

MRC-Net: 6-DoF Pose Estimation with MultiScale Residual Correlation

Yuelong Li *
Amazon Inc.

yuell@amazon.com

Yafei Mao *
Amazon Inc.

yafeimao@amazon.com

Raja Bala
Amazon Inc.

rajabl@amazon.com

Sunil Hadap
Amazon Inc.

hadsunil@lab126.com

Abstract

We propose a single-shot approach to determining 6-DoF pose of an object with available 3D computer-aided design (CAD) model from a single RGB image. Our method, dubbed MRC-Net, comprises two stages. The first performs pose classification and renders the 3D object in the classified pose. The second stage performs regression to predict fine-grained residual pose within class. Connecting the two stages is a novel multi-scale residual correlation (MRC) layer that captures high-and-low level correspondences between the input image and rendering from first stage. MRC-Net employs a Siamese network with shared weights between both stages to learn embeddings for input and rendered images. To mitigate ambiguity when predicting discrete pose class labels on symmetric objects, we use soft probabilistic labels to define pose class in the first stage. We demonstrate state-of-the-art accuracy, outperforming all competing RGB-based methods on four challenging BOP benchmark datasets: T-LESS, LM-O, YCB-V, and ITODD. Our method is non-iterative and requires no complex post-processing. Our code and pretrained models are available at <https://github.com/amzn/mrc-net-6d-pose>.

1. Introduction

Estimating 3D object pose (rotation and translation) relative to the camera from a single image is a fundamental problem in many computer vision applications including robotics, autonomous navigation, and augmented reality. This task is challenging due to the complex shapes of real world objects, and diversity in object appearance due to lighting, surface color, background clutter, object symmetry, and occlusions.

A common solution is to directly regress object poses from images using deep neural networks [28, 53]. Alternatively the problem can be framed as one of classification, predicting a pose in terms of discrete buckets [2, 25]. There have also been attempts to combine the two approaches, predicting a coarse pose class and

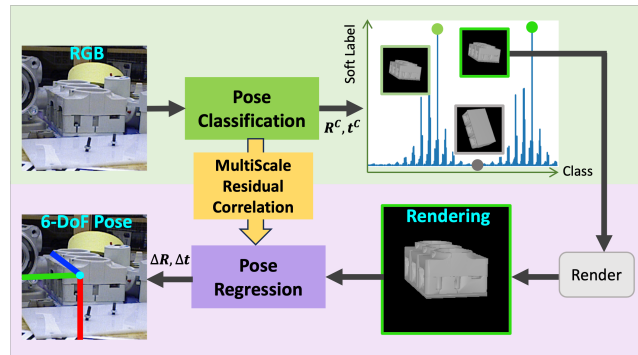


Figure 1. MRC-Net features a *single-shot* sequential Siamese structure of two stages, where the second stage conditions on the first classification stage outcome through multi-scale residual correlation of poses between input and rendered images.

then regressing residual pose within class [37]. While residual regression helps reduce quantization errors from classification to some extent, performance generally falls short of state-of-the-art, especially in challenging scenarios where there is lack of object texture or heavy occlusion. We believe a key reason is that in current approaches, the problem is formulated as multitask learning, where classification and regression tasks are trained *in parallel* with shared top-level features. Such a design does not enable the regression task to receive direct guidance from the classification step.

We approach the problem differently. Since the tasks of classification and residual pose regression are inherently sequential, we hypothesize that it is more effective and natural to also learn them sequentially. We thus propose a two stage deep learning pipeline. In the first stage a classifier predicts pose in terms of a set of pose buckets. In the second stage, a deep regressor predicts residual pose within-bucket using features from the first stage. The classifier and regressor are implemented as Siamese networks with shared weights. Bridging the two stages is a novel multiscale residual correlation (MRC) layer that draws on the well known render-and-compare paradigm to capture correspondences between the input pose and the pose rendered from the first classification

*Equal contribution

stage. The multiscale architecture enables both local and global correspondences to inform pose estimation. To reduce ambiguity in discrete pose classification for objects with symmetry, we define class membership in terms of soft probabilistic labels. The overall network, dubbed MRC-Net is shown in Figure 1. We validate our hypothesis in experiments, showing that the simple concept of sequential pose estimation with the MRC layer produces a major boost in performance and outperforms state-of-the-art techniques on the BOP Challenge [18] datasets without the need for pre-initialization, iterative refinement, and post-processing.

To summarize, the main contributions of this work are:

- MRC-Net, a novel single-shot approach to directly estimate the 6-DoF pose of objects with known 3D models from monocular RGB images. Unlike prior methods, our approach performs classification and regression sequentially, guiding residual pose regression by conditioning on classification outputs. Moreover, we introduce a custom classification design based on soft labels, to mitigate symmetry-induced ambiguities.
- A novel MRC layer that implicitly captures correspondences between input and rendered images at both global and local scales. Since MRC-Net is end-to-end trainable, this encourages the correlation features to be more discriminative, and avoids the need for complicated post-processing procedures.
- State-of-the-art accuracy on a variety of BOP benchmark datasets, advancing average recall by 2.4% on average compared to results reported by competing methods.

2. Related Work

Traditionally, 6-DoF pose estimation has been solved via feature correspondences or template matching [6, 15]. We focus our review on more recent learning-based methods relevant to our approach.

Direct pose estimation aims to determine 6-DoF pose in an end-to-end fashion without resorting to PnP solvers. Due to the continuous nature of object poses, the problem can be naturally formulated as one of regression [29, 53]. As object pose spans a large range, a popular approach is to iteratively refine an initial pose hypothesis through incremental updates [26, 28]. Regression methods might get trapped in poor local minima unless properly initialized. To address this difficulty, some works opt for pose classification [2, 25, 49, 50]. Although more robust and reliable at a coarse scale, classification suffers from loss of accuracy due to quantization of the pose continuum. To tackle this, some methods add a parallel residual regression task [37]. However, the regression branch is only indirectly connected to classification outcomes through a shared backbone. In contrast, our pipeline is sequential, where the residual regressor receives explicit guidance from the

classifier via the MRC layer. To handle objects unseen during training, Megapose [27] extends CosyPose [26] by initializing pose via a coarse classifier, followed by a refiner. However, these two stages are independent networks that are separately trained. In contrast, our method strongly couples the two stages via a Siamese architecture and MRC layer, and is trained end-to-end.

Render-and-compare is a strategy widely employed for refining pose estimates by feeding additional rendered images alongside the input image into the deep network. Pioneering this research direction, [28] and [36] introduced an iterative process that progressively renders an image based on the current pose estimate. CosyPose [26] generalizes this to a multi-view setting. CIR [33], built upon [51], estimates optical flow as a dense 2D-to-2D correspondence between real and rendered images and introduces a differentiable PnP solver to perform coupled updates on both pose and correspondence. In addition to 2D-to-2D correspondence, RNNPose [54] optimize a 3D context encoder using Levenberg-Marquadt optimization [38]. PFA [22] propose a non-iterative method by ensembling flow fields from the exemplars towards the input image. SCFlow [11] restricts indexing correlation volume within the projected target’s 3D shape to alleviate false correspondences. All these methods estimate intermediate dense correspondences and require proper initialization using off-the-shelf networks. Our approach uses render-and-compare, but does not require pose initialization, and directly uses MRC features to predict pose in a single pass.

Correspondence-based methods use deep networks to predict an intermediate sparse or dense 3D-to-2D correspondence map, then predict pose through iterative PnP solvers. Earlier works focused on finding sparse corresponding geometric [41] or semantic [31, 40, 46] keypoints with iterative refinement strategies [3]. As sparse keypoints can be easily occluded, more recent techniques predict dense correspondences coupled with robust matching techniques [13, 20, 29, 44]. Instead of per-pixel correspondences, some methods explore alternative surface representations for correspondence matching, such as recursive binary surface encoding [47], surface fragments [17]. Correspondence based methods involve computationally complex PnP solvers and are hard to train end-to-end, which can lead to suboptimal results.

Other works explore diverse avenues, such as knowledge distillation [10], new loss functions [34], integrating object detection [12], generalizing towards unseen objects [56]. Since depth data naturally aids 6-DoF pose estimation, several works [13, 17, 33, 47] explore this direction. While our method can be extended to incorporate depth, in this work we focus solely on single RGB image inputs to maintain broad scope and applicability.

3. Methodology

Given a single RGB image crop (expanded and padded to square) around the object of interest and the 3D model of the object, MRC-Net estimates 3D object rotation \mathbf{R} and translation \mathbf{t} . In practice, the image crop may be generated by an off-the-shelf object detector such as Mask RCNN [14]. We concatenate the crop with the binary encoding of the detection bounding box, dilated to model detection inaccuracies, as input to the network. The architecture is shown in Figure 2.

The classification and residual regression stages solve three concurrent subtasks, namely, 3D rotation $\mathbf{R} \in SO(3)$, 2D in-plane translation $(t_x, t_y) \in \mathbb{R}^2$, and 1D depth $t_z \in \mathbb{R}$. We employ a Siamese architecture in both stages, incorporating the same asymmetrical ResNet34-UNet (with shared weights) to encode image features and discern the visible object mask. The mask is explicitly supervised during training. To capture local and global features, we use Atrous Spatial Pyramid Pooling (ASPP) [5] modules before both classifier and regressor.

In the first stage, we process the RGB input through a pose classifier that predicts pose labels $\mathbf{R}^c, (t_x^c, t_y^c, t_z^c)$, and renders an image with the predicted pose. Our aim is to achieve granular quantization to minimize pose residuals, thereby reducing the burden on the subsequent regression stage. To quantize rotation $SO(3)$, we uniformly generate K rotation prototypes $\{\mathbf{R}_k\}_{k=1}^K$ using the method described in [55]. This allows us to uniformly partition $SO(3)$ into K buckets through nearest neighbor assignment. We choose $K = 4608$ purely from practical consideration that it is the largest number we can fit into GPU memory. Similarly, for the spatial quantization of t_x, t_y , and t_z , we first use the Scale-invariant Translation Estimation (SITE) approach from [29] to transform translation into a well-defined estimation target denoted as (τ_x, τ_y, τ_z) , and then generate spatial uniform grids in 2D and 1D, respectively.

We feed the rendered image from the first stage and its bounding box into the second stage regressor which estimates pose residuals $\Delta\mathbf{R}, (\Delta\tau_x, \Delta\tau_y, \Delta\tau_z)$ between input and rendered object based on render-and-compare. Thus pose is guided by the classification stage, simplifying the regression task and eliminating the need for iterative refinement. Unlike traditional render-and-compare approaches that concatenate real and rendered images as network inputs [28, 33], we calculate correlations between these images in the MRC layer, detailed in the next section. Finally, the predictions from both stages are combined to determine 6-DoF pose: $\hat{\mathbf{R}} = \Delta\mathbf{R}\mathbf{R}^c$, $\hat{\tau}_i = \tau_i^c + \Delta\tau_i, i \in \{x, y, z\}$.

3.1. MultiScale Residual Correlation

MRC computes correlations between real and rendered image feature volumes at three scales. The correlated

features are subsequently fed into the regression head to produce residual pose. As shown in Figure 2, we progressively aggregate real features, correlation features, and downsampled features from the previous finer scale. This process is designed to incorporate multi-scale correlation and image context into the regressor. At the finest level, we also concatenate and input the object visibility mask, making subsequent layers occlusion-aware. We include a 1×1 convolution before the correlation operation, utilizing shared weights for both real and synthetic branches, to obtain a linear projection of correlation features. When computing correlations, we adhere to the convention outlined in [48], constraining the correlation within a local $P \times P$ image window. Mathematically, given a real feature volume \mathbf{f}_r and a synthetic one \mathbf{f}_s , both with dimensions $d \times H \times W$, the resulting correlation volume \mathbf{c} (of shape $P^2 \times H \times W$) has elements defined by:

$$\mathbf{c}(\mathbf{v}, \mathbf{x}) = \frac{1}{\sqrt{d}} \mathbf{f}_s(\mathbf{x})^T \mathbf{f}_r(\mathbf{x} + \mathbf{v}), \quad \forall \mathbf{v} : \|\mathbf{v}\|_\infty \leq P,$$

where \mathbf{x} indexes each pixel location and \mathbf{v} indexes spatial shift. Several techniques exist for formulating correlations that aim to capture long range and higher order correspondences [8, 51]. Our approach is based on the observation that the correspondences are generally within short range thanks to our class conditional framework.

We contrast our render-and-compare approach with the common approach of computing optical flow to obtain correspondences between input and rendered images [11, 22, 33] (Fig. 3a). This flow field may capture inaccurate or redundant correspondences, requiring either an ensemble of flow fields from multiple exemplars [22] or a recurrent coupled refinement of correspondence and pose [11, 33]. Moreover, most approaches involve non-differentiable operations like RANSAC and PnP [22, 54], making the system not end-to-end trainable. In contrast, our approach (Fig. 3b) directly feeds multi-scale correspondence features to the regressor. This encourages the network to learn discriminative image features capturing correlations, and eliminates the need for an extra iterative/recurrent refinement of the correspondence field. Additionally, we explicitly feed the visibility mask to the correlation module to promote occlusion-robust learning. Our approach is end-to-end trainable.

3.2. Technical Details

Soft labels for classification: Assigning object rotations to discrete pose buckets presents a complex challenge. Specifically, object rotations may reside on decision boundaries of $SO(3)$ buckets, having equal geodesic distances towards multiple prototypes, introducing inherent ambiguity. Additionally, object symmetry can introduce

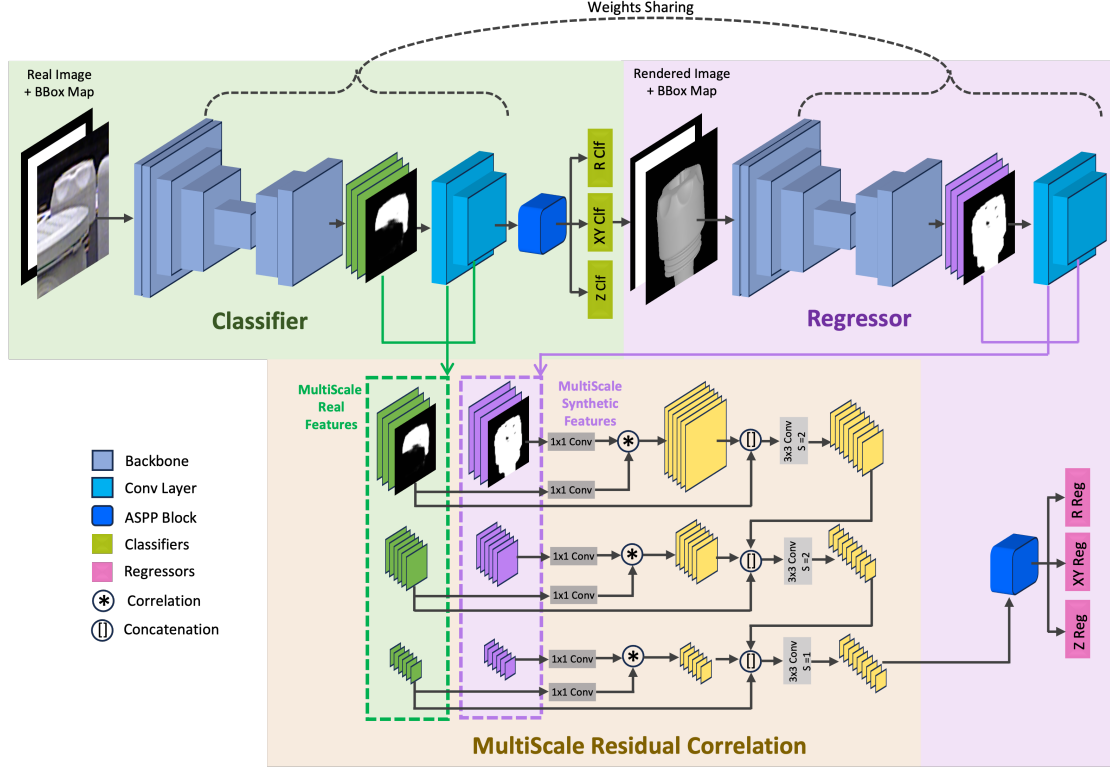


Figure 2. **MRC-Net Architecture.** The classifier and regressor stages employ a Siamese structure with shared weights. Both stages take the object crop and its bounding box map as input, and extract image features to detect the visible object mask, which are concatenated together to estimate object pose. The classifier first predicts pose labels. These predictions, along with the 3D CAD model, are then used to render an image estimate, which serves as input for the second stage. Features from the rendered image are correlated with those from real images in the MRC layer. These correlation features undergo ASPP processing within the rendered branch to regress the pose residuals.

invariance under certain rigid transformations, resulting in multiple valid poses. To address this uncertainty, we employ *soft labels* for the classification task, modeling the *binary* probabilities of whether or not the objects belong to each view bucket based on pose error metrics. A similar concept is explored in [4], where a continuous distribution of object poses is defined based on reprojection error. In contrast, we formulate soft assignments for rotation classes, offering a novel perspective to the problem.

In detail, given the annotated object pose $(\mathbf{R}^*, \mathbf{t}^*)$, we define the rotation labels as:

$$l_k^{\mathbf{R}} = \exp \left\{ -\frac{\rho_{\text{pose-symm}}(\mathbf{R}^*, \mathbf{t}^*; \mathbf{R}_k, \mathbf{t}^*)}{\sigma} \right\}, \quad (1)$$

where $k = 1, \dots, K$, $\rho_{\text{pose-symm}}(\mathbf{R}_1, \mathbf{t}_1; \mathbf{R}_2, \mathbf{t}_2)$ measures the symmetry-aware distance between object poses $(\mathbf{R}_1, \mathbf{t}_1)$ and $(\mathbf{R}_2, \mathbf{t}_2)$ [26], and $\sigma > 0$ is a hyperparameter regulating the concentration of classification labels. Throughout our experiments, we fix $\sigma = 0.03 d_{\text{diam}}$ where d_{diam} is the object diameter defined as the farthest pairwise vertex distances. Essentially, soft labels depict a exponentially-weighted combination of prototype rotations

\mathbf{R}_k centered around the annotated rotation \mathbf{R}^* . The weighting is determined by their symmetry-aware distance.

Similarly, to generate soft labels for translations, we uniformly quantize (τ_x, τ_y) into a 64×64 grid, and τ_z into 1000 bins within its specified range. The translation labels are finally computed using Gaussian functions centered around the ground truth location.

Loss functions: The rotation classification task is trained by minimizing the focal loss $\mathcal{L}_{\text{cls}}^{\mathbf{R}}$ [32]:

$$\mathcal{L}_{\text{cls}}^{\mathbf{R}} = \sum_{k=1}^K -w^+ l_k^{\mathbf{R}} (1 - \hat{p}_k)^2 \log \hat{p}_k - (1 - l_k^{\mathbf{R}}) \hat{p}_k^2 \log (1 - \hat{p}_k),$$

where $\{\hat{p}_k\}_{k=1}^K$ are probabilities predicted by the classifier and $l_k^{\mathbf{R}}$'s are defined in (1). Note that we apply pairwise binary encoding to allow for multiple class affiliations simultaneously and introduce a weighting parameter $w^+ > 0$, which we fix to 100, to address class imbalances.

To solve for translation, both xy and z classification tasks are optimized on multi-class focal loss:

$$\mathcal{L}_{\text{cls}}^i = - \left(1 - \sum_{j=1}^{n_i} l_j^i \hat{p}_j^i \right)^2 \log \left(\sum_{j=1}^{n_i} l_j^i \hat{p}_j^i \right), \quad i \in \{xy, z\},$$

where l_j^{xy} and l_j^z are the translation soft labels for xy and z , \hat{p}_j^{xy} and \hat{p}_j^z are predicted probabilities, $n_{xy} = 64^2$ and $n_z = 1000$. At inference time, we pick the class with highest confidence for all three tasks.

The rotation regression task predicts residual rotation $\Delta \mathbf{R}$ between rendered and real object. This quantity is typically localized with smaller pose angles, and easier to learn via regression. Combining the coarse rotation \mathbf{R}^c from the classifier with $\Delta \mathbf{R}$ yields the fine-grained rotation prediction $\hat{\mathbf{R}} = \Delta \mathbf{R} \mathbf{R}^c$. We adopt the 6D rotation representation suggested by [57] as it generally yields better accuracy and reduces the occurrence of large errors. Similarly, we combine translation estimates from the two stages to arrive at the final object translation: $\hat{\tau}_i = \tau_i^c + \Delta \tau_i$, $i \in \{x, y, z\}$

To train the rotation regression task, disentangled loss [45] is used

$$\mathcal{L}_{\text{reg}}^{\mathbf{R}} = \rho_{\text{pose-sym}}(\Delta \mathbf{R} \mathbf{R}^c, \mathbf{t}^*; \mathbf{R}^*, \mathbf{t}^*),$$

where $(\mathbf{R}^*, \mathbf{t}^*)$ is the annotated object pose. The other two terms $\mathcal{L}_{\text{reg}}^{xy}$ and $\mathcal{L}_{\text{reg}}^z$ are defined similarly. The final loss function is formed via a weighted combination of the individual terms:

$$\begin{aligned} \mathcal{L} = & w_{\text{cls}}^{\mathbf{R}} \mathcal{L}_{\text{cls}}^{\mathbf{R}} + w_{\text{cls}}^{xy} \mathcal{L}_{\text{cls}}^{xy} + w_{\text{cls}}^z \mathcal{L}_{\text{cls}}^z \\ & + w_{\text{reg}}^{\mathbf{R}} \mathcal{L}_{\text{reg}}^{\mathbf{R}} + w_{\text{reg}}^{xy} \mathcal{L}_{\text{reg}}^{xy} + w_{\text{reg}}^z \mathcal{L}_{\text{reg}}^z + w^M \mathcal{L}^M, \end{aligned} \quad (2)$$

where \mathcal{L}^M is the visible mask term defined as the binary cross entropy loss and $w^M > 0$ is its weight.

Perspective correction: As our method operates on a cropped view centered around the object of interest, it may lack global context. Specifically, as highlighted in [30], situations may arise where the network is compelled to predict different poses for image crops with identical appearances, causing confusion in supervision. A common strategy to mitigate this challenge involves converting egocentric rotations into allocentric ones [2, 52]. However, this solution demands accurate prediction of the object center and does not address translation considerations. Instead, we additionally incorporate the global information of the bounding box features into each classifier, similar to [30]. We feed $\left[\frac{b_x - c_x}{f}, \frac{b_y - c_y}{f}, \frac{s_{bb_{ox}}}{f} \right]$, where b_x , b_y , c_x , c_y , and f represent the bounding box center, camera principal point, and focal length, respectively. In our experiments, as a common practice we directly retrieve the camera focal length and principal point from dataset annotations. Different from [30], we obtain the coarse translation estimate (t_x^c, t_y^c, t_z^c) after the classification stage. Subsequently, in the regression render-and-compare stage, we directly feed $\left[\frac{t_x^c}{t_x^c}, \frac{t_y^c}{t_y^c}, \frac{s_{bb_{ox}}}{f} \right]$ to the regression head instead of the bounding box center. This modification ensures that the regressor is informed

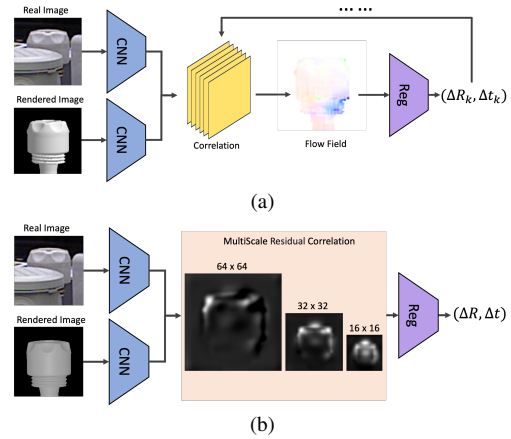


Figure 3. Comparison of (a) conventional approaches [33] and (b) our proposed approach leveraging feature correlations to estimate residual pose. Instead of predicting an intermediate flow field, we directly feed multi-scale feature correlations into the regression head in an end-to-end fashion. These features are discriminative and eliminate the need for post refinement, outlier removal or multiview renderings [11, 22, 33].

about the true projected object center, thereby promoting accurate prediction of the residuals.

4. Experiments

We systematically analyze the novel components of our method and perform comparisons with state of the art.

4.1. Setup

We use AdamW optimizer [35] with cosine annealing learning rate scheduler to train our network. The learning rate is initially set to $2e^{-4}$ and gradually reduced to $5e^{-6}$ with a batch size of 128. We render images using PyTorch3D [42] with constant directional lights along the optical axis, without texturing or antialiasing. We warm-up the learning rate linearly in the first 1K iterations. We choose $w_{\text{cls}}^{\mathbf{R}} = 0.05$, $w_{\text{cls}}^{xy} = 2.0$, $w_{\text{cls}}^z = 2.0$, $w_{\text{reg}}^{\mathbf{R}} = 1.0$, $w_{\text{reg}}^{xy} = 1.0$, $w_{\text{reg}}^z = 0.2$ and $w^M = 10.0$ to balance individual loss terms in (2). These parameters are fixed for all experiments.

Following the BOP Challenge 2022 standard, we employ Mask RCNN detections from [26] to extract image crops. We evaluate our work on four challenging and widely cited BOP benchmark datasets: T-LESS [16], ITODD [9], YCB-V [53], and LM-O [1]. These cover a large variety of difficult cases such as object symmetries, occlusions, absence of textures, etc. A detailed description of these datasets is available on the BOP Challenge website ¹. For each dataset, our model is pre-trained on synthetic

¹<https://bop.felk.cvut.cz/challenges/>

images generated with physically-based rendering (PBR), with a batch size of 128 for a total of 75 epochs on 8 NVIDIA V100 GPUs. We then fine-tune our model using a mixture of real and synthetic images for T-LESS and YCB-V datasets. We train one single model per dataset. The official BOP Challenge website¹ does not provide real images in the training set for ITODD and LM-O; hence, we include results trained exclusively on the synthetic images for these two datasets.

We use the same basic backbone (ResNet34) as many existing works [2, 47, 52] rather than opting for heavier versions for performance. We apply domain randomization techniques similar to [52] to reduce overfitting caused by the domain gap between real and synthetic data. At inference time, we employ the same test-time augmentation (TTA) technique as [2, 13]: we rotate the input image crop by 90°, 180°, 270° and 360°, run inference on these four copies independently, and choose the one with the highest classification score as the final pose prediction.

We adhere to the latest BOP evaluation protocol [18]. Three error metrics are calculated on each object per image: Visible Surface Discrepancy (VSD), Maximum Symmetry-Aware Surface Distance (MSSD), and Maximum Symmetry-Aware Projection Distance (MSPD). An average recall is then computed for each metric by aggregating recall rates at different pre-determined thresholds, to obtain per-metric recalls (AR_{VSD} , AR_{MSSD} , and AR_{MSPD}). Finally, these are averaged to obtain an overall Average Recall (AR). Detailed definitions of these metrics can be found in [18]. Some approaches only report results in ADD(-S), AUC of ADD-S, and AUC of ADD(-S). We show comparisons of these methods on the YCB-V dataset in these metrics, with detailed definitions available in [53].

4.2. Ablation Studies

We explore the significance of each novel component of our design on the challenging T-LESS dataset. The investigation is based solely on using synthetic PBR images for training. Additional ablation studies can be found in the supplementary material.

Hard label vs soft labels: We hypothesize that soft labels enable a more nuanced representation, especially in scenarios where a unique class label is hard to identify due to object symmetries or equiprobable pose hypotheses. To verify its efficacy, we compare training with soft and hard labels, where for the latter, we pick a unique class index with the highest soft label value. Results are summarized in Table 1a. While hard labels are able to maintain reasonable classification performance, soft labels achieve an increase of 2.3% in average AR, clearly validating its effectiveness.

Classification and regression in parallel vs sequential: We start by designing a classification-only baseline model

that comprises only the first stage of MRC-Net. Then, we extend it into a parallel (multitask) architecture by adding an additional regression head atop the pooled feature layer. Next, we implement a simple sequential pipeline that concatenates real and rendered features in 64×64 resolution from stage 1 and passes these to the stage 2 regressor. No MRC layer is included. AR metrics of these models are listed in the first three rows of Table 1b. Clearly, the parallel method leads to significant performance drop (11.3% in average AR), proving the benefit of sequential class-conditioned regression.

An interesting observation is that simple addition of a parallel regression head (second row in Table 1b) offers little or no improvement over the classification-only baseline (first row). As the regression head is unaware of the classification outcome, there is little opportunity for it to correct for classification errors.

MRC: To quantify the benefits of the feature correlation block, we conduct experiments (Table 1b) on single-scale feature correlation, i.e., we only correlate the top-level features of 16×16 resolution. This yields 0.3% additional improvement of AR compared to simple concatenation, implying that correlation features are more discriminative to learn the residual pose. Extending to multiscale correlation further increases AR by 0.6%, confirming that correlation at multiple scales provides complementary information.

Perspective correction: In Table 1c, we examine the impact of perspective correction (PerspCrct) and TTA compared to the complete model. While not being a predominant factor, employing TTA still yields a plausible performance boost of 0.4% AR which demonstrates its effectiveness. Similarly, the inclusion of perspective correction noticeably enhances performance by 0.5%.

4.3. Comparison with State-of-the Art

Quantitative comparisons. We benchmark our method against state-of-the-art techniques spanning a variety of recent approaches: EPOS [17], CDPNv2 [29], DPODv2 [44], PVNet [40], CosyPose [26], SurfEmb [13], SC6D [2], SCFlow [11], CIR [33], PFA [22], SO-Pose [7], NCF [23], CRT-6D [3], GDR-Net [52], ZebraPose [47], DProST [39], RePose [24], SegDriven [19], SingleStage [20], and CheckerPose [31]. We report metrics from the original references.

For all four datasets, we present the AR results using PBR training images. In addition, since real training images are available for T-LESS and YCB-V, we include AR results after fine-tuning on these real images, adhering to the standard BOP protocol. Results are summarized in Table 2. From Table 2a, it can be seen that our method achieves state-of-the-art performance across all datasets when trained on pure synthetic data. We outperform other models on ITODD and YCB-V datasets by a

Method	AR _{VSD} ↑	AR _{MSSD} ↑	AR _{MSPD} ↑	AR↑
Hard Label	67.8	72.2	84.6	74.8
Soft Label	70.6	74.7	86.0	77.1

(a) **Comparison of hard and soft labels.** The table illustrates performance when trained with and without soft labels.

Type	Method	AR _{VSD} ↑	AR _{MSSD} ↑	AR _{MSPD} ↑	AR↑
-	CifOnly	55.6	62.8	77.4	65.3
P	MultiTask	55.3	62.7	76.6	64.9
S	FeatConcat	69.4	73.5	85.7	76.2
S	SSCorr	69.8	73.9	85.8	76.5
S	MRC-Net	70.6	74.7	86.0	77.1

(b) **Comparison of parallel and sequential designs.** **CifOnly** denotes a classification-only model. **MultiTask** is a parallel baseline architecture. **FeatConcat** sends concatenated real and rendered image features to the regressor. **SSCorr** use single-scale feature correlation, while **MRC-Net** uses multi-scale feature correlation. P denotes parallel classification and regression, while S denotes sequential.

Method	AR _{VSD} ↑	AR _{MSSD} ↑	AR _{MSPD} ↑	AR↑
w/o PerspCrrct	69.9	74.1	85.9	76.6
w/o TTA	70.4	74.1	85.7	76.7
Full model	70.6	74.7	86.0	77.1

(c) **Impact of perspective correction (PerspCrrct) and test-time augmentation (TTA).** Dropping perspective correction results in noticeable degradation in performance, while omitting TTA has a more modest impact.

Table 1. Ablation studies of our method on T-LESS dataset [16].

margin of 0.6%, and 0.8% respectively. Particularly on T-LESS, we significantly outperform others by 3%. Surprisingly, our performance on pure synthetic training is even comparable to top-performing methods trained with real data in Table 2b. This suggests that our method has strong potential in applications with limited real data, for example underwater [43] and aerospace [21]. On the other hand, when comparing fine-tuned results over YCB-V dataset, our performance lags behind the top performers by up to 0.9%. Indeed, this dataset carries rich textures on the objects’ surface, facilitating dense correspondence finding and iterative refinement among top-performing methods [11, 26, 33]. Moreover, as discussed in [13], there are inaccurate CAD models and noisy ground truth annotations within this dataset.

Table 3 shows non-BOP metrics (ADD(-S), AUC of ADD-S, and AUC of ADD(-S)) on the YCB-V dataset for methods [7, 11, 19, 20, 24, 26, 31, 39, 47, 52] reporting these metrics. These methods commonly rely on detections from [26] or [29] for their evaluations. Therefore, we

Method	T-LESS	ITODD	YCB-V	LM-O	Avg.
EPOS [17]	46.7	18.6	49.9	54.7	42.5
CDPNv2 [29]	40.7	10.2	39.0	62.4	38.1
DPODv2 [44]	63.6	-	-	58.4	-
PVNet [40]	-	-	-	57.5	-
CosyPose [26]	64.0	21.6	57.4	63.3	51.6
SurfEmb [13]	<u>74.1</u>	<u>38.7</u>	65.3	65.6	<u>60.9</u>
SC6D [2]	73.9	30.3	61.0	-	-
SCFlow [11]	-	-	65.1	<u>68.2</u>	-
PFA [22]	-	-	61.5	67.4	-
CIR [33]	-	-	-	65.5	-
SO-Pose [7]	-	-	-	61.3	-
NCF [23]	-	-	<u>67.3</u>	63.2	-
CRT-6D [3]	-	-	-	66.0	-
MRC-Net	77.1	39.3	68.1	68.5	63.3

(a)

Method	T-LESS	YCB-V	Avg.
CDPNv2 [29]	47.8	53.2	50.5
CosyPose [26]	72.8	82.1	77.4
SurfEmb [13]	77.0	71.8	74.4
SC6D [2]	<u>78.0</u>	78.8	<u>78.4</u>
SCFlow [11]	-	82.6	-
CIR [33]	71.5	<u>82.4</u>	77.0
SO-Pose [7]	-	71.5	-
NCF [23]	-	77.5	-
CRT-6D [3]	-	75.2	-
MRC-Net	79.8	81.7	80.8

(b)

Table 2. **Comparison with state-of-the-art RGB methods on the BOP benchmarks.** We report Average Recall in % on (a) **T-LESS, ITODD, YCB-V, and LM-O** datasets trained with purely **synthetic** images, and (b) **T-LESS and YCB-V** datasets also trained with **real** images. We highlight the best results in **bold** and underline the second best results. “-” denotes results missing from the original paper.

present results for both cases. More details are available in the supp. materials. Notably, MRC-Net offers a clear improvement of 2.2%, 5.4%, and 7.9% in the three metrics.

Qualitative comparison. Figure 4 depicts the qualitative performance of our model on T-LESS test images, compared with three representative state-of-the-art models: (a) PFA initialized with CosyPose (a 2D-to-2D correspondence model) [22], (b) SC6D (a parallel classification and regression model) [2], and (c) ZebraPose (a 2D-to-3D indirect method) [47]. For each object in the scene, we render the corresponding CAD model using the predicted pose to visualize its accuracy. Our method is capable of accurately predicting object poses even under




Figure 4. Qualitative comparison of results on T-LESS: (a) Original RGB image, (b) MRC-Net, (c) CosyPose initialized PFA [22], (d) SC6D [2], and (e) ZebraPose [47]. The object’s 3D model is projected with estimated 6D pose and overlaid on original images with distinct colors. Red boxes denote cases where pose predictions are distinctly different across the methods. MRC-Net outperforms the state-of-art models particularly under heavy occlusion. (Best viewed when zoomed in.)

Method	ADD(-S) ↑	AUC ADD-S ↑	AUC ADD(-S) ↑
SegDriven [19]	39.0	-	-
SingleStage [20]	53.9	-	-
CosyPose [26]	-	89.8	84.5
RePose [24]	62.1	88.5	82.0
GDR-Net [52]	60.1	91.6	84.4
SO-Pose [7]	56.8	90.9	83.9
ZebraPose [47]*	80.5	90.1	85.3
SCFlow [11]	70.5	-	-
DProST [39]	65.1	-	77.4
CheckerPose [31]*	<u>81.4</u>	91.3	86.4
MRC-Net	81.2	<u>95.0</u>	<u>92.3</u>
MRC-Net*	83.6	97.0	94.3

Table 3. **Comparison on the YCB-V Dataset.** We report the ADD(-S), AUC of ADD-S, and AUC of ADD(-S) metrics in %. We highlight the best results in **bold** and underline the second best results. “-” denotes results missing from the original paper and * denotes results obtained using the FCOS detector [29].

heavy occlusions, and is able to discern different objects or multiple instances of the same object even when they are nearby. In contrast, other methods predict inaccurate poses under these challenges. More visual results and failure cases are included in the supplementary material.

Runtime analysis. Table 4 compares run time with recent methods for single object pose estimation from a 256×256 image crop. As expected, our method is significantly faster than iterative refinement methods, even without accounting for their initialization step. MRC-Net is slower than SC6D and GDR-Net due to its

Methods	[33]*	[22]*	[13]	[52]	[2]	MRC-Net
Time (ms)	2542	88	1121	26	25	61

Table 4. **Runtime comparisons.** We run CIR [33], PFA [22], SurfEmb [13], GDR-Net [52], SC6D [2] and MRC-Net (Ours) on the same AWS EC2 P3.2Xlarge instance, and report the average runtime in milliseconds to infer one object pose. For methods with *, the time required for pose initialization, typically involving the execution of an additional off-the-shelf network, is not included.

on-the-fly rendering and two stage inference; however it significantly outperforms these two methods on accuracy, and can be considered viable for real-time operation at 16 frames-per-second.

5. Conclusions and Limitations

We introduced a novel architecture for single-shot 6-DoF pose estimation from single view RGB images, wherein pose classification acts as a conditioning input to residual pose regression. A critical part of the conditioning is a multi-scale correlation layer that captures and transfers residual image features from classification to regression. Extensive experiments on four challenging datasets demonstrate the superior performance of our method. As with other techniques [13, 28], our method suffers from degraded performance when the CAD models are inaccurate as it relies on the render-and-compare technique. Potential solutions involve enforcing stronger augmentations or developing model-agnostic techniques. Although our focus in this work is exclusively on RGB inputs, extending the framework to incorporate depth maps is an interesting future research direction to pursue.

References

- [1] 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. [5](#)
- [2] Dingding Cai, Janne Heikkilä, and Esa Rahtu. SC6D: Symmetry-agnostic and correspondence-free 6d object pose estimation. In *International Conference on 3D Vision*, pages 536–546. IEEE, 2022. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [3] Pedro Castro and Tae-Kyun Kim. CRT-6D: Fast 6d object pose estimation with cascaded refinement transformers. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, pages 5746–5755, 2023. [2](#), [6](#), [7](#)
- [4] Hansheng Chen, Pichao Wang, Fan Wang, Wei Tian, Lu Xiong, and Hao Li. EPro-PnP: Generalized end-to-end probabilistic perspective-n-points for monocular object pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2781–2790, 2022. [4](#)
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2017. [3](#)
- [6] Alvaro Collet and Siddhartha S Srinivasa. Efficient multi-view object recognition and full pose estimation. In *IEEE International Conference on Robotics and Automation*, pages 2050–2055. IEEE, 2010. [2](#)
- [7] Yan Di, Fabian Manhardt, Gu Wang, Xiangyang Ji, Nassir Navab, and Federico Tombari. SO-Pose: Exploiting self-occlusion for direct 6d pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 12396–12405, 2021. [6](#), [7](#), [8](#)
- [8] 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](#)
- [9] Bertram Drost, Markus Ulrich, Paul Bergmann, Philipp Hartinger, and Carsten Steger. Introducing MVTEC ITODD-a dataset for 3d object recognition in industry. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 2200–2208, 2017. [5](#)
- [10] Shuxuan Guo, Yinlin Hu, Jose M. Alvarez, and Mathieu Salzmann. Knowledge distillation for 6d pose estimation by aligning distributions of local predictions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 18633–18642, 2023. [2](#)
- [11] Yang Hai, Rui Song, Jiaojiao Li, and Yinlin Hu. Shape-constraint recurrent flow for 6d object pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4831–4840, 2023. [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [12] Yang Hai, Rui Song, Jiaojiao Li, Mathieu Salzmann, and Yinlin Hu. Rigidity-aware detection for 6d object pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8927–8936, 2023. [2](#)
- [13] Rasmus Laurvig Haugaard and Anders Glent Buch. Surfemb: Dense and continuous correspondence distributions for object pose estimation with learnt surface embeddings. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6749–6758, 2022. [2](#), [6](#), [7](#), [8](#)
- [14] 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. [3](#)
- [15] Stefan Hinterstoisser, Stefan Holzer, Cedric Cagniart, Slobodan Ilic, Kurt Konolige, Nassir Navab, and Vincent Lepetit. Multimodal templates for real-time detection of texture-less objects in heavily cluttered scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 858–865. IEEE, 2011. [2](#)
- [16] Tomáš Hodan, Pavel Haluza, Štěpán Obdržálek, Jiri Matas, Manolis Lourakis, and Xenophon Zabulis. T-LESS: An RGB-D dataset for 6d pose estimation of texture-less objects. In *IEEE Winter Conference on Applications of Computer Vision*, pages 880–888. IEEE, 2017. [5](#), [7](#)
- [17] Tomas Hodan, Daniel Barath, and Jiri Matas. EPOS: Estimating 6d pose of objects with symmetries. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11703–11712, 2020. [2](#), [6](#), [7](#)
- [18] Tomáš Hodaň, Martin Sundermeyer, Bertram Drost, Yann Labbé, Eric Brachmann, Frank Michel, Carsten Rother, and Jiří Matas. BOP challenge 2020 on 6d object localization. In *European Conference on Computer Vision*, pages 577–594. Springer, 2020. [2](#), [6](#)
- [19] Yinlin Hu, Joachim Hugonot, Pascal Fua, and Mathieu Salzmann. Segmentation-driven 6d object pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3385–3394, 2019. [6](#), [7](#), [8](#)
- [20] Yinlin Hu, Pascal Fua, Wei Wang, and Mathieu Salzmann. Single-stage 6d object pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2930–2939, 2020. [2](#), [6](#), [7](#), [8](#)
- [21] Yinlin Hu, Sebastien Speierer, Wenzel Jakob, Pascal Fua, and Mathieu Salzmann. Wide-depth-range 6d object pose estimation in space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 15870–15879, 2021. [7](#)
- [22] Yinlin Hu, Pascal Fua, and Mathieu Salzmann. Perspective flow aggregation for data-limited 6d object pose estimation. In *European Conference on Computer Vision*, pages 89–106. Springer, 2022. [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [23] Lin Huang, Tomas Hodan, Lingni Ma, Linguang Zhang, Luan Tran, Christopher Twigg, Po-Chen Wu, Junsong Yuan, Cem Keskin, and Robert Wang. Neural correspondence field for object pose estimation. In *European Conference on Computer Vision*, pages 585–603. Springer, 2022. [6](#), [7](#)

- [24] 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 International Conference on Computer Vision*, pages 3303–3312, 2021. 6, 7, 8
- [25] 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. 1, 2
- [26] Yann Labbé, Justin Carpentier, Mathieu Aubry, and Josef Sivic. CosyPose: Consistent multi-view multi-object 6d pose estimation. In *European Conference on Computer Vision*, pages 574–591. Springer, 2020. 2, 4, 5, 6, 7, 8
- [27] Yann Labbé, Lucas Manuelli, Arsalan Mousavian, Stephen Tyree, Stan Birchfield, Jonathan Tremblay, Justin Carpentier, Mathieu Aubry, Dieter Fox, and Josef Sivic. MegaPose: 6d pose estimation of novel objects via render & compare. In *Proceedings of the 6th Conference on Robot Learning*, 2022. 2
- [28] 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*, pages 683–698, 2018. 1, 2, 3, 8
- [29] Zhigang Li, Gu Wang, and Xiangyang Ji. CDPN: coordinates-based disentangled pose network for real-time rgb-based 6-DoF object pose estimation. In *IEEE International Conference on Computer Vision*, pages 7677–7686, 2019. 2, 3, 6, 7, 8
- [30] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. CLIFF: Carrying location information in full frames into human pose and shape estimation. In *European Conference on Computer Vision*, pages 590–606. Springer, 2022. 5
- [31] Ruyi Lian and Tae-Kyun Kim. CheckerPose: Progressive dense keypoint localization for object pose estimation with graph neural network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 14022–14033, 2023. 2, 6, 7, 8
- [32] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2980–2988, 2017. 4
- [33] Lahav Lipson, Zachary Teed, Ankit Goyal, and Jia Deng. Coupled iterative refinement for 6d multi-object pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6728–6737, 2022. 2, 3, 5, 6, 7, 8
- [34] Fulin Liu, Yinlin Hu, and Mathieu Salzmann. Linear-covariance loss for end-to-end learning of 6d pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 14107–14117, 2023. 2
- [35] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 5
- [36] 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*, pages 800–815, 2018. 2
- [37] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Kosecka. 3D bounding box estimation using deep learning and geometry. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7074–7082, 2017. 1, 2
- [38] Jorge Nocedal and Stephen J Wright. *Numerical optimization*. Springer, 1999. 2
- [39] Jaewoo Park and Nam Ik Cho. DProST: Dynamic projective spatial transformer network for 6d pose estimation. In *European Conference on Computer Vision*, pages 363–379. Springer, 2022. 6, 7, 8
- [40] 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 Conference on Computer Vision and Pattern Recognition*, pages 4561–4570, 2019. 2, 6, 7
- [41] 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
- [42] 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. 5
- [43] Petter Risholm, Peter Ørnulf Ivarsen, Karl Henrik Haugholt, and Ahmed Mohammed. Underwater marker-based pose-estimation with associated uncertainty. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3713–3721, 2021. 7
- [44] Ivan Shugurov, Sergey Zakharov, and Slobodan Ilic. DPODv2: Dense correspondence-based 6 dof pose estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2, 6, 7
- [45] Andrea Simonelli, Samuel Rota Buló, Lorenzo Porzi, Manuel López-Antequera, and Peter Kontschieder. Disentangling monocular 3d object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1991–1999, 2019. 5
- [46] Chen Song, Jiaru Song, and Qixing Huang. HybridPose: 6d object pose estimation under hybrid representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 431–440, 2020. 2
- [47] Yongzhi Su, Mahdi Saleh, Torben Fetzner, Jason Rambach, Nassir Navab, Benjamin Busam, Didier Stricker, and Federico Tombari. ZebraPose: Coarse to fine surface encoding for 6dof object pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6738–6748, 2022. 2, 6, 7, 8
- [48] 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. 3
- [49] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit

- 3d orientation learning for 6d object detection from rgb images. In *European Conference on Computer Vision*, pages 699–715, 2018. [2](#)
- [50] Martin Sundermeyer, Maximilian Durner, En Yen Puang, Zoltan-Csaba Marton, Narunas Vaskevicius, Kai O Arras, and Rudolph Triebel. Multi-path learning for object pose estimation across domains. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 13916–13925, 2020. [2](#)
- [51] Zachary Teed and Jia Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *European Conference on Computer Vision*, pages 402–419. Springer, 2020. [2](#), [3](#)
- [52] 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 Conference on Computer Vision and Pattern Recognition*, pages 16611–16621, 2021. [5](#), [6](#), [7](#), [8](#)
- [53] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. PoseCNN: A convolutional neural network for 6d object pose estimation in cluttered scenes. *Robotics: Science and Systems*, 2018. [1](#), [2](#), [5](#), [6](#)
- [54] Yan Xu, Kwan-Yee Lin, Guofeng Zhang, Xiaogang Wang, and Hongsheng Li. RNNPose: Recurrent 6-dof object pose refinement with robust correspondence field estimation and pose optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 14880–14890, 2022. [2](#), [3](#)
- [55] Anna Yershova, Swati Jain, Steven M Lavalley, and Julie C Mitchell. Generating uniform incremental grids on $SO(3)$ using the Hopf fibration. *The International Journal of Robotics Research*, 29(7):801–812, 2010. [3](#)
- [56] Heng Zhao, Shenxing Wei, Dahu Shi, Wenming Tan, Zheyang Li, Ye Ren, Xing Wei, Yi Yang, and Shiliang Pu. Learning symmetry-aware geometry correspondences for 6d object pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 14045–14054, 2023. [2](#)
- [57] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2019. [5](#)