

RNNPose: Recurrent 6-DoF Object Pose Refinement with Robust Correspondence Field Estimation and Pose Optimization

Yan Xu¹ Kwan-Yee Lin¹ Guofeng Zhang² Xiaogang Wang¹ Hongsheng Li¹

¹The Chinese University of Hong Kong ²State Key Lab of CAD&CG, Zhejiang University

Abstract

6-DoF object pose estimation from a monocular image is challenging, and a post-refinement procedure is generally needed for high-precision estimation. In this paper, we propose a framework based on a recurrent neural network (RNN) for object pose refinement, which is robust to erroneous initial poses and occlusions. During the recurrent iterations, object pose refinement is formulated as a non-linear least squares problem based on the estimated correspondence field (between a rendered image and the observed image). The problem is then solved by a differentiable Levenberg-Marquardt (LM) algorithm enabling end-to-end training. The correspondence field estimation and pose refinement are conducted alternatively in each iteration to recover the object poses. Furthermore, to improve the robustness to occlusion, we introduce a consistency-check mechanism based on the learned descriptors of the 3D model and observed 2D images, which downweights the unreliable correspondences during pose optimization. Extensive experiments on LINEMOD, Occlusion-LINEMOD, and YCB-Video datasets validate the effectiveness of our method and demonstrate state-of-the-art performance.

1. Introduction

6-DoF object pose estimation is of crucial importance in various applications, including augmented reality, robotic manipulation, and autonomous driving. Influenced by varying illuminations and occlusions, appearances of the differently posed objects may vary significantly from different views, which poses great challenges for 6-DOF object pose estimation from a single color image.

The recent top-performing methods [23,27,41,52,58] additionally include a pose refinement procedure which substantially improves the performance. Some of these frameworks [41, 52] rely on depth sensors and refine the poses with the ICP algorithm [2]. To avoid the expensive depth sensor, Li *et al.* [23] and Manhardt *et al.* [27] pioneered the RGB-based pose refinement. During refinement, these methods first render a reference color image according to the coarse pose estimate. This rendered image along with

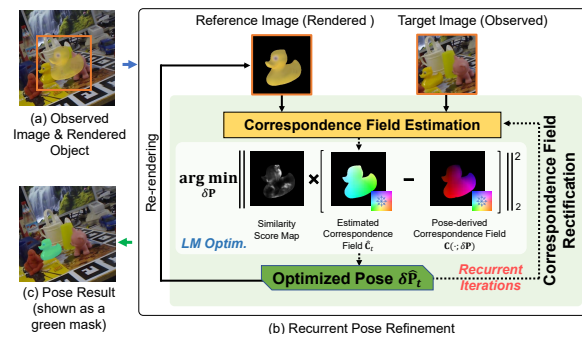


Figure 1. The basic idea. (a) Before refinement, a reference image is rendered according to the object initial pose (shown in a fused view). (b) Our RNN-based framework recurrently refines the object pose based on the estimated correspondence field between the reference and target images. The pose is optimized to be consistent with the reliable correspondence estimations highlighted by the similarity score map (built from learned 3D-2D descriptors) via differentiable LM optimization. (c) The output refined pose.

the observed image is then fed to a CNN to directly predict the residual pose for refining the coarse pose [23, 27, 58]. While these methods perform well in ideal scenarios based on massive training data, the pose regression becomes less stable in practice. More recently, Iwase *et al.* [20] formulated the object pose refinement as an optimization problem based on feature alignment, and reported significant performance improvements. In their work, the encoded features of a 3D model by a neural network are projected to the 2D image plane according to the pose parameters. Thereafter, the pose optimization is conducted by aligning the projected features with the observed target image features. As the pose optimization depends on the gradients from the pixel-level feature differences, the feature alignment based methods are only applicable to small inter-frame pose variations [57] and are not quite robust with erroneous initial poses. Moreover, Iwase *et al.* [20] still have a limited design for occlusion handling, which might limit the deployment scope.

In this work, we propose a recurrent object pose refinement framework, dubbed RNNPose, which is robust to erroneous initial poses and occlusions. The overall pipeline

K. Lin and H. Li are the co-corresponding authors.

is illustrated in Fig. 1. Before refinement, a reference image of the object is rendered according to the initial pose estimation. Our refinement module refines the initial pose based on this rendered image and the observed image. To increase the tolerance to erroneous initial poses, our refinement is conducted within a recurrent framework, where the pose optimization is formulated as a non-linear least squares problem based on estimated correspondence fields. In each recurrent iteration, the dense correspondences between the rendered image and observed image are estimated, and the object pose is then optimized to be consistent with the correspondence field estimation. The architecture of our correspondence estimation is inspired by the recent optical flow estimation techniques [39, 44], which is integrated with our pose optimization recurrently. To suit our task where unpatterned objects and illumination variations are ubiquitous, we further include a correspondence field rectification step in each recurrent iteration based on the currently optimized pose. The inconsistent correspondences are rectified by enforcing rigid-transformation constraints. The rectified correspondence field is also used to initialize the next recurrent iteration to improve the robustness further.

For occlusion handling, we introduce a 3D-2D hybrid network trained with a contrastive loss, which generates distinctive point-wise descriptors for the 3D object model and observed 2D images. A similarity score is constructed for each estimated correspondence pair based on the learned descriptors, with which to downweight the unreliable correspondences during pose optimization. The pose optimization is conducted by a differentiable Levenberg-Marquardt (LM) algorithm (sharing the ideas of [42, 43]) for end-to-end training.

Our contributions are three-fold: 1) We propose an RNN-based 6-DoF pose refinement framework that is robust to large initial pose errors and occlusions. During recurrent iterations, the pose optimization is formulated as a non-linear least squares problem based on the estimated correspondence field. Meanwhile, the correspondence field is also being rectified and improved by the optimized pose for robustness. 2) To handle the occlusions, a 3D-2D hybrid network is introduced to learn point-wise descriptors which are used to downweight unreliable correspondence estimations during pose optimization. 3) We achieve new state-of-the-art performances on LINEMOD, Occlusion LINEMOD, and YCB-Video datasets. Our code is public at <https://github.com/DecaYale/RNNPose>.

2. Related Work

6-DoF object pose estimation. 6-DoF object pose estimation systems (usually going after an object detector [5, 11, 28, 47, 54–56, 59]) aim to estimate the 3-DoF orientations and 3-DoF locations of rigid objects. The boom of deep learning has significantly improved object pose estimation in recent years. Some methods proposed to directly regress object poses from monocular color images [7, 16, 21, 24, 51, 52] or with the aid from depth sen-

sors [13, 49, 50]. They leveraged CNNs’ regression ability to directly map the observed images to object poses. More recently, correspondence-based methods [6, 15, 33–35, 45] become more popular. They employed CNNs to estimate the corresponding 3D model point for each observed object pixel, and then solve for poses with PnP [10]. These methods may estimate the object’s bounding box corners [35, 45], predict dense 2D-3D correspondence maps [33] or vote the keypoints by all object pixels [34]. More recently, EPOS [15] proposed to handle symmetric objects by segmenting 3D models into patches and estimating the patch centers. The above direct object pose estimation methods usually become less stable when varying illuminations and occlusions exist. Several methods [20, 23, 27, 41, 52, 58] hence conducted pose refinement based on the estimated coarse initial pose above, which achieved significant performance gains. Some of these methods [41, 52] relied on depth data from costly sensors and utilized ICP to align the known object model to the observed depth image. While [23, 27, 58] first rendered a 2D object image according to the initial pose and then compared the rendered image with the observed image via a CNN to estimate the residual pose. These RGB-based methods are especially attractive due to their economical nature. However, most of these methods need massive training data and are not quite robust in practical scenarios. Moreover, they need a cumbersome CNN for pose regression, which sacrifices efficiency. Iwase *et al.* [20] proposed to alleviate such dilemma by reusing the images features extracted by CNN and attained real-time processing. Concretely, they employed the CNN as an image feature encoder, based on which to formulate a non-linear optimization problem to align the features from the inference and target images for pose refinement inspired by BA-Net [42]. Though efficient, their formulation is built upon overlapped object regions across the reference image and target image, which may thus be less stable with erroneous initial pose inputs. The previous work [1] proposed to refine the pose based on the correspondences, but their method is still limited to ideal scenarios.

Non-linear least squares optimization with deep learning. Non-linear least squares optimization algorithms, such as Gauss-Newton [30] and Levenberg-Marquardt [31], are widely used in computer vision [17, 25, 32, 53], given their efficient and effective nature. Recently, the differentiability of the optimization algorithm itself has been widely studied and several works [37, 42, 43, 48] have included the differentiable optimization algorithm during the network training for localization systems and visual SLAMs. These inspire our formulation for object pose refinement.

3. Method

Given an observed object image \mathbf{I}_{obs} , an initial object pose estimate \mathbf{P}_{init} and the object’s CAD model \mathcal{M} as inputs, a 6-DoF pose refinement system aims to further improve the object pose estimation. In this paper, we propose a recurrent pose refinement method, dubbed RNNPose,

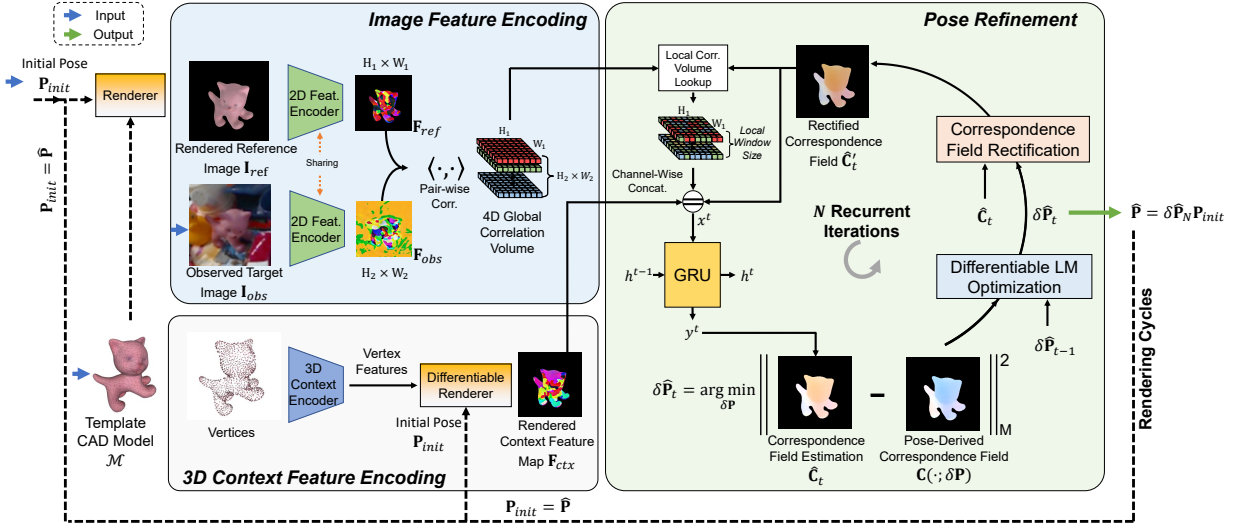


Figure 2. Overview of the proposed method. For pose refinement, a reference image \mathbf{I}_{ref} is rendered with the object CAD model and its initial pose \mathbf{P}_{init} . The **image feature encoding** module encodes the rendered image \mathbf{I}_{ref} and the observed image \mathbf{I}_{obs} to feature maps and build a 4D global correlation volume. In parallel, the **3D context feature encoding** module encodes the 3D model geometry and render the 3D features to a 2D context feature map \mathbf{F}_{ctx} according to the initial pose estimation. During **pose refinement**, the correspondence field $\hat{\mathbf{C}}_t$ and the residual pose $\delta \hat{\mathbf{P}}_t$ are alternately estimated in a recurrent framework. After the LM pose optimization, correspondence field estimation $\hat{\mathbf{C}}_t$ is rectified as $\hat{\mathbf{C}}'_t$ by enforcing rigid-transformation constraints with the currently optimized pose $\delta \hat{\mathbf{P}}_t$ to further improve next-iteration estimations. After N recurrent iterations, the reference image \mathbf{I}_{ref} is re-rendered with the current pose estimation.

which is robust to erroneous initial poses and occlusions. Our method is based on a rendering pipeline and may have several rendering cycles as illustrated by Fig. 2. At the beginning of the first rendering cycle, a reference image \mathbf{I}_{ref} is rendered with the object’s CAD model according to its initial pose \mathbf{P}_{init} (estimated by any direct methods [34,52]). Then, the rendered reference image, the observed target image, and the vertices of the CAD model are encoded as high-dimensional features which will be used to estimate the correspondences (between the rendered image and observed image) in the follow-up pose refinement module. The pose refinement module constitutes our major contribution, where we formulate an optimization problem based on the correspondence estimations. We integrate correspondence field estimation and pose refinement into a recurrent framework for robustness and efficiency. To handle occlusions, we generate point-wise distinctive descriptors for the 3D object model and observed images with a 3D-2D hybrid network, with which to downweight the unreliable correspondences during pose optimization. After every several recurrent iterations, the reference image \mathbf{I}_{ref} is re-rendered with the currently optimized pose to decrease the pose gap to the target for the next cycle.

In the ensuing subsections, we will detail the feature extraction (Sec. 3.1), recurrent pose refinement (Sec. 3.2), and the loss functions for training (Sec. 3.3).

3.1. 2D-3D Feature Encoding and Rendering

The rendered reference image \mathbf{I}_{ref} and observed target image \mathbf{I}_{obs} first need to be encoded into high-dimensional feature maps \mathbf{F}_{ref} and \mathbf{F}_{obs} for the follow-up feature corre-

lation volume construction [9,18,19,39,44]. The correlation volume encodes the appearance similarities between image pixels, which is essential for correspondence reasoning. In our work, we adopt several residual blocks [12] for image feature encoding, and the pair-wise correlations of the encoded features are calculated to create a global correlation volume. The global correlation volume will be frequently queried for correspondence field estimation in the follow-up pose refinement module.

Besides the pair-wise correlation volume, popular dense correspondence estimation methods also incorporate context features of the reference image for guidance. As shown in Fig. 2, to better encode the geometric contexts, unlike previous methods encoding the context features from 2D images, we directly encode the features from 3D object point clouds with a *3D context feature encoder* based on KPConv [46]. The point-wise geometric features are then rendered as a 2D context feature map \mathbf{F}_{ctx} according to the initial object pose estimation. Here, we adopt a differentiable renderer [36] for feature rendering to enable geometric feature learning. We empirically found that encoding the context features from point clouds brings more robustness. Besides, the vertex features only need to be extracted once per object model and archived for inference after training, which is quite efficient.

3.2. Recurrent Correspondence Field Estimation and 6-DoF Pose Refinement

Based on the constructed correlation volume and encoded context features, we propose a 6-DoF object pose refinement system by integrating the correspondence estima-

tion and pose optimization as a recurrent framework. The correspondence field estimation and pose optimization rely on each other and improve recurrently for robust pose refinement. The basic pipeline is illustrated in the pose refinement module in Fig. 2.

3.2.1 Correspondence Field Estimation

For correspondence field estimation, we adopt a network architecture similar to RAFT [44] but make major modifications to suit our task, *i.e.*, including the 3D context feature encoding (Sec. 3.1) and correspondence rectification (Sec. 3.2.2). At the beginning of each recurrent iteration, for each pixel of the reference image, we first look up and collect (from the global correlation volume) its correlation values with the candidate pixels in the target image. The candidate pixels are within a square local window centered at the estimated correspondences from the previous iteration. The collected correlations are then reshaped as a local correlation volume (a 2D map) spatially aligned with the reference image. In the first iteration, we use an all-zeros correspondence field to bootstrap correlation candidate identification, while in the later iterations, the rectified correspondence field (to be elaborated in Sec. 3.2.2) is used.

After the correlation lookup, the collected local correlation volume, the rectified correspondence field, and the previously encoded context feature map $\mathbf{F}_{ct,x}$ are concatenated as inputs to a GRU network to estimate the correspondence field $\hat{\mathbf{C}}_t$ for the current (t -th) recurrent iteration.

3.2.2 6-DoF Pose Refinement

Basic Formulation. Given a reference image (with depth map) and a target image, the ground-truth correspondence field of the reference image can be derived based on the ground-truth residual pose $\delta\mathbf{P}^{gt}$ point-wisely:

$$\mathbf{C}(\mathbf{x}^i; \delta\mathbf{P}^{gt}) = \pi(\delta\mathbf{P}^{gt} \pi^{-1}(\mathbf{x}^i, z^i)), \quad (1)$$

where $\mathbf{C}(\mathbf{x}^i; \delta\mathbf{P}^{gt}) \in \mathbb{R}^2$ denotes the ground-truth correspondence field value of point \mathbf{x}^i , and z^i denotes the associated rendered depth value. Here, $\pi(\cdot)$ and $\pi^{-1}(\cdot; z^i)$ are the projection (3D-to-2D) and inverse projection (2D-to-3D) functions of a pinhole camera model.

To estimate the residual pose, we take the correspondence field $\hat{\mathbf{C}}_t$ estimated by the GRU as an approximation of its ground-truth, *i.e.*, $\hat{\mathbf{C}}_t(\mathbf{x}^i) \approx \mathbf{C}(\mathbf{x}^i; \delta\mathbf{P}^{gt})$, and push the correspondence field derived by the pose argument $\delta\mathbf{P}$, *i.e.*, $\mathbf{C}(\mathbf{x}^i; \delta\mathbf{P})$, close to the GRU’s estimation by optimizing $\delta\mathbf{P}$. In this way, the residual pose parameter $\delta\mathbf{P}$ will approximate the ground-truth $\delta\mathbf{P}^{gt}$ after the optimization. The specific formulation is a non-linear least squares problem and the objective function is expressed as

$$E(\boldsymbol{\xi}) = \sum_{i=1}^M (\hat{\mathbf{C}}_t(\mathbf{x}^i) - \mathbf{C}(\mathbf{x}^i; \boldsymbol{\xi}))^T (\hat{\mathbf{C}}_t(\mathbf{x}^i) - \mathbf{C}(\mathbf{x}^i; \boldsymbol{\xi})), \quad (2)$$




Figure 3. With the rectified correspondences, the related local correlation windows are accordingly shifted to better locations, which improves the estimation in the next recurrent iteration.

where the residual pose argument $\delta\mathbf{P}$ is parameterized as its minimal representation $\boldsymbol{\xi} \in \mathfrak{se}(3)$ (of the associated Lie-algebra) during optimization. $\hat{\mathbf{C}}_t$ is the GRU-estimated correspondence field at the t -th recurrent iteration, and $\mathbf{C}(\mathbf{x}^i; \boldsymbol{\xi})$ denotes the correspondence of point \mathbf{x}^i derived with the pose parameter argument $\boldsymbol{\xi}$, and M is the total number of object points in the rendered reference image.

Handling Unreliable Correspondences with Similarity Scores. The formulation of Eq. (2) is based on an impractical assumption that the correspondence field $\hat{\mathbf{C}}_t$ can be reliably estimated for all foreground regions, which is extremely difficult considering ubiquitous occlusions. We further propose to incorporate a consistency-check mechanism to downweight the unreliable values in $\hat{\mathbf{C}}_t$ during pose optimization. To model the reliability of estimated correspondence, one option is to adopt a forward-and-backward consistency check [26, 29]. However, the bidirectional consistency check doubles the computational cost, and the domain gap between the rendered images and the real images increases the learning difficulty.

We therefore propose a descriptor-based consistency check to alleviate the dilemma. The basic idea is to represent the 3D object model \mathcal{M} and the observed 2D target image \mathbf{I}_{obs} as two sets of distinctive descriptors point-wisely via a 3D-2D hybrid network (with KPConv [46] and a key-point description net [8] as backbones). The corresponding descriptors of the object model and object images are enforced to be similar, while the non-corresponding descriptors are enforced to be dissimilar (by training with a contrastive descriptor loss function being described in Sec. 3.3). The learned 3D model descriptors are rendered as 2D feature maps, denoted as $\mathbf{D}_{\mathcal{M}}$, according to the object pose of the reference image for fast indexing. The encoded target image descriptor map is denoted as $\mathbf{D}_{\mathcal{I}}$.

With these high-dimensional distinctive descriptors, for each estimated correspondence pair $(\mathbf{x}^i, \hat{\mathbf{C}}_t(\mathbf{x}^i))$, we measure its reliability according to the similarity between their associated 3D and 2D descriptors $(\mathbf{d}_{\mathcal{M}}^i, \mathbf{d}_{\mathcal{I}}^i)$. $\mathbf{d}_{\mathcal{M}}^i$ and $\mathbf{d}_{\mathcal{I}}^i$ here are collected from the above descriptor maps: $\mathbf{d}_{\mathcal{M}}^i = \mathbf{D}_{\mathcal{M}}(\mathbf{x}^i)$ and $\mathbf{d}_{\mathcal{I}}^i = \mathbf{D}_{\mathcal{I}}(\hat{\mathbf{C}}_t(\mathbf{x}^i))$, where bilinear interpolation may be applied for non-integer correspondence coordinates. The reliability of this correspondence pair is modeled

with a similarity score:

$$w^i = \exp\left(-\frac{|1 - \mathbf{d}_{\mathcal{M}}^{i\top} \mathbf{d}_{\mathcal{I}}^i|}{\sigma}\right), \quad (3)$$

where σ is a learnable parameter (initialized with 1) adjusting the sharpness. The similarity scores are used as the weights of the Mahalanobis distance measurements in Eq. (2), which effectively downweight unreliable correspondences during optimization. By introducing a diagonal weighting matrix $\mathbf{w}^i = \begin{pmatrix} w^i & 0 \\ 0 & w^i \end{pmatrix}$, the weighted version of Eq. (2) is written as

$$E(\boldsymbol{\xi}) = \sum_{i=1}^M (\hat{\mathbf{C}}_t(\mathbf{x}^i) - \mathbf{C}(\mathbf{x}^i; \boldsymbol{\xi}))^T \mathbf{w}^i (\hat{\mathbf{C}}_t(\mathbf{x}^i) - \mathbf{C}(\mathbf{x}^i; \boldsymbol{\xi})). \quad (4)$$

The pose optimization is thus formulated as

$$\hat{\boldsymbol{\xi}} = \arg \min_{\boldsymbol{\xi}} E(\boldsymbol{\xi}), \quad (5)$$

where the pose parameter $\boldsymbol{\xi} \in \mathfrak{se}(3)$ is optimized by minimizing the objective function defined by Eq. (4).

Differentiable Residual Pose Optimization. We solve the non-linear least squares problem (Eq. (5)) with Levenberg-Marquardt (LM) algorithm. For the optimization in the t -th recurrent iteration, the pose parameter is initialized with the estimated pose from the previous iteration *i.e.*, $\boldsymbol{\xi}_0 = \log(\delta \hat{\mathbf{P}}_{t-1})$. Continuing from the parameter $\boldsymbol{\xi}_{p-1}$ of the previous LM iteration, the left-multiplied increment $\Delta \boldsymbol{\xi}_p$ is computed by

$$\Delta \boldsymbol{\xi}_p = (\mathbf{J}^T \mathbf{W} \mathbf{J} + \lambda \mathbf{I})^{-1} \mathbf{J}^T \mathbf{W} \mathbf{r}(\boldsymbol{\xi}_{p-1}), \quad (6)$$

with which we update the parameter as $\boldsymbol{\xi}_p \leftarrow \Delta \boldsymbol{\xi}_p \circ \boldsymbol{\xi}_{p-1}$, to approach the optimal solution. Here, $\mathbf{J} = -\frac{\partial \mathbf{r}}{\partial \boldsymbol{\xi}}$ is the Jacobian matrix containing the derivative of the stacked residual vector $\mathbf{r} = (r_1, r_2, \dots, r_{2M})^T$ (established from Eq. (4)) with regard to a left-multiplied increment. We unroll the parameter update procedure and make the LM optimization layer differentiable to enable end-to-end network training. The differentiable optimization procedure enhances the feature learning for correspondence field estimation, which is essential to high performance. After LM optimization, the residual pose of the t -th recurrent iteration is estimated as $\delta \hat{\mathbf{P}}_t = \exp(\hat{\boldsymbol{\xi}})$, where $\hat{\boldsymbol{\xi}}$ denotes the optimized parameter after several updates with Eq. (6).

Correspondence Field Rectification. The erroneous initial poses usually produce large offsets between the rendered reference object and the observed object, which poses challenges for correspondence estimation. Moreover, unlike the standard scenarios of optical flow estimation [18, 19, 39, 44], unpatterned objects and varying illuminations are ubiquitous in our task, which further increases the difficulty. Considering the optimized pose by Eq. (5) is mainly supported by the reliable correspondence estimations with our weight-

ing mechanism Eq. (3), we rectify the correspondence field as $\hat{\mathbf{C}}'_t(\mathbf{x}) = \pi(\delta \hat{\mathbf{P}}_t \pi^{-1}(\mathbf{x}; z))$ based on the currently optimized pose $\delta \hat{\mathbf{P}}_t$. The rectification enforces the rigid-transformation constraints among the correspondence field, which improves the overall correspondence quality for the correlation volume lookup in the following recurrent iteration. A toy example is shown in Fig. 3 for better understanding.

Object Pose Estimation Update. After every N recurrent iterations, the residual pose is estimated as $\delta \hat{\mathbf{P}}_N$ by the RNN. We update the object pose estimation with the estimated residual pose $\delta \hat{\mathbf{P}}_N$ as $\hat{\mathbf{P}} \leftarrow \delta \hat{\mathbf{P}}_N \mathbf{P}_{init}$, and we re-render the reference image \mathbf{I}_{ref} based on this updated pose to start the next N -recurrent-iteration refinement, as illustrated in Fig. 2. We refer to the N -recurrent-iteration refinement as a *rendering cycle*, and the initial pose \mathbf{P}_{init} for the next cycle is set to $\hat{\mathbf{P}}$ accordingly. The performance and efficiency with different rendering cycles and recurrent iterations will be discussed in Sec. 4.2.

3.3. Loss Functions

Model Alignment Loss. To supervise the residual pose estimations $\{\delta \hat{\mathbf{P}}_t | t = 1 \dots N\}$ generated in each rendering cycle (including N recurrent iterations), we apply these residual poses as the left-multiplied increments to the initial pose \mathbf{P}_{init} , having the corresponding object pose estimations $\{\hat{\mathbf{P}}_t | t = 1 \dots N\}$, where $\hat{\mathbf{P}}_t = \delta \hat{\mathbf{P}}_t \mathbf{P}_{init}$. Thereafter, we adopt a 3D model alignment loss to supervise these pose estimations for each rendering cycle:

$$L_{ma} = \sum_{t=1}^N \|\hat{\mathbf{P}}_t \mathbf{X}_{model} - \mathbf{P}^{gt} \mathbf{X}_{model}\|_1, \quad (7)$$

where $\hat{\mathbf{P}}_t$ is the object pose estimation mentioned above and \mathbf{P}^{gt} denotes the ground-truth pose. Here, $\mathbf{X}_{model} \in \mathbb{R}^{4 \times M}$ contains homogeneous coordinates of the M model points. This loss function encourages the pose estimation to be close to the ground-truth so that the transformed model points can be well aligned.

Correspondence Loss. We adopt $L1$ loss [44] for correspondence field supervision, where the ground-truth correspondence fields are derived with Eq. (1) based on ground-truth poses.

Descriptor Loss. We use circle loss \mathcal{L}_{cir} [40] as the contrastive loss to supervise the point-wise descriptor learning of the 3D object model and the target images for similarity score calculation Eq. (3). Concretely, we view the target image \mathbf{I}_{obs} as two parts, *i.e.*, the foreground region (object region) denoted as $fg(\mathbf{I}_{obs})$ and the background region denoted as $bg(\mathbf{I}_{obs})$. For each foreground descriptor $\mathbf{d}_{\mathcal{I}}^i \in fg(\mathbf{I}_{obs})$, we first find a set of its corresponding 3D descriptors $\{\mathbf{d}_{\mathcal{M}}^j\}_+$ of object model via KNN searching (see supplementary materials for details). Then, $\mathbf{d}_{\mathcal{I}}^i \in fg(\mathbf{I}_{obs})$ is enforced to be similar to $\{\mathbf{d}_{\mathcal{M}}^j\}_+$ and dissimilar to the remaining non-corresponding descriptors

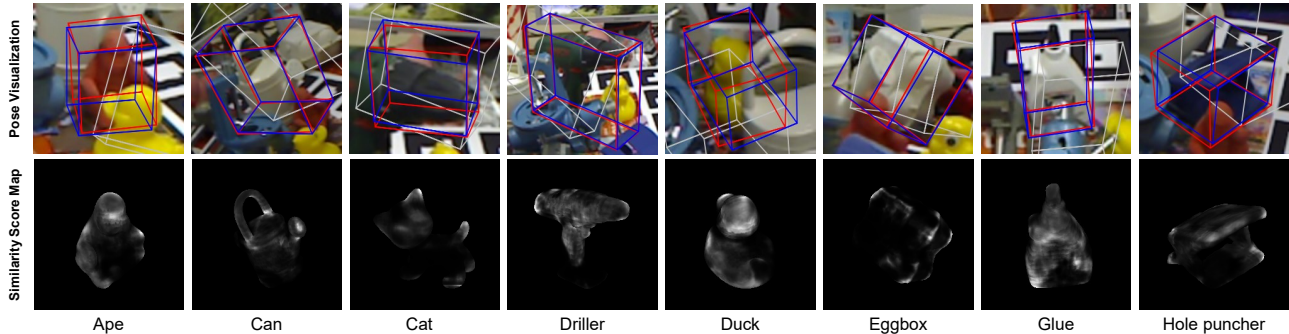


Figure 4. Visualization of our pose estimations (first row) on Occlusion LINEMOD dataset and the similarity score maps (second row) for downweighting unreliable correspondences during pose optimization. For pose visualization, the white boxes represent the erroneous initial poses, the red boxes are estimated by our algorithm and the ground-truth boxes are in blue. Here, the initial poses for pose refinement are originally from PVNet [34] but added with significant disturbances for robustness testing.

$\{\mathbf{d}_{\mathcal{M}}^k\}_-$ with circle loss \mathcal{L}_{cir} [40], which is expressed as $\mathcal{L}_{cir}(\mathbf{d}_{\mathcal{T}}^i, \{\mathbf{d}_{\mathcal{M}}^j\}_+, \{\mathbf{d}_{\mathcal{M}}^k\}_-)$. Moreover, for background descriptors $\mathbf{d}_{\mathcal{T}}^i \in bg(\mathbf{I}_{obs})$, we constrain them to be similar to each other in the background, while to be dissimilar to the foreground descriptor set $fg(\mathbf{I}_{obs})$ with loss $\mathcal{L}_{cir}(\mathbf{d}_{\mathcal{T}}^i, bg(\mathbf{I}_{obs}), fg(\mathbf{I}_{obs}))$. Traversing all target image descriptors $\mathbf{d}_{\mathcal{T}}^i$, the descriptor loss is calculated as

$$L_d = \sum_{\mathbf{d}_{\mathcal{T}}^i \in fg(\mathbf{I}_{obs})} \mathcal{L}_{cir}(\mathbf{d}_{\mathcal{T}}^i, \{\mathbf{d}_{\mathcal{M}}^j\}_+, \{\mathbf{d}_{\mathcal{M}}^k\}_-) + \sum_{\mathbf{d}_{\mathcal{T}}^i \in bg(\mathbf{I}_{obs})} \mathcal{L}_{cir}(\mathbf{d}_{\mathcal{T}}^i, bg(\mathbf{I}_{obs}), fg(\mathbf{I}_{obs})) \quad (8)$$

to supervise the descriptor learning. With the contrastive learning, the corresponding 2D-3D descriptors would be similar while the noncorresponding ones would be dissimilar, which provides the foundation for unreliable correspondence handling with similarity scores Eq. (3).

4. Experiments

4.1. Experimental Setup

Implementation Details. We train all of our networks end-to-end using the Adam [22] optimizer with an initial learning rate of 10^{-4} and adjust it with a cosine annealing strategy. The weights of model alignment loss L_{ma} and descriptor loss L_d are set to 1, while the correspondence loss weight is set to 0.5. During training, we conduct 3 rendering cycles, each of which performs 4 recurrent refinement iterations for pose refinement. All our models are trained agnostic to the initial pose sources where disturbed ground-truth poses are taken as initial poses for training following [23]. For testing, we conduct the same numbers of rendering cycles and refinement iterations as those during training for most experiments if without further declaration, though more iterations could produce better results. Please refer to the supplementary materials for more details.

Datasets. We evaluate our method on three datasets, including LINEMOD [14], Occlusion LINEMOD [3] and YCB-Video [52]. LINEMOD is a standard benchmark for 6D

object pose estimation. This dataset contains texture-less objects in cluttered scenes captured with challenging illuminance variations. The Occlusion LINEMOD is a subset of LINEMOD dataset with additional annotations for occluded objects, which is suitable for testing the robustness to severe occlusions. Besides, the YCB-Video dataset contains the images of the YCB object set [4] where strong occlusions, clutters are exhibited. It includes more than 110k real images captured for 21 objects with or without textures. We follow similar conventions in data processing and synthetic data generation as the previous works [20,34]. For the initial poses, we mainly rely on PoseCNN [52] and PVNet [34], two typical direct estimation methods, following [23] and [20]. We also create a set of extremely erroneous initial poses by adding random Gaussian noise to the original initial pose estimations to evaluate the robustness to large initial poses errors.

Evaluation Metrics. We evaluate our method with the metrics ADD(-S) [14] and AUC of ADD(-S) [52]. For the ADD(-S) metric, the mean distance between the model points transformed with the pose estimation and the ground-truth is calculated. With the standard ADD(-S) metric, if the mean distance is less than 10% of the model diameter, the pose estimation is regarded as correct. In some of our experiments, we also test the performances when the threshold is set to 2% or 5% of the model diameter for stricter testing. For symmetric objects, the mean distance is computed based on closest point distances [14]. When evaluating on the YCB-Video dataset, we also compute the AUC (Area Under Curve) of ADD(-S) by varying the distance threshold from 0 cm to 10 cm following [52].

4.2. Ablation Study

We conduct a thorough ablation study on LINEMOD and Occlusion LINEMOD datasets to evaluate the effectiveness of the components in our framework.

Correspondence Field Supervision. We first remove the correspondence loss to verify the influence of correspondence field quality on the pose estimation. The results ‘w/o correspondence loss’ in Table 1 correspond to this ablation study, and the performance degrades significantly. Since

Table 1. (a) Ablation study on LINEMOD dataset. (b) Validation of effectiveness of similarity score on the Occlusion LINEMOD dataset with the ADD(-S) metric. For more detailed comparison, the evaluations with different thresholds of ADD(-S) metric are conducted. The thresholds are set to s 2%, 5% and 10% of the model diameter, denoted as 0.01d, 0.05d, 0.1d respectively.

	w/o correspondence loss			w/o L_{ma}			w/o \hat{C}'_t rect.			w/o 3D context F_{ctx}			w/ 2D context			Full(Ours)		
	0.02d	0.05d	0.1d	0.02d	0.05d	0.1d	0.02d	0.05d	0.1d	0.02d	0.05d	0.1d	0.02d	0.05d	0.1d	0.02d	0.05d	0.1d
Object	1.29	17.60	61.23	8.65	36.03	70.35	4.40	35.76	74.51	14.86	50.48	80.10	12.19	52.22	82.33	18.76	57.14	88.19
Ape	31.60	87.08	99.32	58.14	94.30	99.71	79.56	98.72	100.0	72.26	99.13	100.0	75.26	98.37	99.81	75.17	98.25	100.0
Benchvis	19.37	70.89	94.90	45.13	82.31	95.95	56.72	90.09	97.91	53.63	90.69	98.73	56.90	91.68	97.78	55.39	89.12	98.04
Camera	8.95	77.88	96.83	32.65	86.71	98.76	47.13	94.37	99.31	53.25	95.28	99.80	53.21	95.62	99.72	54.53	94.69	99.31
Can	4.59	28.39	71.64	25.24	62.60	92.81	31.74	75.76	97.98	32.34	74.55	96.71	36.81	79.15	98.55	36.43	74.85	96.41
Cat	40.25	84.04	92.57	49.88	88.50	98.22	59.81	96.43	99.70	60.46	95.34	99.70	60.69	95.54	99.41	62.44	95.44	99.70
Driller	5.62	22.44	69.08	16.66	47.46	79.69	19.18	55.68	87.01	16.71	57.37	85.92	25.19	63.62	88.01	25.82	61.13	89.30
Duck	43.45	89.81	99.65	46.40	87.12	98.12	52.64	83.45	97.65	50.05	81.03	95.59	54.51	86.38	96.36	59.06	93.80	99.53
Eggbox	44.08	93.57	69.83	10.67	52.84	92.29	51.83	93.95	99.87	55.12	94.40	99.52	54.14	95.71	99.87	60.14	95.56	99.71
Glue	6.26	15.89	51.95	31.55	65.55	95.04	32.81	70.22	96.53	24.26	66.51	93.91	20.61	56.03	91.04	35.68	75.26	97.43
Holep.	42.33	96.09	99.08	52.14	95.48	99.69	62.46	97.32	99.59	63.74	97.45	100.00	63.07	98.24	100.0	68.03	98.16	100.0
Lamp	33.87	88.57	98.98	30.18	81.90	99.17	45.95	94.75	99.81	46.35	93.76	99.81	60.71	94.43	99.00	61.32	94.91	99.81
Phone	2.33	18.53	55.59	31.08	72.81	95.36	36.55	82.74	98.13	39.66	82.06	97.26	42.68	83.85	98.39	42.30	83.95	98.39
Average	21.85	60.83	85.27	33.72	73.35	93.47	44.68	82.25	96.00	44.82	82.93	95.93	47.44	83.91	96.17	50.39	85.56	97.37




Figure 5 & Table 2. **Left:** ADD(-S) accuracies w.r.t. different recurrent iterations and rendering cycles on LINEMOD. (a) Results based on the initial poses from PoseCNN [52] (b) Results based on the disturbed PoseCNN poses (with Gaussian noise $\sigma_t = 15\text{cm}$, $\sigma_r = 10^\circ$). **Right:** runtime analysis of individual modules.

Module	Runtime (ms)
Ref. Image Rendering	8.88
3D Context Encoding (run once per sequence)	35.20
3D Feat. Rendering (context&descriptor)	5.39
Image Feat. Encoding	6.39
2D-3D Hybrid Net (2D part)	2.99
CF Estimation	6.21
Pose Optim.	6.23
CF Rectification	1.48

our pose optimization is directly based on the correspondence field estimation, solid supervision on correspondence field estimation is essential to the overall system.

Effectiveness of the Pose Supervision and End-to-end Learning. We further remove the supervision to the pose estimation by setting the weight of the model alignment loss L_{ma} to 0. This is equivalent to adopting a typical non-differentiable LM optimizer because no gradient is back-propagated through the LM layer during training. It can be found that the object pose can still be reasonably estimated (denoted as ‘w/o L_{ma} ’ in Table 1), but with humble performance, especially with stricter evaluation criteria, *i.e.*, by setting a smaller threshold 0.01d or 0.05d. The performance degradation reflects the importance of end-to-end pose learning. The differentiable LM layer enables the pose supervision to affect the feature learning for more robust correspondence field estimation, which is essential to our formulation.

Correspondence Field Rectification. Another key procedure in our recurrent pose refinement is the correspondence field rectification. To validate the effectiveness, we ablate this step and directly use the correspondence estimation \hat{C}_t from the GRU as the initialization for the next iteration (denoted as ‘w/o \hat{C}'_t rect.’ in Table 1). We find that the performance drops significantly compared with our full framework, especially on more strict metrics, *i.e.*, 0.01d and 0.05d. This phenomenon demonstrates that the corrected correspondence field with the rigid-transformation constraints from the optimized pose can facilitate the re-




Figure 6. Robustness comparison with RePOSE by degrading the initial poses (from PVNet [34]) with Gaussian noise on LINEMOD dataset.

finement in the following iterations.

3D Context Encoder. To verify the effectiveness of our 3D context encoder, we test the system without the context encoder (denoted as ‘w/o 3D context F_{ctx} ’) or with a commonly used 2D context encoder (denoted as ‘w/ 2D context’). The performances of these two versions both degrade compared to that with a 3D context encoder. The degradation not only reveals the importance of context information as indicated by previous works [39,44], but also proves that our 3D context encoder is a more effective choice than the 2D counterpart in our task. We reckon that the more robust performance may be attributed to the finer granularity of dense 3D point cloud features (compared with the low-resolution 2D image features). The finer-granularity features could provide more detailed geometric contexts.

Similarity Scores for Occlusion Handling. In Table 1(b), we evaluate the effectiveness of similar scores in occlusion handling on the Occlusion LINEMOD dataset. The version ‘w/ similarity score’ performs better for severely occluded objects. By including similarity scores during pose optimization, flawed correspondence estimations in the occluded unreliable regions are effectively downweighted. Some similarity score map examples are exhibited in Fig. 4 for better understanding.

Recurrent Iterations vs Rendering Cycles. The number of refinement iterations affects the system performance, especially when erroneous initial pose estimations exist. We analyze the performances with different recurrent iterations and rendering cycles in Fig. 5. From Fig. 5(a), it can be found that, by solely increasing the recurrent iterations while rendering the reference object image only once, we have achieved a high accuracy of 96.05% which is com-

Table 3. The comparison of estimation accuracy with competitive direct methods (PoseCNN [52], PVNet [34] and HybridPose [38]) and refinement methods (DPOD [58], DeepIM [23] and RePOSE [20]) on LINEMOD dataset in terms of the ADD(-S) metric.

Method	PoseCNN	PVNet	HybridPose	DeepIM [23]	DPOD [58]	RePOSE [20]	Ours		
Init. pose	-	-	-	PoseCNN	self-designed	PoseCNN	PVNet		
Ape	25.62	43.62	63.1	76.95	87.73	47.4	79.5	88.19	85.62
Benchvise	77.11	99.90	99.9	97.48	98.45	88.5	100.0	100.0	100.0
Camera	47.25	86.86	90.4	93.53	96.07	67.0	99.2	98.04	98.43
Can	69.98	95.47	98.5	96.46	99.71	88.0	99.8	99.31	99.51
Cat	56.09	79.34	89.4	82.14	94.71	80.6	97.9	96.41	96.41
Driller	64.92	96.43	98.5	94.95	98.80	78.5	99.0	99.70	99.50
Duck	41.78	52.58	65.0	77.65	86.29	66.1	80.3	89.30	89.67
Eggbox	98.50	99.15	100.0	97.09	99.91	98.6	100.0	99.53	100.0
Glue	94.98	95.66	98.8	99.42	96.82	95.6	98.3	99.71	97.30
Holep.	52.24	81.92	89.7	52.81	86.87	62.7	96.9	97.43	97.15
Iron	70.17	98.88	100.0	98.26	100.0	80.3	100.0	100.0	100.0
Lamp	70.73	99.33	99.5	97.50	96.84	87.8	99.8	99.81	100.0
Phone	53.07	92.41	94.9	87.72	94.69	74.3	98.9	98.39	98.68
Average	63.26	86.27	91.3	88.61	95.15	78.1	96.1	97.37	97.10

Table 4. Accuracy comparison with the state of the art on OCCLUSION LINEMOD dataset in terms of the ADD(-S) metric.

Object	PoseCNN [52]	PVNet [34]	HybridPose [38]	GDR-Net [51]	DPOD [58]	RePOSE [20]	Ours
Ape	9.60	15.8	20.9	39.3	-	31.1	37.18
Can	45.2	63.3	75.3	79.2	-	80.0	88.07
Cat	0.93	16.7	24.9	23.5	-	25.6	29.15
Driller	41.4	65.7	70.2	71.3	-	73.1	88.14
Duck	19.6	25.2	27.9	44.4	-	43.0	49.17
Eggbox	22.0	50.2	52.4	58.2	-	51.7	66.98
Glue	38.5	49.6	53.8	49.3	-	54.3	63.79
Holep.	22.1	39.7	54.2	58.7	-	53.6	62.76
Average	24.9	40.8	47.5	53.0	47.3	51.6	60.65

parable to RePOSE [20]. If conducting refinement with more recurrent iterations and rendering cycles, steady improvements are reported, which reflect good convergence of our method. To further validate the robustness to erroneous initial poses, we add Gaussian noise to the initial poses. Specifically, we randomly disturb translation components and rotation Euler angles with Gaussian noise. For the rotation, we add angular noise with standard deviation (STD, denoted as σ_r) of 10° in all three axes. For the translational disturbance, we apply noise with a STD of 15 cm along the z axis (the axis perpendicular to the image plane) and STDs of 3 cm in x and y directions ($\frac{1}{5}\times$) considering current methods usually have larger variances on depth estimations. From Fig. 5(b), we find that the necessity of recurrent refinement becomes more noticeable.

Though more rendering cycles bring performance gains as well, the extra costs are significant, since most of the input features need re-encoding. Based on the runtime analysis (Table 2), increasing the recurrent iterations is more economical for better performance as only the CF (correspondence field) estimation, pose optimization and CF rectification modules are activated for a recurrent iteration.

4.3. Comparison with State-of-the-Art Methods

We compare with the cutting-edge methods on LINEMOD, Occlusion LINEMOD, and YCB-Video.

For the LINEMOD dataset, we compare with the recent pose refinement methods RePOSE [20], DPOD [58] and DeepIM [23] as well as some direct estimation baselines [34, 38, 52]. Table 3 contains the comparison results and we achieve a state-of-the-art performance. Interestingly, we achieve slightly better average performance when using PoseCNN [52] as the initial pose generator rather than

Table 5. Comparison with the refinement methods based on single images on the YCB-Video dataset. The performance of our initial pose generator, *i.e.*, PoseCNN, is also included.

Metric	PoseCNN [52]	DeepIM [23]	DPOD [58]	RePOSE [20]	Ours
AUC, ADD(-S)	61.3	81.9	76.3	80.8	83.1
ADD(-S)	21.3	53.6	50.4	60.3	66.4

the PVNet [34], although the pose accuracy of PVNet is much better as exhibited in Table 3. This phenomenon reveals the good tolerance of our system to erroneous initial poses. To test our robustness to even larger initial pose errors, we add random Gaussian pose noises to the initial rotation and translation components separately for accuracy evaluation similar to those in Sec. 4.2. Fig. 6 plots the accuracy variations w.r.t. the disturbance magnitudes. Our method exhibits strong robustness and works reasonably even with extremely noisy initial poses.

We also conduct comparisons on Occlusion LINEMOD. As shown in Table 4, we outperform the cutting-edge method [20] by a significant margin (51.6 \nearrow 60.65), which manifests the system robustness to occlusions. We visualize some of our pose estimates from severely occluded images in the first row of Fig. 4, where the initial poses from PVNet are disturbed with Gaussian noise like before ($\sigma_t = 15$ cm, $\sigma_r = 10^\circ$) to pose more challenges. It is shown that our system is capable of handling large initial pose errors even in highly occluded scenarios.

Our additional evaluation on the YCB-Video dataset uses PoseCNN as the pose initializer, following the settings of RePOSE [20]. We compare with the refinement methods based on monocular color images. Our system still performs well on this large-scale complex dataset. We consistently improve the initial poses provided by PoseCNN [52], and outperform the cutting-edge pose refinement method RePOSE in both metrics, as shown in Table 5.

5. Conclusions and Limitations

We have presented a recurrent framework for 6-DOF object pose refinement. A non-linear least squares problem is formulated for pose optimization based on the estimated correspondence field between the rendered image and observed image. Descriptor-based consistency checking is included to downweight unreliable correspondences for occlusion handling. Our method performs robustly against erroneous pose initializations and severe occlusions, which achieves state-of-the-art performances on public datasets.

One limitation of our method is that the trained model is object-specific similar to many other works [20, 23, 58]. Although for a novel object, only the pose refinement module needs further finetuning, the limited generalization ability to unknown objects is still undeniable. More detailed discussions are in the supplementary material. In the future, we plan to extend our method to handle unknown objects for better generality.

Acknowledgement. This work is supported in part by Centre for Perceptual and Interactive Intelligence Limited, in part by the General Research Fund through the Research Grants Council of Hong Kong under Grants (Nos. 14204021, 14207319, 14203118, 14208619), in part by Research Impact Fund Grant No. R5001-18, in part by CUHK Strategic Fund.

References

- [1] Geometric correspondence fields: Learned differentiable rendering for 3D pose refinement in the wild. 2020. [2](#)
- [2] K Somani Arun, Thomas S Huang, and Steven D Blostein. Least-squares fitting of two 3-D point sets. *IEEE Transactions on pattern analysis and machine intelligence*, (5):698–700, 1987. [1](#)
- [3] Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother. Learning 6D object pose estimation using 3D object coordinates. In *European conference on computer vision*, pages 536–551. Springer, 2014. [6](#)
- [4] Berk Calli, Arjun Singh, Aaron Walsman, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. The YCB object and model set: Towards common benchmarks for manipulation research. In *2015 international conference on advanced robotics (ICAR)*, pages 510–517. IEEE, 2015. [6](#)
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. [2](#)
- [6] Bo Chen, Alvaro Parra, Jiewei Cao, Nan Li, and Tat-Jun Chin. End-to-end learnable geometric vision by backpropagating pnp optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8100–8109, 2020. [2](#)
- [7] Dengsheng Chen, Jun Li, Zheng Wang, and Kai Xu. Learning canonical shape space for category-level 6D object pose and size estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11973–11982, 2020. [2](#)
- [8] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018. [4](#)
- [9] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2758–2766, 2015. [3](#)
- [10] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. [2](#)
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. [2](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [3](#)
- [13] Yisheng He, Wei Sun, Haibin Huang, Jianran Liu, Haoqiang Fan, and Jian Sun. PVN3D: A deep point-wise 3D keypoints voting network for 6DoF pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11632–11641, 2020. [2](#)
- [14] Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary Bradski, Kurt Konolige, and Nassir Navab. Model based training, detection and pose estimation of texture-less 3D objects in heavily cluttered scenes. In *Asian conference on computer vision*, pages 548–562. Springer, 2012. [6](#)
- [15] Tomas Hodan, Daniel Barath, and Jiri Matas. EPOS: Estimating 6D pose of objects with symmetries. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11703–11712, 2020. [2](#)
- [16] Yinlin Hu, Pascal Fua, Wei Wang, and Mathieu Salzmann. Single-stage 6D object pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2930–2939, 2020. [2](#)
- [17] Zhaoyang Huang, Han Zhou, Yijin Li, Bangbang Yang, Yan Xu, Xiaowei Zhou, Hujun Bao, Guofeng Zhang, and Hongsheng Li. VS-Net: Voting with segmentation for visual localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6101–6111, 2021. [2](#)
- [18] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. Lite-FlowNet: A lightweight convolutional neural network for optical flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8981–8989, 2018. [3](#), [5](#)
- [19] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017. [3](#), [5](#)
- [20] Shun Iwase, Xingyu Liu, Rawal Khirodkar, Rio Yokota, and Kris M Kitani. RePOSE: Fast 6d object pose refinement via deep texture rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3303–3312, 2021. [1](#), [2](#), [6](#), [8](#)
- [21] Wadim Kehl, Fabian Manhardt, Federico Tombari, Slobodan Ilic, and Nassir Navab. SSD-6D: Making rgb-based 3D detection and 6D pose estimation great again. In *Proceedings of the IEEE international conference on computer vision*, pages 1521–1529, 2017. [2](#)
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [23] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. DeepIM: Deep iterative matching for 6D pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 683–698, 2018. [1](#), [2](#), [6](#), [8](#)
- [24] Zhigang Li, Gu Wang, and Xiangyang Ji. CDPN: Coordinates-based disentangled pose network for real-time rgb-based 6-DoF object pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7678–7687, 2019. [2](#)
- [25] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. BARF: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5741–5751, 2021. [2](#)
- [26] Pengpeng Liu, Michael Lyu, Irwin King, and Jia Xu. Self-low: Self-supervised learning of optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4571–4580, 2019. [4](#)

- [27] Fabian Manhardt, Wadim Kehl, Nassir Navab, and Federico Tombari. Deep model-based 6D pose refinement in rgb. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 800–815, 2018. 1, 2
- [28] Jiageng Mao, Minzhe Niu, Haoyue Bai, Xiaodan Liang, Hang Xu, and Chunjing Xu. Pyramid R-CNN: Towards better performance and adaptability for 3D object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2723–2732, 2021. 2
- [29] Simon Meister, Junhwa Hur, and Stefan Roth. UnFlow: Unsupervised learning of optical flow with a bidirectional census loss. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. 4
- [30] Milan Meloun and Jiří Militký. *Statistical data analysis: A practical guide*. Woodhead Publishing Limited, 2011. 2
- [31] Jorge J Moré. The levenberg-marquardt algorithm: implementation and theory. In *Numerical analysis*, pages 105–116. Springer, 1978. 2
- [32] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. ORB-SLAM: a versatile and accurate monocular SLAM system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015. 2
- [33] Kiru Park, Timothy Patten, and Markus Vincze. Pix2Pose: Pixel-wise coordinate regression of objects for 6D pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7668–7677, 2019. 2
- [34] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. PVNet: Pixel-wise voting network for 6DoF pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4561–4570, 2019. 2, 3, 6, 7, 8
- [35] Mahdi Rad and Vincent Lepetit. BB8: A scalable, accurate, robust to partial occlusion method for predicting the 3D poses of challenging objects without using depth. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3828–3836, 2017. 2
- [36] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3D deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 3
- [37] Paul-Edouard Sarlin, Ajaykumar Unagar, Mans Larsson, Hugo Germain, Carl Toft, Viktor Larsson, Marc Pollefeys, Vincent Lepetit, Lars Hammarstrand, Fredrik Kahl, et al. Back to the feature: Learning robust camera localization from pixels to pose. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3247–3257, 2021. 2
- [38] Chen Song, Jiaru Song, and Qixing Huang. Hybridpose: 6D object pose estimation under hybrid representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 431–440, 2020. 8
- [39] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018. 2, 3, 5, 7
- [40] Yifan Sun, Changmao Cheng, Yuhan Zhang, Chi Zhang, Liang Zheng, Zhongdao Wang, and Yichen Wei. Circle loss: A unified perspective of pair similarity optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6398–6407, 2020. 5, 6
- [41] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3D orientation learning for 6D object detection from rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 699–715, 2018. 1, 2
- [42] Chengzhou Tang and Ping Tan. BA-Net: Dense bundle adjustment networks. In *International Conference on Learning Representations*, 2018. 2
- [43] Zachary Teed and Jia Deng. DeepV2D: Video to depth with differentiable structure from motion. In *International Conference on Learning Representations*, 2019. 2
- [44] Zachary Teed and Jia Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *16th European Conference on Computer Vision, ECCV 2020*, pages 402–419. Springer Science and Business Media Deutschland GmbH, 2020. 2, 3, 4, 5, 7
- [45] Bugra Tekin, Sudipta N Sinha, and Pascal Fua. Real-time seamless single shot 6D object pose prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 292–301, 2018. 2
- [46] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotequi, François Goulette, and Leonidas J Guibas. KPConv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6411–6420, 2019. 3, 4
- [47] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9627–9636, 2019. 2
- [48] Lukas Von Stumberg, Patrick Wenzel, Qadeer Khan, and Daniel Cremers. GN-Net: The gauss-newton loss for multi-weather relocalization. *IEEE Robotics and Automation Letters*, 5(2):890–897, 2020. 2
- [49] Kentaro Wada, Edgar Sucar, Stephen James, Daniel Lenton, and Andrew J Davison. Morefusion: Multi-object reasoning for 6D pose estimation from volumetric fusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14540–14549, 2020. 2
- [50] Chen Wang, Danfei Xu, Yuke Zhu, Roberto Martín-Martín, Cewu Lu, Li Fei-Fei, and Silvio Savarese. Densefusion: 6D object pose estimation by iterative dense fusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3343–3352, 2019. 2
- [51] Gu Wang, Fabian Manhardt, Federico Tombari, and Xiangyang Ji. GDR-Net: Geometry-guided direct regression network for monocular 6D object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16611–16621, 2021. 2, 8
- [52] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes. In Hadas Kress-Gazit, Siddhartha S. Srinivasa, Tom Howard, and Nikolay Atanasov, editors, *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018*, 2018. 1, 2, 3, 6, 7, 8
- [53] Yan Xu, Junyi Lin, Jianping Shi, Guofeng Zhang, Xiaogang Wang, and Hongsheng Li. Robust self-supervised lidar odometry via representative structure discovery and 3D inherent error modeling. *IEEE Robotics and Automation Letters*, 2022. 2

- [54] Li Yang, Yan Xu, Shaoru Wang, Chunfeng Yuan, Ziqi Zhang, Bing Li, and Weiming Hu. PDNet: Towards better one-stage object detection with prediction decoupling. *arXiv preprint arXiv:2104.13876*, 2021. [2](#)
- [55] Li Yang, Yan Xu, Chunfeng Yuan, Wei Liu, Bing Li, and Weiming Hu. Improving visual grounding with visual-linguistic verification and iterative reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. [2](#)
- [56] Hongwei Yi, Shaoshuai Shi, Mingyu Ding, Jiankai Sun, Kui Xu, Hui Zhou, Zhe Wang, Sheng Li, and Guoping Wang. SegVoxelNet: Exploring semantic context and depth-aware features for 3D vehicle detection from point cloud. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2274–2280. IEEE, 2020. [2](#)
- [57] Georges Younes. The advantages of a joint direct and indirect VSLAM in AR. 2019. [1](#)
- [58] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. DPOD: 6D pose object detector and refiner. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1941–1950, 2019. [1](#), [2](#), [8](#)
- [59] Xinge Zhu, Yuexin Ma, Tai Wang, Yan Xu, Jianping Shi, and Dahua Lin. SSN: Shape signature networks for multi-class object detection from point clouds. In *European Conference on Computer Vision*, pages 581–597. Springer, 2020. [2](#)