbed 3D pose \mathcal{P}_{curr} , and compute the geometric attention mask w^{att} . Finally, we predict a geometric correspondence field using our network $f(I, R(\mathcal{P}_{curr}); \theta)$ given the RGB image I and the renderings $R(\mathcal{P}_{curr})$, and optimize for the network parameters θ .

In this way, we train a network that performs three tasks: First, it maps RGB images and multi-representation 3D model renderings to a common feature space. Second, it compares features in this space. Third, it predicts geometric correspondence fields which serve as pixel-level gradients that enable us to approximate the rasterization backward pass of our differentiable renderer.

3.4 Learned Differentiable Rendering

In the classic rendering pipeline, the only non-differentiable operation is the rasterization [33] that determines which pixels of a rendering have to be filled, solves the visibility of projected triangles, and fills the pixels using a shading computation. This discrete operation raises one main challenge: Its gradient is zero, which prevents gradient flow [23]. However, we must flow non-zero gradients from pixels to projected 3D model vertices to perform differentiable rendering.




Fig. 3: To approximate the rasterization backward pass, we predict a geometric correspondence field (*left*), disperse the predicted 2D displacement of each pixel among the vertices of its corresponding visible triangle (*middle*), and normalize the contributions of all pixels. In this way, we obtain gradients for projected 3D model vertices (*right*).

We solve this problem using geometric correspondence fields. Instead of actually differentiating a loss in the image space and relying on hand-crafted comparisons between pixel intensities to approximate the gradient flow from pixels to projected 3D model vertices [16, 23], we first use a network to predict per-pixel 2D displacement vectors in the form of a geometric correspondence field, as shown in Figure 2. We then compute gradients for projected 3D model vertices by simply accumulating the predicted 2D displacement vectors using our knowledge of the projected 3D model geometry, as illustrated in Figure 3.

Formally, we compute the gradient of a projected 3D model vertex \mathbf{m}_i as

$$\nabla \mathbf{m}_i = \frac{1}{\sum_{x,y} w_{x,y}^{att} w_{x,y}^{bar,i}} \sum_{x,y} w_{x,y}^{att} w_{x,y}^{bar,i} f(I, R(\mathcal{P}_{curr}); \theta)_{x,y} \quad (5)$$

$$\forall x, y : \mathbf{m}_i \in \Delta_{\text{IndexMap}_{x,y}} .$$

In this case, $f(I, R(\mathcal{P}_{curr}); \theta)$ is a geometric correspondence field predicted by our network $f(\cdot)$ with frozen parameters θ given an RGB image I and concatenated 3D model renderings $R(\mathcal{P}_{curr})$ under the current 3D pose estimate \mathcal{P}_{curr} , $w_{x,y}^{att}$ is a geometric attention weight, $w_{x,y}^{bar,i}$ is a barycentric weight for \mathbf{m}_i , and (x, y) is a pixel position. We accumulate predicted 2D displacement vectors for all positions (x, y) for which \mathbf{m}_i is a vertex of the triangle $\Delta_{\text{IndexMap}_{x,y}}$ visible at (x, y) . For this task, we generate an `IndexMap` which stores the index of the visible triangle for each pixel during the forward pass of our differentiable renderer.

Inference. Our computed $\nabla \mathbf{m}_i$ approximate the second term in Eq. (2) that cannot be computed analytically. In this way, our approach combines local per-pixel 2D displacement vectors into per-vertex gradients and further computes accurate global 3D pose gradients considering the 3D model geometry and the rendering pipeline. Our experiments show that this approach generalizes better to unseen data than predicting 3D pose updates with a network [29, 60].

During inference of our system, we perform iterative updates to refine \mathcal{P}_{curr} . In each iteration, we compute a 3D pose gradient by evaluating our refinement loop presented in Figure 2. For our implementation, we use the Adam optimizer [25] with a small learning rate and perform multiple updates to account for noisy correspondences and achieve the best accuracy.

4 Experimental Results

To demonstrate the benefits of our 3D pose refinement approach, we evaluate it on the challenging Pix3D [46] dataset which provides in-the-wild images for objects of different categories. In particular, we quantitatively and qualitatively compare our approach to state-of-the-art refinement methods in Sec. 4.1, perform an ablation study in Sec. 4.2, and combine our refinement approach with feed-forward 3D pose estimation [10] and 3D model retrieval [11] methods to predict fine-grained 3D poses without providing initial 3D poses or ground truth 3D models in Sec. 4.3. We follow the evaluation protocol of [10] and report the median (*MedErr*) of rotation, translation, pose, and projection distances. Details on evaluation setup, datasets, and metrics as well as extensive results and further experiments are provided in our **supplementary material**.

4.1 Comparison to the State of the Art

We first quantitatively compare our approach to state-of-the-art refinement methods. For this purpose, we perform 3D pose refinement on top of an initial 3D pose estimation baseline. In particular, we predict initial 3D poses using the feed-forward approach presented in [10] which is the state of the art for single image 3D pose estimation on the Pix3D dataset (*Baseline*). We compare our refinement approach to traditional image-based refinement without differentiable rendering [60] (*Image Refinement*) and mask-based refinement with differentiable rendering [23] (*Mask Refinement*).

RGB-based refinement with differentiable rendering [33] is not possible in our setup because all available 3D models lack textures and materials. This approach even fails if we compare grey-scale images and renderings because the image intensities at corresponding locations do not match. As a consequence, 2D-3D alignment using a photo-metric loss is impossible.

For *Image Refinement*, we use grey-scale instead of RGB renderings because all available 3D models lack textures and materials. In addition, we do not perform a single full update [60] but perform up to 1000 iterative updates with a small learning rate of $\eta = 0.05$ using the Adam optimizer [25] for all methods.

For *Mask Refinement*, we predict instance masks from the input RGB image using Mask R-CNN [13]. To achieve maximum accuracy, we internally predict masks at four times the original spatial resolution proposed in Mask R-CNN and fine-tune a model pre-trained on COCO [30] on Pix3D.

Table 1 (**upper part**) summarizes our results. In this experiment, we provide the ground truth 3D model of the object in the image for refinement. Compared to the baseline, *Image Refinement* only achieves a small improvement in the rotation, translation, and pose metrics. There is almost no improvement in the projection metric (*MedErr_P*), as this method does not minimize the reprojection error. Traditional refinement methods are not aware of the rendering pipeline and the underlying 3D scene geometry and can only provide coarse 3D pose updates [60]. In our in-the-wild scenario, the number of 3D models, possible 3D pose perturbations, and category-level appearance variations is too large to

Table 1: Quantitative 3D pose refinement results on the Pix3D dataset. In the case of provided **ground truth 3D models**, our refinement significantly outperforms previous refinement methods across all metrics by **up to 55%** relative. In the case of automatically **retrieved 3D models** (+ Retrieval [11]), we reduce the 3D pose error ($MedErr_{R,t}$) compared to the state of the art for single image 3D pose estimation on Pix3D (*Baseline*) by **55%** relative **without using additional inputs**.

Method	Input	Rotation	Translation	Pose	Projection
		$MedErr_R$ -1	$MedErr_t$ -10 ²	$MedErr_{R,t}$ -10 ²	$MedErr_P$ -10 ²
Baseline [10]	RGB Image	6.75	6.21	4.76	3.71
Image Refinement [60]	RGB Image + 3D Model	6.46	5.43	4.31	3.67
Mask Refinement [23]	RGB Image + 3D Model	3.56	4.06	2.96	1.90
Our Refinement	RGB Image + 3D Model	2.56	1.74	1.34	1.27
Image Refinement [60] + Retrieval [11]	RGB Image	6.47	5.51	4.33	3.74
Mask Refinement [23] + Retrieval [11]	RGB Image	5.47	5.25	4.15	3.12
Our Refinement + Retrieval [11]	RGB Image	3.79	2.65	2.14	2.18

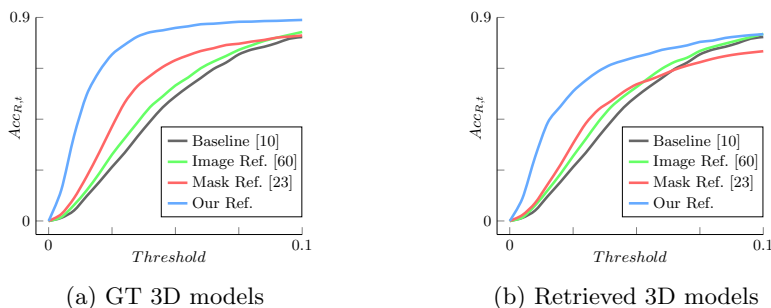


Fig. 4: Evaluation on 3D pose accuracy at different thresholds. We significantly outperform other methods on strict thresholds using both GT and retrieved 3D models.

simulate all permutations during training. As a consequence, this method cannot generalize to examples which are far from the ones seen during training and only achieves small improvements.

Additionally, we observe that after the first couple of refinement steps, the predicted updates are not accurate enough to refine the 3D pose but start to jitter without further improving the 3D pose. Moreover, for many objects, the prediction fails and the iterative updates cause the 3D pose to drift off. Empirically, we obtain the best overall results for this method using only 20 iterations. For all other methods based on differentiable rendering, we achieve the best accuracy after the full 1000 iterations. A detailed analysis on this issue is presented in our supplementary material.

Next, *Mask Refinement* improves upon *Image Refinement* by a large margin across all metrics. Due to the 2D-3D alignment with differentiable rendering, this method computes more accurate 3D pose updates and additionally reduces the reprojection error ($MedErr_P$). However, we observe that *Mask Refinement* fails in two common situations: First, when the object has holes and the mask is not a




Fig. 5: Qualitative 3D pose refinement results for objects of different categories. We project the ground truth 3D model on the image using the 3D pose estimated by different methods. Our approach overcomes the limitations of previous methods and predicts fine-grained 3D poses for objects in the wild. The last example shows a failure case (indicated by the **red frame**) where the initial 3D pose is too far from the ground truth 3D pose and no refinement method can converge. More qualitative results are presented in our **supplementary material**. Best viewed in **digital zoom**.

single blob the refinement fails (see Figure 5, *e.g.*, 1st row - right example). In the presence of holes, the hand-crafted approximation for the rasterization backward pass accumulates gradients with alternating signs while traversing the image. This results in unreliable per-vertex motion gradients. Second, simply aligning the silhouette of objects is ambiguous as renderings under different 3D poses can have similar masks. The interior structure of the object is completely ignored. As a consequence, the refinement gets stuck in poor local minima. Finally, the performance of *Mask Refinement* is limited by the quality of the target mask predicted from the RGB input image [13].

In contrast, our refinement overcomes these limitations and significantly outperforms the baseline as well as competing refinement methods across all metrics by **up to 70%** and **55%** relative. Using our geometric correspondence fields, we bridge the domain gap between real-world images and synthetic renderings and align both the object outline as well as interior structures with high accuracy.

Our approach performs especially well in the fine-grained regime, as shown in Figure 4a. In this experiment, we plot the 3D pose accuracy $Acc_{R,t}$ which gives the percentage of samples for which the 3D pose distance $e_{R,t}$ is smaller than a varying threshold. For strict thresholds close to zero, our approach outperforms other refinement methods by a large margin. For example, at the threshold 0.015, we achieve more than 55% accuracy while the runner-up *Mask Refinement* achieves only 19% accuracy.

This significant performance improvement is also reflected in our qualitative examples presented in Figure 5. Our approach precisely aligns 3D models to objects in RGB images and computes 3D poses which are in many cases visually indistinguishable from the ground truth. Even if the initial 3D pose estimate (*Baseline*) is significantly off, our method can converge towards the correct 3D pose (see Figure 5, *e.g.*, 1st row - left example). Finally, Figure 6 illustrates the high quality of our predicted geometric correspondence fields.




Fig. 6: Qualitative examples of our predicted geometric correspondence fields. Our predicted 2D displacement vectors are highly accurate. Best viewed in **digital zoom**.

4.2 Ablation Study

To understand the importance of individual components in our system, we conduct an ablation study in Table 2. For this purpose, we modify a specific system component, retrain our approach, and evaluate the performance impact.

If we use smaller kernels with less receptive field (3×3 vs 7×7) or fewer layers (2 vs 4) in our correspondence branch, the performance drops significantly. Also, using shallow mapping branches which only employ a single Conv-BN-ReLU block to simulate simple edge and ridge features results in low accuracy because the computed features are not discriminative enough. If we perform refinement without our geometric attention mechanism, the accuracy degrades due to unreliable correspondence predictions in homogeneous regions.

Next, the choice of the rendered representation is important for the performance of our approach. While using masks only performs poorly, depth, normal, and object coordinate renderings increase the accuracy. Finally, we achieve the best accuracy by exploiting complementary information from multiple different renderings by concatenating depth, normal, and object coordinate renderings.

By inspecting failure cases, we observe that our method does not converge if the initial 3D pose is too far from the ground truth 3D pose (see Figure 5, last example). In this case, we cannot predict accurate correspondences because our computed features are not robust to large viewpoint changes and the receptive field of our correspondence branch is limited. In addition, occlusions cause our refinement to fail because there are no explicit mechanisms to address them. We plan to resolve this issue in the future by predicting occlusion masks and correspondence confidences. However, other refinement methods also fail in these scenarios (see Figure 5, last example).

4.3 3D Model Retrieval

So far, we assumed that the ground truth 3D model required for 3D pose refinement is given at runtime. However, we can overcome this limitation by automatically retrieving 3D models from single RGB images. For this purpose, we combine all refinement approaches with the retrieval method presented in [11], where the 3D model database essentially becomes a part of the trained model. In this way, we perform initial 3D pose estimation, 3D model retrieval, and 3D pose refinement given only a single RGB image. This setting allows us to benchmark refinement methods against feed-forward baselines in a fair comparison.

Table 2: Ablation study of our method. Using components which increase the discriminability of learned features is important for the performance of our approach. Also, our geometric attention mechanism and the chosen type of rendering effect the accuracy.

Method	Rotation	Translation	Pose	Projection
	$MedErr_R$	$MedErr_t$	$MedErr_{R,t}$	$MedErr_P$
	$\cdot 1$	$\cdot 10^2$	$\cdot 10^2$	$\cdot 10^2$
Ours less receptive	3.01	2.12	1.75	1.56
Ours fewer layers	2.84	1.98	1.58	1.41
Ours shallow features	2.76	2.00	1.55	1.42
Ours without attention	2.70	1.92	1.45	1.37
Ours only MASK	2.98	1.85	1.44	1.41
Ours only OBJ COORD	2.57	1.80	1.39	1.31
Ours only DEPTH	2.60	1.77	1.38	1.29
Ours only NORMAL	2.58	1.76	1.36	1.30
Ours	2.56	1.74	1.34	1.27

The corresponding results are presented in Table 1 (lower part) and Figure 4b. Because the retrieved 3D models often differ from the ground truth 3D models, the refinement performance decreases compared to given ground truth 3D models. Differentiable rendering methods lose more accuracy than traditional refinement methods because they require 3D models with accurate geometry.

Still, all refinement approaches perform remarkably well with retrieved 3D models. As long as the retrieved 3D model is reasonably close to the ground truth 3D model in terms of geometry, our refinement succeeds. Our method achieves even lower 3D pose error ($MedErr_{R,t}$) with retrieved 3D models than *Mask Refinement* with ground truth 3D models. Finally, using our joint 3D pose estimation-retrieval-refinement pipeline, we reduce the 3D pose error ($MedErr_{R,t}$) compared to the state of the art for single image 3D pose estimation on Pix3D (*Baseline*) by **55% relative without using additional inputs**.

5 Conclusion

Aligning 3D models to objects in RGB images is the most accurate way to predict 3D poses. However, there is a domain gap between real-world images and synthetic renderings which makes this alignment challenging in practice. To address this problem, we predict deep cross-domain correspondences in a feature space optimized for 3D pose refinement and combine local 2D displacement vectors into global 3D pose updates using our novel differentiable renderer. Our method outperforms existing refinement approaches by up to 55% relative and can be combined with feed-forward 3D pose estimation and 3D model retrieval to predict fine-grained 3D poses for objects in the wild given only a single RGB image. Finally, our novel learned differentiable rendering framework can be used for other tasks in the future.

References

1. Azinovic, D., Li, T.M., Kaplanyan, A., Niessner, M.: Inverse Path Tracing for Joint Material and Lighting Estimation. In: Conference on Computer Vision and Pattern Recognition. pp. 2447–2456 (2019)
2. Brachmann, E., Krull, A., Michel, F., Gumhold, S., Shotton, J., Rother, C.: Learning 6D Object Pose Estimation using 3D Object Coordinates. In: European Conference on Computer Vision. pp. 536–551 (2014)
3. Brachmann, E., Michel, F., Krull, A., Ying Yang, M., Gumhold, S., Rother, C.: Uncertainty-Driven 6D Pose Estimation of Objects and Scenes from a Single RGB Image. In: Conference on Computer Vision and Pattern Recognition. pp. 3364–3372 (2016)
4. Calli, B., Singh, A., Walsman, A., Srinivasa, S., Abbeel, P., Dollar, A.: The YCB Object and Model Set: Towards Common Benchmarks for Manipulation Research. In: International Conference on Advanced Robotics. pp. 510–517 (2015)
5. Choy, C.B., Gwak, J., Savarese, S., Chandraker, M.: Universal Correspondence Network. In: Advances in Neural Information Processing Systems. pp. 2414–2422 (2016)
6. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: ScanNet: Richly-Annotated 3D Reconstructions of Indoor Scenes. In: Conference on Computer Vision and Pattern Recognition. pp. 5828–5839 (2017)
7. Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T.: FlowNet: Learning Optical Flow with Convolutional Networks. In: Conference on Computer Vision and Pattern Recognition. pp. 2758–2766 (2015)
8. Genova, K., Cole, F., Maschinot, A., Sarna, A., Vlasic, D., Freeman, W.T.: Unsupervised Training for 3D Morphable Model Regression. In: Conference on Computer Vision and Pattern Recognition. pp. 8377–8386 (2018)
9. Grabner, A., Roth, P.M., Lepetit, V.: 3D Pose Estimation and 3D Model Retrieval for Objects in the Wild. In: Conference on Computer Vision and Pattern Recognition. pp. 3022–3031 (2018)
10. Grabner, A., Roth, P.M., Lepetit, V.: GP²C: Geometric Projection Parameter Consensus for Joint 3D Pose and Focal Length Estimation in the Wild. In: International Conference on Computer Vision. pp. 2222–2231 (2019)
11. Grabner, A., Roth, P.M., Lepetit, V.: Location Field Descriptors: Single Image 3D Model Retrieval in the Wild. In: International Conference on 3D Vision. pp. 583–593 (2019)
12. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press (2003)
13. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: International Conference on Computer Vision. pp. 2980–2988 (2017)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. In: Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
15. He, K., Zhang, X., Ren, S., Sun, J.: Identity Mappings in Deep Residual Networks. In: European Conference on Computer Vision. pp. 630–645 (2016)
16. Henderson, P., Ferrari, V.: Learning to Generate and Reconstruct 3D Meshes with only 2D Supervision. In: British Machine Vision Conference. pp. 139:1–139:13 (2018)
17. Hinterstoisser, S., Cagniart, C., Ilic, S., Sturm, P., Navab, N., Fua, P., Lepetit, V.: Gradient Response Maps for Real-Time Detection of Textureless Objects. IEEE Transactions on Pattern Analysis and Machine Intelligence **34**(5), 876–888 (2011)

18. Hinterstoisser, S., Lepetit, V., Ilic, S., Holzer, S., Bradski, G., Konolige, K., Navab, N.: Model Based Training, Detection and Pose Estimation of Texture-Less 3D Objects in Heavily Cluttered Scenes. In: Asian Conference on Computer Vision. pp. 548–562 (2012)
19. Hodan, T., Haluza, P., Obdržálek, Š., Matas, J., Lourakis, M., Zabulis, X.: T-LESS: An RGB-D Dataset for 6D Pose Estimation of Texture-Less Objects. In: IEEE Winter Conference on Applications of Computer Vision. pp. 880–888 (2017)
20. Hodaň, T., Matas, J., Obdržálek, Š.: On Evaluation of 6D Object Pose Estimation. In: European Conference on Computer Vision. pp. 606–619 (2016)
21. Jafari, O.H., Mustikovela, S.K., Pertsch, K., Brachmann, E., Rother, C.: iPose: Instance-Aware 6D Pose Estimation of Partly Occluded Objects. In: Asian Conference on Computer Vision. pp. 477–492 (2018)
22. Kanazawa, A., Tulsiani, S., Efros, A.A., Malik, J.: Learning Category-Specific Mesh Reconstruction from Image Collections. In: European Conference on Computer Vision. pp. 371–386 (2018)
23. Kato, H., Ushiku, Y., Harada, T.: Neural 3D Mesh Renderer. In: Conference on Computer Vision and Pattern Recognition. pp. 3907–3916 (2018)
24. Kehl, W., Manhardt, F., Tombari, F., Ilic, S., Navab, N.: SSD-6D: Making RGB-Based 3D Detection and 6D Pose Estimation Great Again. In: International Conference on Computer Vision. pp. 1530–1538 (2017)
25. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. arXiv:1412.6980 (2014)
26. Kundu, A., Li, Y., Rehg, J.M.: 3D-RCNN: Instance-Level 3D Object Reconstruction via Render-and-Compare. In: Conference on Computer Vision and Pattern Recognition. pp. 3559–3568 (2018)
27. Li, C., Bai, J., Hager, G.D.: A Unified Framework for Multi-View Multi-Class Object Pose Estimation. In: European Conference on Computer Vision. pp. 1–16 (2018)
28. Li, T.M., Aittala, M., Durand, F., Lehtinen, J.: Differentiable Monte Carlo Ray Tracing through Edge Sampling. In: ACM SIGGRAPH Asia. pp. 222:1–222:11 (2018)
29. Li, Y., Wang, G., Ji, X., Xiang, Y., Fox, D.: DeepIM: Deep Iterative Matching for 6D Pose Estimation. In: European Conference on Computer Vision. pp. 683–698 (2018)
30. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common Objects in Context. In: European Conference on Computer Vision. pp. 740–755 (2014)
31. Liu, L., Lu, J., Xu, C., Tian, Q., Zhou, J.: Deep Fitting Degree Scoring Network for Monocular 3D Object Detection. In: Conference on Computer Vision and Pattern Recognition. pp. 1057–1066 (2019)
32. Liu, S., Li, T., Chen, W., Li, H.: Soft Rasterizer: A Differentiable Renderer for Image-based 3D Reasoning. In: International Conference on Computer Vision. pp. 7708–7717 (2019)
33. Loper, M.M., Black, M.J.: OpenDR: An Approximate Differentiable Renderer. In: European Conference on Computer Vision. pp. 154–169 (2014)
34. Mahendran, S., Ali, H., Vidal, R.: A Mixed Classification-Regression Framework for 3D Pose Estimation from 2D Images. In: British Machine Vision Conference. pp. 238:1–238:12 (2018)
35. Manhardt, F., Kehl, W., Navab, N., Tombari, F.: Deep Model-Based 6D Pose Refinement in RGB. In: European Conference on Computer Vision. pp. 800–815 (2018)

36. Massa, F., Marlet, R., Aubry, M.: Crafting a Multi-Task CNN for Viewpoint Estimation. In: British Machine Vision Conference. pp. 91:1–91:12 (2016)
37. Mottaghi, R., Xiang, Y., Savarese, S.: A Coarse-To-Fine Model for 3D Pose Estimation and Sub-Category Recognition. In: Conference on Computer Vision and Pattern Recognition. pp. 418–426 (2015)
38. Mousavian, A., Anguelov, D., Flynn, J., Kosecka, J.: 3D Bounding Box Estimation Using Deep Learning and Geometry. In: Conference on Computer Vision and Pattern Recognition. pp. 7074–7082 (2017)
39. Nguyen-Phuoc, T.H., Li, C., Balaban, S., Yang, Y.: RenderNet: A Deep Convolutional Network for Differentiable Rendering from 3D Shapes. In: Advances in Neural Information Processing Systems. pp. 7891–7901 (2018)
40. Oberweger, M., Rad, M., Lepetit, V.: Making Deep Heatmaps Robust to Partial Occlusions for 3D Object Pose Estimation. In: European Conference on Computer Vision. pp. 119–134 (2018)
41. Palazzi, A., Bergamini, L., Calderara, S., Cucchiara, R.: End-to-End 6-DoF Object Pose Estimation through Differentiable Rasterization. In: European Conference on Computer Vision Workshops. pp. 1–14 (2018)
42. Pavlakos, G., Zhou, X., Chan, A., Derpanis, K., Daniilidis, K.: 6-DoF Object Pose from Semantic Keypoints. In: International Conference on Robotics and Automation. pp. 2011–2018 (2017)
43. Peng, S., Liu, Y., Huang, Q., Zhou, X., Bao, H.: 3D Object Class Detection in the Wild. In: Conference on Computer Vision and Pattern Recognition. pp. 4561–4570 (2019)
44. Pepik, B., Stark, M., Gehler, P., Ritschel, T., Schiele, B.: 3D Object Class Detection in the Wild. In: Conference on Computer Vision and Pattern Recognition Workshops. pp. 1–10 (2015)
45. Rad, M., Lepetit, V.: BB8: A Scalable, Accurate, Robust to Partial Occlusion Method for Predicting the 3D Poses of Challenging Objects Without Using Depth. In: International Conference on Computer Vision. pp. 3828–3836 (2017)
46. Sun, X., Wu, J., Zhang, X., Zhang, Z., Zhang, C., Xue, T., Tenenbaum, J., Freeman, W.T.: Pix3D: Dataset and Methods for Single-Image 3D Shape Modeling. In: Conference on Computer Vision and Pattern Recognition. pp. 2974–2983 (2018)
47. Tekin, B., Sinha, S.N., Fua, P.: Real-Time Seamless Single Shot 6D Object Pose Prediction. In: Conference on Computer Vision and Pattern Recognition. pp. 292–301 (2018)
48. Tulsiani, S., Carreira, J., Malik, J.: Pose Induction for Novel Object Categories. In: International Conference on Computer Vision. pp. 64–72 (2015)
49. Tulsiani, S., Malik, J.: Viewpoints and Keypoints. In: Conference on Computer Vision and Pattern Recognition. pp. 1510–1519 (2015)
50. Wang, H., Sridhar, S., Huang, J., Valentin, J., Song, S., Guibas, L.: Normalized Object Coordinate Space for Category-Level 6D Object Pose and Size Estimation. In: Conference on Computer Vision and Pattern Recognition. pp. 2642–2651 (2019)
51. Wang, Y., Tan, X., Yang, Y., Liu, X., Ding, E., Zhou, F., Davis, L.S.: 3D Pose Estimation for Fine-Grained Object Categories. In: European Conference on Computer Vision Workshops (2018)
52. Wu, J., Zhang, C., Xue, T., Freeman, W.T., Tenenbaum, J.: Learning A Probabilistic Latent Space of Object Shapes via 3D Generative-Adversarial Modeling. In: Advances in Neural Information Processing Systems. pp. 82–90 (2016)
53. Xiang, Y., Kim, W., Chen, W., Ji, J., Choy, C., Su, H., Mottaghi, R., Guibas, L., Savarese, S.: ObjectNet3D: A Large Scale Database for 3D Object Recognition. In: European Conference on Computer Vision. pp. 160–176 (2016)

54. Xiang, Y., Mottaghi, R., Savarese, S.: Beyond Pascal: A Benchmark for 3D Object Detection in the Wild. In: IEEE Winter Conference on Applications of Computer Vision. pp. 75–82 (2014)
55. Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes. In: Robotics: Science and Systems Conference. pp. 1–10 (2018)
56. Xiao, Y., Qiu, X., Langlois, P.A., Aubry, M., Marlet, R.: Pose from Shape: Deep Pose Estimation for Arbitrary 3D Objects. In: British Machine Vision Conference. pp. 120:1–120:14 (2019)
57. Yao, S., Hsu, T.M., Zhu, J.Y., Wu, J., Torralba, A., Freeman, W.T., Tenenbaum, J.: 3D-Aware Scene Manipulation via Inverse Graphics. In: Advances in Neural Information Processing Systems. pp. 1887–1898 (2018)
58. Zabulis, X., Lourakis, M.I.A., Stefanou, S.S.: 3D Pose Refinement Using Rendering and Texture-Based Matching. In: International Conference on Computer Vision and Graphics. pp. 672–679 (2014)
59. Zagoruyko, S., Komodakis, N.: Learning to Compare Image Patches via Convolutional Neural Networks. In: Conference on Computer Vision and Pattern Recognition. pp. 4353–4361 (2015)
60. Zakharov, S., Shugurov, I., Ilic, S.: DPOD: Dense 6D Pose Object Detector in RGB Images. In: International Conference on Computer Vision. pp. 1941–1950 (2019)

Geometric Correspondence Fields: Learned Differentiable Rendering for 3D Pose Refinement in the Wild

—

Supplementary Material

Alexander Grabner^{1,2}, Yaming Wang², Peizhao Zhang², Peihong Guo²,
Tong Xiao², Peter Vajda², Peter M. Roth¹, and Vincent Lepetit¹

¹ Graz University of Technology, Austria ² Facebook Inc., USA
{alexander.grabner, pmroth, lepetit}@icg.tugraz.at
{wym, stzpz, phg, xiaot, vajdap}@fb.com

In the following, we provide additional details and experimental results of our novel 3D pose refinement approach. In Sec. 6, we present our evaluation setup and discuss different datasets. In Sec. 7, we formally describe our evaluated metrics. In Sec. 8, we give specific details on the implementation of our approach. In Sec. 9, we analyze the iterative refinement behavior of different methods. In Sec. 10, we present detailed quantitative 3D pose refinement results for individual object categories. In Sec. 11, we show additional qualitative results. Finally, we analyze failure cases of our approach in Sec. 12.

6 Datasets and Evaluation Setup

We evaluate our proposed 3D pose refinement approach on the challenging Pix3D [46] dataset. The Pix3D dataset provides in-the-wild RGB images with 3D pose, 3D model, and focal length annotations for objects of different categories. We follow the evaluation protocol of previous work [10] and perform experiments on categories which have more than 300 non-occluded and non-truncated samples (*bed, chair, sofa, table*). Further, we restrict the training and evaluation to samples marked as non-occluded and non-truncated because all evaluated refinement methods lack explicit mechanisms to deal with occlusions. In addition, this dataset does not provide information on which objects parts are occluded nor the extent of the occlusion. For each category, we use 50% of the samples for training and the other 50% for testing as in [10].

Other category-level datasets do not provide annotations with sufficient accuracy to both train and evaluate fine-grained 3D pose refinement methods. For example, the Comp [51] and Stanford [51] datasets only provide coarse 3D pose annotations. In addition to coarse 3D pose annotations, the ScanNet [6], Pascal3D+ [54], and ObjectNet3D [53] datasets also just provide approximate 3D model annotations. Moreover, the latter two datasets assume constant camera intrinsics for images captured with different cameras which further decreases

the annotation quality [10]. As a consequence of this label noise, the training of refinement methods results in models with poor accuracy, while quantitative evaluations are not representative of the true refinement performance due to the lack of accuracy in the annotations.

In contrast, instance-level datasets like LineMOD [17], YCB [4], T-LESS [19], or NOCS [50] provide accurate annotations but have many images with strong occlusions. Neither traditional refinement methods [29, 35, 60], nor differentiable rendering based refinement methods [23, 33, 41], nor our approach employ explicit mechanisms to deal with occlusions. For example, one issue across all evaluated methods is that they align renderings to real-world images but occlusions are only present in the real-world images while the renderings are always un-occluded. Also, simply training methods based on feed-forward CNNs on occluded objects is in practice not sufficient to handle occlusions [40]. However, we specifically plan to address occlusions in the future by predicting occlusion masks and correspondence confidences.

In addition, our approach is specifically designed for category-level 3D pose refinement using untextured 3D models. This task is very different from instance-level 3D pose estimation where exactly matching colored and textured 3D models are available. In this case, methods which leverage color and texture information of 3D models have a clear advantage and should be used instead.

For these two reasons, we did not evaluate our method on instance-level datasets like LineMOD or YCB which have many images with strong occlusions and provide 3D models with colors and textures that exactly match those of the objects in the RGB images.

7 Metrics

We follow the evaluation protocol of previous work [10] and report the median error (*MedErr*) of multiple geometric distances:

Rotation: The 3D rotation distance

$$e_R = \frac{\|\log(\mathbf{R}_{\text{gt}}^T \mathbf{R}_{\text{pred}})\|_F}{\sqrt{2}} \quad (6)$$

represents the minimal angle between the ground truth rotation matrix \mathbf{R}_{gt} and the predicted rotation matrix \mathbf{R}_{pred} [49].

Translation: The 3D translation distance

$$e_t = \frac{\|\mathbf{t}_{\text{gt}} - \mathbf{t}_{\text{pred}}\|_2}{\|\mathbf{t}_{\text{gt}}\|_2} \quad (7)$$

gives the relative error between the ground truth translation \mathbf{t}_{gt} and the predicted translation \mathbf{t}_{pred} [20].

Pose: The 3D pose distance

$$e_{R,t} = \operatorname{avg}_{\mathbf{M}_i \in \mathcal{M}} \frac{d_{\text{bbox}}}{d_{\text{img}}} \cdot \frac{\|\operatorname{transf}(\mathbf{M}_i, \mathcal{P}_{\text{gt}}) - \operatorname{transf}(\mathbf{M}_i, \mathcal{P}_{\text{pred}})\|_2}{\|\mathbf{t}_{\text{gt}}\|_2} \quad (8)$$

represents the average normalized Euclidean distance of all transformed 3D model points in 3D space [18, 20]. Each 3D point \mathbf{M}_i of the ground truth 3D model \mathcal{M} is transformed using the ground truth 3D pose \mathcal{P}_{gt} and the predicted 3D pose $\mathcal{P}_{\text{pred}}$. This distance is normalized by the relative size of the object in the image using the ratio between the ground truth 2D bounding box diagonal d_{bbox} and the image diagonal d_{img} , and the L2-norm of the ground truth translation $\|\mathbf{t}_{\text{gt}}\|_2$.

Projection: The 2D projection distance

$$e_P = \operatorname{avg}_{\mathbf{M}_i \in \mathcal{M}} \frac{\|\operatorname{proj}(\mathbf{M}_i, \mathcal{P}_{\text{gt}}) - \operatorname{proj}(\mathbf{M}_i, \mathcal{P}_{\text{pred}})\|_2}{d_{\text{bbox}}} \quad (9)$$

is the average reprojection error normalized by the ground truth 2D bounding box diagonal d_{bbox} [51]. In this case, each 3D point \mathbf{M}_i of the ground truth 3D model \mathcal{M} is projected to the 2D image plane using the ground truth 3D pose \mathcal{P}_{gt} and the predicted 3D pose $\mathcal{P}_{\text{pred}}$ subject to a camera model. In this work, we assume the camera intrinsics to be known.

8 Implementation and Training Details

To train our refinement network, we resize and pad images to a spatial resolution of 256×256 while maintaining the aspect ratio. In this way, we are able to combine images with different aspect ratios in the same training batch. For our mapping branches, we adapt a ResNet-50 [14, 15] architecture. We utilize all layers up to the end of the first stage but use a stride of 1 for all convolutional layers and discard max pooling layers. For our correspondence branch, we use a channel dimensionality of 64 for all convolutional layers except the output layer.

During training of our network, we additionally employ different forms of data augmentation like mirroring, affine transformations, and independent pixel augmentations like additive or multiplicative noise for the RGB image I .

To regularize our network, we use L2 weight decay with a factor of $1e^{-5}$. We train our network $f(\cdot)$ for 1500 epochs using the Adam optimizer [25] with an initial learning rate of $\eta = 1e^{-3}$. We use a batch size of 8 and decrease the learning rate by a factor of 5 after 1000 and 1400 epochs.

During inference of our system, we compute geometry-level gradients $\nabla \mathbf{m}_i$ from our predicted geometric correspondence fields. In this way, vertices belonging to self-occluded triangles [23] or to visible triangles which are masked out by our geometric attention module do not receive gradients. However, since the geometry is fixed, providing gradients for a subset of all vertices is sufficient to perform 3D pose refinement.

9 Iterative Refinement

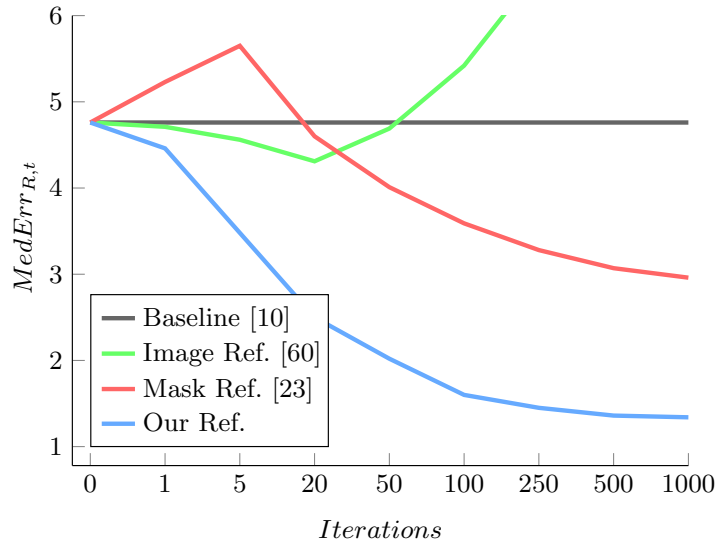


Fig. 7: Evaluation on 3D pose refinement after varying numbers of iterations. In contrast to other methods, our refinement achieves a consistent improvement upon the baseline, increasing with the number of iterations.

Figure 7 shows the performance of different refinement methods after varying numbers of iterations. In this experiment, we report the 3D pose metric $MedErr_{R,t}$ as a function of the number of 3D pose updates. The baseline does not perform iterative updates, thus, its $MedErr_{R,t}$ score is constant.

For *Image Refinement* [60], the accuracy increases until 20 iterations but then starts to decrease. After the first couple of coarse refinement steps, the predicted updates are not accurate enough to refine the 3D pose but start to jitter without further improving the 3D pose. Moreover, for many objects the prediction fails and the iterative updates cause the 3D pose to drift off which results in high $MedErr_{R,t}$ for large numbers of iterations.

For *Mask Refinement* [23], we observe an opposite effect. In the beginning, the accuracy decreases but then the performance increases. This is due to degenerated masks predicted from RGB images by Mask R-CNN [13]. The mask prediction often fails to capture fine-grained and thin structures, *e.g.*, ornaments and legs of a bed (see Figure 8). These degenerated masks cause large gradients during refinement and quickly pull the 3D pose away from the reasonable initial estimate predicted by the baseline (also see Figure 9, 2nd row). After five iterations the refinement on samples with correctly predicted masks counteracts this effect and the performance improves.

Table 3: Average computation times of different refinement methods for a single iteration using a Titan X GPU. For all evaluated methods, the execution time for computing a single 3D pose update is within the same order of magnitude.

Method	Time per Iteration
Image Refinement [60]	86.5 ms
Mask Refinement [23]	29.7 ms
Our Refinement	36.0 ms

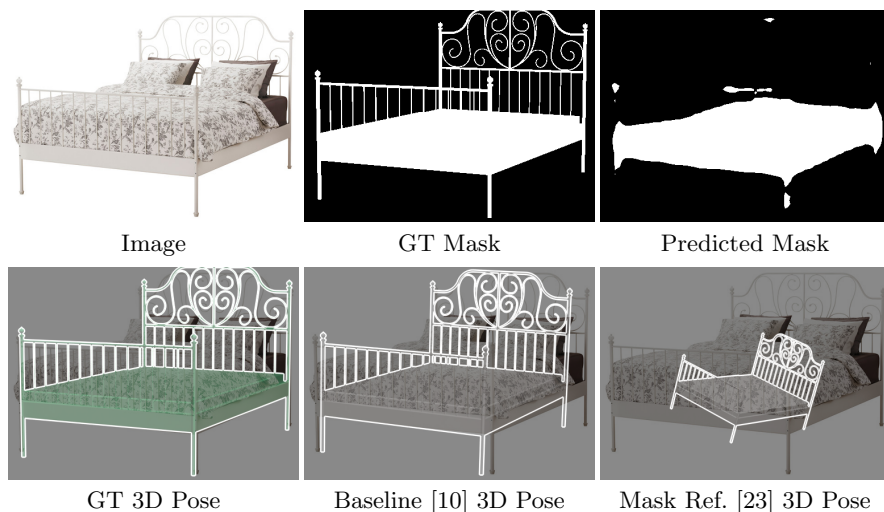


Fig. 8: Typical failure case of *Mask Refinement* [23]. Degenerated predicted masks (*top right*) cause large gradients during refinement and quickly pull the 3D pose away from reasonable initial estimates by the baseline (*bottom middle*). As a consequence, the 3D pose refinement using *Mask Refinement* fails (*bottom right*). Also see Figure 9, 2nd row.

In contrast to other refinement methods, our approach achieves a consistent improvement upon the baseline, increasing with the number of iterations. As expected, the accuracy saturates for large numbers of iterations. Empirically, we achieve maximum accuracy by performing 1000 3D pose updates using the Adam optimizer [25] with a learning rate of $\eta = 0.05$.

Finally, Table 3 compares the computation times of different refinement methods. *Image Refinement* evaluates two large ResNet-style networks [14] during inference and, thus, this method is the slowest in our evaluation. *Mask Refinement* generates a mask rendering, computes a loss in the mask space, and performs the backward pass of its differentiable renderer in each iteration. This method is the fastest in our evaluation since the target mask predicted from the input RGB image by Mask R-CNN [13] does not change during refinement. It is only predicted once before the iterative process and the inference time of Mask R-CNN is not considered in this experiment.

Our refinement is only marginally slower than *Mask Refinement* as we just evaluate our efficient network branches in addition to the forward and backward pass of our differentiable renderer in each iteration. However, all evaluated refinement methods show comparable execution time within the same order of magnitude for computing a single 3D pose update.

10 Detailed Quantitative Results

Tables 4 (*GT 3D models*) and 5 (*retrieved 3D models*) show detailed quantitative results for individual object categories on Pix3D. For completeness, we additionally report the detection accuracy $Acc_{D_{0.5}}$ which gives the percentage of objects for which the intersection over union between the ground truth 2D bounding box and the predicted 2D bounding box is larger than 50% [54]. We do not make 3D pose predictions for objects which are not detected by the baseline [10]. However, the reported *MedErr* metrics are computed over all samples, both detected and not detected.

Our refinement outperforms the baseline as well as competing refinement methods across all metrics. In fact, we do not only increase the *mean* performance over all categories but also achieve state-of-the-art results for each individual category. Using both ground truth (see Table 4) and retrieved (see Table 5) 3D models, we improve the performance compared to other methods by a large margin for each evaluated category.

Table 4: Detailed 3D pose refinement results for individual categories on Pix3D. In this experiment, we provide the **ground truth 3D model** for refinement. We outperform existing methods across all categories by a large margin.

Method	Category	Detection	Rotation	Translation	Pose	Projection
		$Acc_{D_{0.5}}$	$MedErr_R$ -1	$MedErr_t$ -10 ²	$MedErr_{R,t}$ -10 ²	$MedErr_P$ -10 ²
Baseline [10]	bed	99.0	5.07	6.68	5.18	3.42
Image Refinement [60]			4.65	5.45	4.60	3.38
Mask Refinement [23]			3.03	4.04	3.07	2.00
Our Refinement			2.40	1.84	1.45	1.28
Baseline [10]	chair	95.2	7.36	5.49	3.90	3.32
Image Refinement [60]			7.10	5.31	3.68	3.34
Mask Refinement [23]			4.42	4.89	3.17	1.79
Our Refinement			2.96	1.77	1.23	1.17
Baseline [10]	sofa	96.5	4.40	4.96	3.78	2.57
Image Refinement [60]			4.30	3.87	3.15	2.54
Mask Refinement [23]			2.97	2.89	2.25	1.54
Our Refinement			2.28	1.36	1.19	1.08
Baseline [10]	table	94.0	10.18	7.72	6.17	5.54
Image Refinement [60]			9.81	7.07	5.80	5.40
Mask Refinement [23]			3.81	4.44	3.34	2.27
Our Refinement			2.59	2.00	1.48	1.55
Baseline [10]	<i>mean</i>	96.2	6.75	6.21	4.76	3.71
Image Refinement [60]			6.46	5.43	4.31	3.67
Mask Refinement [23]			3.56	4.06	2.96	1.90
Our Refinement			2.56	1.74	1.34	1.27

Table 5: Detailed 3D pose refinement results for individual categories on Pix3D. In this experiment, we automatically **retrieve 3D models** for refinement using the method presented in [11]. We outperform existing methods across all categories by a large margin.

Method	Category	Detection	Rotation	Translation	Pose	Projection
		$Acc_{D_{0.5}}$	$MedErr_R$.1	$MedErr_t$.10 ²	$MedErr_{R,t}$.10 ²	$MedErr_P$.10 ²
Baseline [10]	bed	99.0	5.07	6.68	5.18	3.42
Image Refinement [60]			4.65	5.86	4.41	3.38
Mask Refinement [23]			4.40	5.32	4.21	2.69
Our Refinement			2.95	2.75	2.18	1.83
Baseline [10]	chair	95.2	7.36	5.49	3.90	3.32
Image Refinement [60]			7.15	5.21	3.67	3.35
Mask Refinement [23]			7.22	6.32	4.53	3.31
Our Refinement			4.89	2.87	2.04	2.19
Baseline [10]	sofa	96.5	4.40	4.96	3.78	2.57
Image Refinement [60]			4.34	3.75	3.02	2.54
Mask Refinement [23]			3.33	3.00	2.34	1.63
Our Refinement			2.60	1.60	1.42	1.19
Baseline [10]	table	94.0	10.18	7.72	6.17	5.54
Image Refinement [60]			9.73	7.23	6.22	5.68
Mask Refinement [23]			6.92	6.34	5.52	4.85
Our Refinement			4.73	3.38	2.93	3.49
Baseline [10]	<i>mean</i>	96.2	6.75	6.21	4.76	3.71
Image Refinement [60]			6.47	5.51	4.33	3.74
Mask Refinement [23]			5.47	5.25	4.15	3.12
Our Refinement			3.79	2.65	2.14	2.18

11 Additional Qualitative Results

Figures 9, 10, and 11 show additional qualitative 3D pose refinement results for different methods complementary to those presented in the main paper. While other methods fail to predict fine-grained 3D poses in the wild, our approach precisely aligns 3D models to objects in RGB images which results in significantly improved 3D poses for objects of different categories. In many cases, our predicted 3D pose is visually indistinguishable from the ground truth 3D pose.

Figures 12 and 13 show additional qualitative results of our predicted geometric correspondence fields. Our predicted 2D displacement vectors are highly accurate for many different objects and scales in the wild. The illustrations also show our computed geometric attention masks, outlined in white underneath the 2D displacement vectors.

12 Additional Failure Cases

Figure 14 shows additional failure cases of our approach. In the presence of strong image noise, we cannot predict accurate geometric correspondence fields and, thus, our refinement fails. Also, if there are duplicate or ambiguous structures in the image our method sometimes predicts wrong correspondences and aligns the 3D model to unintended image parts.




Fig. 9: Additional qualitative 3D pose refinement results for objects of different categories. We project the ground truth 3D model on the image using the predicted 3D pose. Our approach overcomes the limitations of previous methods and predicts fine-grained 3D poses which are in many cases visually indistinguishable from the ground truth. Best viewed in **digital zoom**.




Fig. 10: Additional qualitative 3D pose refinement results for objects of different categories. We project the ground truth 3D model on the image using the predicted 3D pose. Our approach overcomes the limitations of previous methods and predicts fine-grained 3D poses which are in many cases visually indistinguishable from the ground truth. Best viewed in **digital zoom**.





Fig. 11: Additional qualitative 3D pose refinement results for objects of different categories. We project the ground truth 3D model on the image using the predicted 3D pose. Our approach overcomes the limitations of previous methods and predicts fine-grained 3D poses which are in many cases visually indistinguishable from the ground truth. Best viewed in **digital zoom**.



Initial 3D Pose Ground Truth Prediction

Fig. 12: Additional qualitative examples of our predicted geometric correspondence fields. Our predicted 2D displacement vectors are highly accurate for many different objects and scales. Best viewed in **digital zoom**.



Fig. 13: Additional qualitative examples of our predicted geometric correspondence fields. Our predicted 2D displacement vectors are highly accurate for many different objects and scales. Best viewed in **digital zoom**.

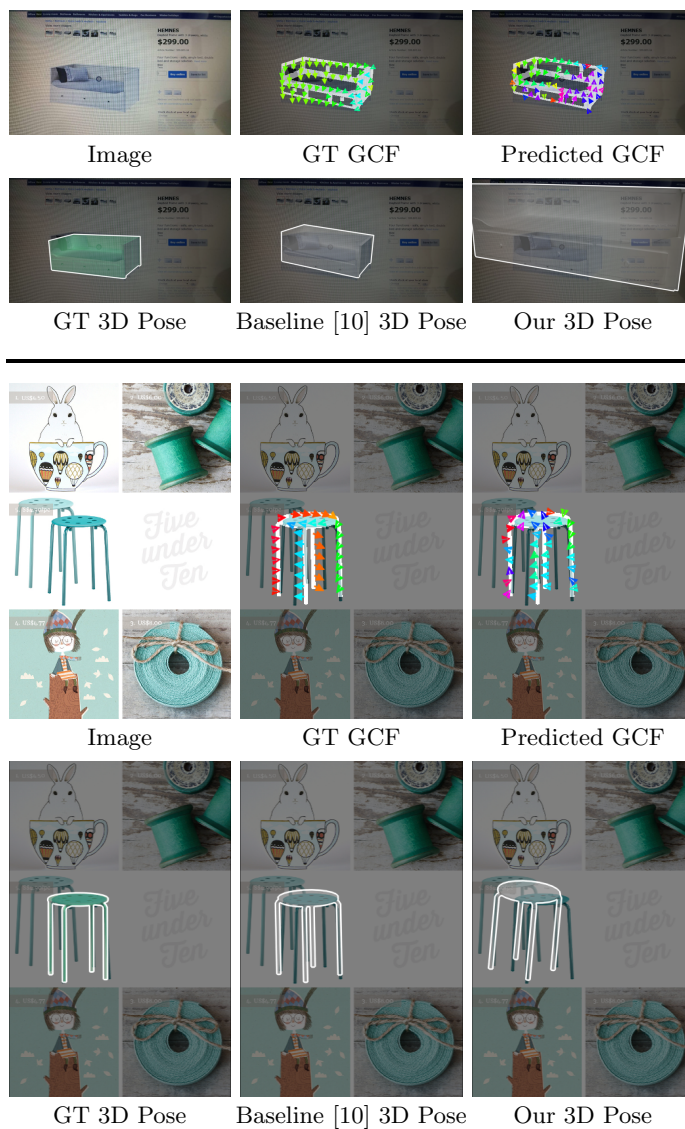


Fig. 14: Failure cases of our approach. In the presence of strong image noise (*top example*), we cannot predict accurate geometric correspondence fields (GCF) and, thus, our refinement fails. Also, if there are duplicate or ambiguous structures in the image our method sometimes predicts wrong correspondences and aligns the 3D model to unintended image parts (*bottom example*). Best viewed in **digital zoom**.