

Sim2Real Object-Centric Keypoint Detection and Description

Chengliang Zhong^{*1,2}, Chao Yang^{*2}, Fuchun Sun², Jinshan Qi⁴,
Xiaodong Mu¹, Huaping Liu², Wenbing Huang³

¹Xi'an Research Institute of High-Tech, Xi'an 710025, China

²Beijing National Research Center for Information Science and Technology (BNRist),
State Key Lab on Intelligent Technology and Systems,

Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China

³Institute for AI Industry Research, Tsinghua University

⁴Shandong University of Science and Technology, Qingdao 266590, China
zhongcl19@mails.tsinghua.edu.cn, fcsun@tsinghua.edu.cn

Abstract

Keypoint detection and description play a central role in computer vision. Most existing methods are in the form of *scene-level* prediction, without returning the object classes of different keypoints. In this paper, we propose the *object-centric* formulation, which, beyond the conventional setting, requires further identifying which object each interest point belongs to. With such fine-grained information, our framework enables more downstream potentials, such as object-level matching and pose estimation in a clustered environment. To get around the difficulty of label collection in the real world, we develop a sim2real contrastive learning mechanism that can generalize the model trained in simulation to real-world applications. The novelties of our training method are three-fold: (i) we integrate the uncertainty into the learning framework to improve feature description of hard cases, e.g., less-textured or symmetric patches; (ii) we decouple the object descriptor into two output branches, intra-object salience and inter-object distinctness, resulting in a better pixel-wise description; (iii) we enforce cross-view semantic consistency for enhanced robustness in representation learning. Comprehensive experiments on image matching and 6D pose estimation verify the encouraging generalization ability of our method from simulation to reality. Particularly for 6D pose estimation, our method significantly outperforms typical unsupervised/sim2real methods, achieving a closer gap with the fully supervised counterpart.

Introduction

Extracting and describing points of interest (*keypoints*) from images are fundamental problems for many geometric computer vision tasks such as image matching (Lowe 2004), camera calibration (Strecha et al. 2008), and visual localization (Piasco et al. 2018). Particularly for image matching, it requires searching the same and usually sparse keypoints for a pair of images that record the same scene but under a different viewpoint.

^{*}These authors contributed equally. Corresponding author: Fuchun Sun.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

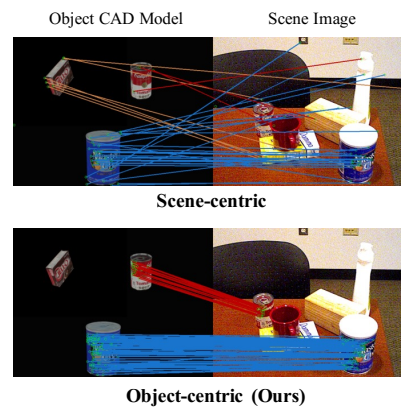


Figure 1: Keypoint matching processing for scene-centric method, R2D2 (Revaud et al. 2019) and our object-centric method. Our method accurately matches the keypoints on different objects, while R2D2 predicts some unwanted points located in the background.

A variety of works have been done towards keypoint detection and description, ranging from traditional hand-crafted methods (Lowe 2004; Bay, Tuytelaars, and Van Gool 2006) to current data-driven approaches (DeTone, Malisiewicz, and Rabinovich 2018; Revaud et al. 2019; Tyszkiewicz, Fua, and Trulls 2020). Despite the fruitful progress, most existing methods are initially targeted on image-level/scene-centric tasks, making them less sophisticated for other more fine-grained problems, e.g., the object-level matching or pose estimation. Object-level tasks are crucial in many applications. A typical example is in robotic grasp manipulation, where for planning a better grasp posture, the robot requires to compare the keypoints between the scene image and the CAD model followed by estimating the pose of the object according to the keypoint correspondence (Sadran, Wurm, and Burschka 2013). If we apply the previous methods (such as R2D2 (Revaud et al. 2019)) straightly, the detected keypoints from the scene image usually contain not only the desired points on the target object, but also those unwanted points located in the background

that share a similar local texture with the object CAD model, as illustrated in Figure 1 (top row). Such failure is natural as R2D2 never tells which keypoints are on the same object, and it solely captures the local similarity therein.

In this paper, we propose a novel conception: *object-centric* keypoint detection and description, in contrast to the conventional *scene-centric* setting. Beyond keypoint detection and description, the proposed object-centric formulation further teaches the algorithm to identify which object each keypoint belongs to. Via this extra supervision, our formulation emphasizes the similarity between keypoints in terms of both local receptive field, and more importantly, objectness. Figure 1 depicts that the object-centric method (bottom row) accurately predicts the object correspondence (different colors) and matches the keypoints on different objects between the scene image and the CAD model, thanks to the object-wise discrimination.

Obtaining object annotations, of course, is resource-consuming in practice. Yet, we find it well addressable if leveraging the sim2real training and domain randomization mechanism (Wen et al. 2020). By this strategy, we can easily obtain rich supervision and manipulate the training samples arbitrarily to serve our goal. Specifically in our scenario, we can access the transformation projection between different views of the same scene image as well as the object label of each pixel. Based on the richly annotated simulation data, we develop a contrastive learning framework to jointly learn keypoint detection and description. To be specific, it contains three main components: (i) A novel uncertainty term is integrated into object keypoints detection, which is expected to handle challenging objects with similar local textures or geometries. (ii) As for the keypoint description, it is divided into two output branches, one to obtain intra-object salience for the keypoints on the same object and the other to enforce inter-object distinctness across the keypoints on different objects. (iii) For each target object in the scene image, a contrastive sample containing only the object of the same view but with a clean background is also adopted to derive better cross-view semantic consistency. Once the model is trained in simulation, it can be applied to real applications.

We summarize our contributions below.

- To the best of our knowledge, we are the first to raise the notion of object-centric keypoint detection and description, which better suits the object-level tasks.
- To address the proposed task, we develop a novel sim2real training method, which enforces uncertainty, intra-object salience/inter-object distinctness, and semantic consistency.
- Comprehensive experiments on image matching and 6D pose estimation verify the encouraging generalization ability of our method from simulation to reality.

Related Work

Here, we focus on local keypoint detection and description from a 2D image.

Scene-centric method. The hand-crafted methods often employ corners (Smith and Brady 1997; Trajković and Hedley 1998) or blobs (Bay, Tuytelaars, and Van Gool 2006) as

keypoints whose associated descriptions are based on histograms of local gradients, including the famous SIFT descriptor (Lowe 2004). In part, this is due to the increased complexity of scene semantics, which exacerbates the reliance of keypoint detection and description to modern deep learning approaches. Such data-driven methods can be categorized into learned detectors (Verdie et al. 2015; Barroso-Laguna et al. 2019), learned descriptors (Han et al. 2015; Balntas et al. 2016), and the combination of them (Yi et al. 2016; Dusmanu et al. 2019; Revaud et al. 2019).

Object-centric keypoint detection. According to the different levels of supervision, object-centric keypoint detection can be divided into fully supervised (Rad and Lepetit 2017; Peng et al. 2019), semi-supervised (Vecerik et al. 2020), and self-supervised learning methods (Zhao et al. 2020; Kulkarni et al. 2019). To distinguish the objectness of keypoints, most researchers utilize pre-trained object detectors to focus on small patches of different objects and find keypoints for each. Also, the number of keypoints of each object would be fixed. Although some works (Vecerik et al. 2020) did not involve the object detector, there is only one object in each image. These methods are difficult to generalize to new objects because of the specialized object detectors and fixed number of keypoints. In the keypoints detection part, the unsupervised learning method (Kulkarni et al. 2019) will predict the keypoints of multiple objects at once. However, their methods are limited to simple scenarios.

Object-centric keypoints descriptors. Dense-Object-Net (DON) (Florence, Manuelli, and Tedrake 2018) is the first work to learn the dense descriptors of objects in a self-supervised manner. Based on DON, MCDONs (Chai, Hsu, and Tsao 2019) introduces several contrastive losses to maintain inter-class separation and intra-class variation. However, a non-disentangled descriptor with two kinds of distinctiveness requires more supervision and increases the difficulty of model convergence. BIND (Chan, Addison Lee, and Kemao 2017) uses multi-layered binary nets to encode edge-based descriptions, which is quite different from the point-based methods. Although the work of semantic correspondence (Yang et al. 2017; J. Lee 2019) learns object-centric semantic representation, they directly predict dense correspondence from different views instead of outputting the high-level descriptors. To the best of our knowledge, jointly learned object-centric keypoint detection and description has not been explored before.

Object-centric Detection and Description

This section presents the problem formulation of the object-centric keypoint detection and description. The details of the sim2real contrastive training method are provided as well.

Formulation and Overall Architecture

Keypoint detection and description is actually an image-based dense prediction problem. It needs to detect whether each pixel (or local patch) in the input image corresponds to the interest point or not. Besides detection, the predicted description vector of each pixel is essential, upon which we

can, for example, compute the similarity between different keypoints in different images. Moreover, this paper studies the object-centric formulation; hence we additionally associate each descriptor with the objectness—the keypoints on the same object are clustered while those on different objects are detached from each other. The formal definitions are as below.

Keypoint detector. Given an input image $I \in \mathbb{R}^{3 \times H \times W}$, the detector outputs a non-negative confidence map $\sigma(I) \in \mathbb{R}_{\geq 0}^{H \times W}$, where H and W are respectively the height and width. The pixel of the i -th row and j -th column denoted as $I[i, j]$ will be considered as a keypoint if $\sigma(I)[i, j] > r_{\text{thr}}$ for a certain threshold $r_{\text{thr}} > 0$.

Keypoint descriptor with objectness. The descriptor is formulated as $\eta(I) \in \mathbb{R}^{C \times H \times W}$, where C represents the dimensionality of the descriptor vector for each pixel. As mentioned above, there are two subparts of each description vector, one for intra-object salience and the other one for inter-object distinctness. We denote them as $\eta_s(I) \in \mathbb{R}^{C_1 \times H \times W}$ and $\eta_c(I) \in \mathbb{R}^{C_2 \times H \times W}$, respectively, where $C_1 + C_2 = C$.

Similar to previous works (such as R2D2), the detector and descriptor share a major number of layers, called an encoder. For the implementation of the encoder, we apply Unet (Ronneberger, Fischer, and Brox 2015) plus ResNet (He et al. 2016) blocks and upsampling layers as the backbone, inspired by Monodepth2 (Godard et al. 2019) which is originally for depth estimation. We have also made some minor modifications regarding the encoder to deliver better expressivity. The output of the encoder serves as input to i) the detector after element-wise square operation and ii) the descriptor after ℓ_2 normalization layer, motivated by design in R2D2 (Revaud et al. 2019). More implementation details are provided in the appendix. The overall architecture is illustrated in Figure 2.

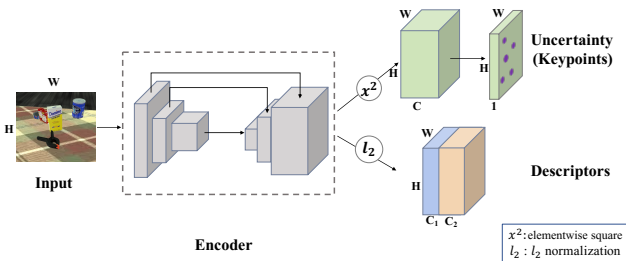


Figure 2: Overview of our network. The uncertainty map is associated with the detector and the descriptor is disentangled into two parts which respectively pursue inter-object salience and inter-object distinctness.

The training of our model is conducted in a simulated environment. In general, for each query image I_1 , we collect its viewpoint-varying version I_2 (which is indeed the adjacent frame of I_1 as our simulation data are videos). Besides, we generate a rendered version of I_1 by retaining the target object only and cleaning all other objects and background; we call this rendered image I_0 . With these three images, we

have the overall training loss as follows.

$$\mathcal{L}(I_1, I_2, I_0) = \mathcal{L}_r(I_1, I_2) + \lambda_1 \mathcal{L}_{d_s}(I_1, I_2, I_0) + \lambda_2 \mathcal{L}_{d_c}(I_1, I_2, I_0), \quad (1)$$

where \mathcal{L}_r stands for enforcing the repeatability of keypoint detector, \mathcal{L}_{d_s} and \mathcal{L}_{d_c} are the objectives of the descriptor for the intra-object salience and inter-object distinctness, respectively, λ_1 and λ_2 are the associated trade-off weights.

Now we introduce each loss. The motivation of minimizing \mathcal{L}_r is to enforce the activation of the detector to be invariant with respect to the change of viewpoint, which is dubbed as repeatability by (Revaud et al. 2019). Here, we introduce a metric operator r different from that used in (Revaud et al. 2019), which combines SSIM (Wang et al. 2004) and ℓ_1 normalization:

$$r(x, y) = \frac{\alpha}{2} (1 - \text{SSIM}(x, y)) + (1 - \alpha) \|x - y\|_1, \quad (2)$$

where $\alpha = 0.85$ by default. We thereby compute $\mathcal{L}_r(I_1, I_2)$:

$$\mathcal{L}_r(I_1, I_2) = \frac{1}{|\mathbb{U}_1|} \sum_{u_1 \in \mathbb{U}_1} r(\sigma(I_1)[u_1], \sigma(I_2)[T_{12}(u_1)]), \quad (3)$$

where \mathbb{U}_1 refers to all $N \times N$ patches around each coordinate in I_1 , $|\mathbb{U}_1|$ is the size of \mathbb{U}_1 , T_{12} is the coordinate transformation between I_1 and I_2 , and thus $T_{12}(u_1)$ returns the corresponding coordinate of u_1 in I_2 .

Both \mathcal{L}_{d_s} and \mathcal{L}_{d_c} are formulated by the contrastive learning strategy (van den Oord, Li, and Vinyals 2019). The only difference lies in the different construction of positive-negative training samples. By omitting the involvement of the rendered sample I_0 , we first discuss the general form with uncertainty and then specify the difference between \mathcal{L}_{d_s} and \mathcal{L}_{d_c} . At last, we further consider the training by adding I_0 .

Contrastive Learning with Uncertainty

We assume the general form of $\mathcal{L}_{d_s}(I_1, I_2)$ and $\mathcal{L}_{d_c}(I_1, I_2)$ (without I_0) to be $\mathcal{L}_c(I_1, I_2)$. A crucial property of keypoint descriptor is that it should be invariant to image transformations between I_1 and I_2 like viewpoint or illumination changes. We thus treat the descriptor learning as a contrastive learning task (Wang et al. 2021). To be specific, we define the query vector in $\eta(I_1)$ as d_1 , the positive and negative vectors in $\eta(I_2)$ as d_2^+ and d_2^- , respectively. According to the definition in (van den Oord, Li, and Vinyals 2019), the contrastive loss with one positive sample d_2^+ and the negative set \mathbb{D}_2^- is given by

$$\mathcal{L}_c(d_1, d_2^+, \mathbb{D}_2^-) = -\log \frac{\exp(d_1 \cdot d_2^+ / \tau)}{\exp(d_1 \cdot d_2^+ / \tau) + \sum_{d_2^- \in \mathbb{D}_2^-} \exp(d_1 \cdot d_2^- / \tau)}, \quad (4)$$

where τ is the temperature parameter.

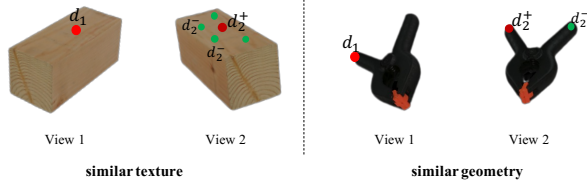


Figure 3: Less-textured and symmetric regions.

The quality of the contrastive samples influences the training performance greatly. Considering some objects with texture-less or symmetry geometries in Figure 3, if the coordinate of d_2^- is distant to d_1 after view projection, then d_2^- is considered as a negative sample (the detailed negative sampling is in next subsection). Then, under the contrastive learning Eq. (4), d_2^- is enforced to be dissimilar to d_1 as much as possible. However, this conflicts with the texture-/geometry distribution of the objects, as the local fields of d_1 and d_2^- are indeed similar. We should design a certain mechanism to avoid such inconsistency. Recall that our focus is on selecting sparse points of interest. The keypoint detector $\sigma(I_1)$ returns the confidence to determine which point should be selected. If we use this confidence to weight the importance of the query sample, we are more potential to filter the unexpected cases as mentioned above.

More specifically, we borrow the uncertainty estimation from (Poggi et al. 2020; Kendall and Gal 2017), which has been proved to improve the robustness of deep learning in many applications. Starting by predicting a posterior $p(\mu|\bar{\mu}, \gamma)$ for the descriptor of each pixel parameterized with its mean μ and variance γ over ground-truth labels μ (Yang et al. 2020), the negative log-likelihood becomes:

$$-\log p(\mu|\bar{\mu}, \gamma) = \frac{|\mu - \bar{\mu}|}{\gamma} + \log \gamma. \quad (5)$$

We adjust this formula into our contrastive learning framework. We first regard the reciprocal of the detection confidence as the uncertainty variance for each pixel, *i.e.* $\gamma = \sigma(I_1)^{-1}$. It is reasonable as the larger confidence it outputs, the smaller uncertainty it exhibits. Second, we replace the error $|\mu - \bar{\mu}|$ with our loss \mathcal{L}_c in Eq. (4), since our goal is to refine the contrastive learning in the first place.

By summation over all queries in I_1 , we derive:

$$\mathcal{L}_d(I_1, I_2) = \frac{1}{M} \sum_{i=1}^M \frac{\mathcal{L}_c(d_1^i, d_2^{i+}, \mathbb{D}_2^{i-})}{(\sigma_1^i)^{-1}} + \log(\sigma_1^i)^{-1}, \quad (6)$$

where, for the i -th query in I_1 , d_1^i indicates the description query vector, d_2^{i+} and \mathbb{D}_2^{i-} are the corresponding positive sample and negative sample set in I_2 , σ_1^i is the detection value, M is the number of all queries.

Disentangled Descriptor Learning

As depicted in Figure 2, the descriptor is learned for two goals: it should not only distinguish different keypoints on the same object but also classify those across different objects. We realise this via two disentangled losses $\mathcal{L}_{d_s}(I_1, I_2)$

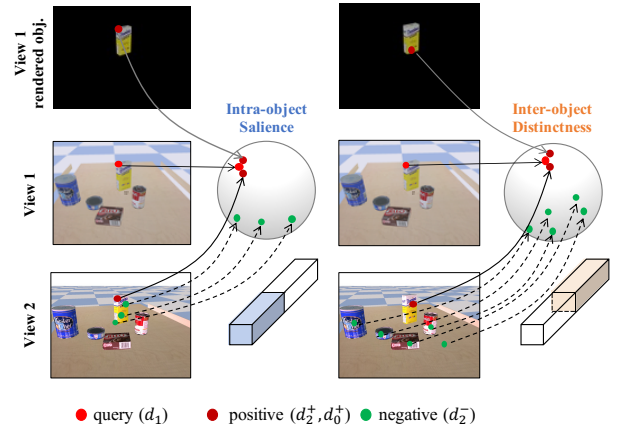


Figure 4: Illustration of the positive and negative samples for a given query. The middle and bottom rows denote the synthetic scenes from two different viewpoints. The top line renders the object with a clean background at view 1. We decouple the descriptor into two parts for learning the intra-object saliency (first column) and the inter-object distinctness (second column).

and $\mathcal{L}_{d_c}(I_1, I_2)$, which respectively follow the general form of $\mathcal{L}_d(I_1, I_2)$ in Eq. (6) and \mathcal{L}_c in Eq. (4). The two losses also employ distinct constructions of the training samples d_2^+ and \mathbb{D}_2^- for any given query d_1 . For better readability, we refer to their constructions as: $d_{2,s}^+$ and $\mathbb{D}_{2,s}^-$, $d_{2,c}^+$ and $\mathbb{D}_{2,c}^-$, respectively.

Intra-object saliency. The loss $\mathcal{L}_{d_s}(I_1, I_2)$ is for intra-object saliency. Suppose the coordinate of the query d_1 in image I_1 to be u_1 , indicating $d_1 = \eta_s(I_1)[u_1]$ where η_s outputs the saliency part of the descriptor as defined before. The positive sample $d_{2,s}^+$ is chosen as the projection from I_1 to I_2 via the view transformation T_{12} ; in other words $d_{2,s}^+ = \eta_s(I_2)[T_{12}(u_1)]$. As for the negative candidates $\mathbb{D}_{2,s}^-$, we pick the points from I_2 on the same object as the query but out of the δ -neighbourhood, namely, $\mathbb{D}_{2,s}^- = \{\eta_s(I_2)[u] \mid \|u - T_{12}(u_1)\|_2 > \delta, l(u) = l(T_{12}(u_1))\}$ where $l(u)$ returns the object label at pixel u . Figure 4 illustrates the sampling process (first column). By iterating over all possible queries, we arrive at the similar form to Eq. (6):

$$\mathcal{L}_{d_s}(I_1, I_2) = \frac{1}{M} \sum_{i=1}^M \frac{\mathcal{L}_c(d_1^i, d_{2,s}^{i+}, \mathbb{D}_{2,s}^{i-})}{(\sigma_1^i)^{-1}} + \log(\sigma_1^i)^{-1}, \quad (7)$$

where $d_{2,s}^{i+}$ and $\mathbb{D}_{2,s}^{i-}$ are the positive sample and negative set for the i -th query. Via Eq. (7), our hope is to accomplish the distinctness between keypoints on the same object.

Inter-object distinctness. We now introduce how to create the samples for $\mathcal{L}_{d_c}(I_1, I_2)$. Different from the above inter-object loss, here any point on the same object as the query in I_2 is considered as the positive sample, that is, $d_{2,c}^+ \in \{\eta_c(I_2)[u] \mid l(u) = l(T_{12}(u_1))\}$ where η_c denotes the inter-object output branch of the descriptor. In terms of the negative samples, the points on other objects or

the background are selected, implying $\mathbb{D}_{2,c}^- = \{\eta_c(I_2)[u] \mid l(u) \neq l(T_{12}(u_1))\}$. The illustration is displayed in Figure 4 (second column). In form, we have the summation over all queries as follows:

$$\mathcal{L}_{d_c}(I_1, I_2) = \frac{1}{M} \sum_{i=1}^M \mathcal{L}_c(d_1^i, d_{2,c}^{i+}, \mathbb{D}_{2,c}^{i-}). \quad (8)$$

By making use of Eq. (7) and (8) together, we obtain more fine-grained information: we can tell if any two keypoints are on the same object, and if yes, we can further know if they correspond to different parts of the object by comparing their intra-descriptors.

Semantic Consistency

This subsection presents how to involve the rendered image I_0 into our contrastive training. Selvaraju et al. (Selvaraju et al. 2021) found the existing contrastive learning models often cheat by exploiting low-level visual cues or spurious background correlations, hindering the expected ability in semantics understanding. This actually happens in our case when the descriptor may use the spatial relationship with other objects or the background. To illustrate this issue, in Figure 5, we assume the left (red) and right (green) edges of the cup’s mouth are a pair of negative samples, their main difference lies in the distance from the handle of the cup. Nevertheless, the local region of the sugar box (yellow box) behind the red dot could be used as a shot-cut reference for the distinctness between the red and green points, which is NOT what we desire. The ideal learning of the descriptor for the object is to make it focus on the object itself.



Figure 5: Illustration of semantic consistency paradigm.

For this purpose, we render the cup according to its pose in I_1 and remove all other things, leading to the image I_0 . We then perform contrastive learning by further taking the positive from I_0 into account following the similar process in Eq. (7) and (8). The whole pipeline is demonstrated in Figure 4. The losses $\mathcal{L}_{d_s}(I_1, I_2)$ and $\mathcal{L}_{d_c}(I_1, I_2)$ are rewritten as follows.

$$\mathcal{L}_{d_s}(I_1, I_2, I_0) = \frac{1}{M} \sum_{i=1}^M \frac{\mathcal{L}_c(d_1^i, d_{2\&0,s}^{i+}, \mathbb{D}_{2,s}^{i-})}{(\sigma_1^i)^{-1}} + \log(\sigma_1^i)^{-1}, \quad (9)$$

$$\mathcal{L}_{d_c}(I_1, I_2, I_0) = \frac{1}{M} \sum_{i=1}^M \mathcal{L}_c(d_1^i, d_{2\&0,c}^{i+}, \mathbb{D}_{2,c}^{i-}), \quad (10)$$

where $d_{2\&0,s}^{i+}$ denotes the union of the positive samples $d_{2,s}^{i+}$ from I_2 and $d_{0,s}^{i+}$ from I_0 for the i -th query, and $d_{2\&0,c}^{i+}$ is defined similarly.

Experiments

Training data generation. To bootstrap the object-centric keypoint detection and description, we first create a large-scale object-clustered synthetic dataset that consists of 21 objects from the YCB-Video dataset (Xiang et al. 2018). Furthermore, to align the synthetic and real domains, we refer to the idea of physically plausible domain randomization (PPDR) (Wen et al. 2020) to generate the scenes where objects can be fallen onto the table/ground with preserving physical properties. The viewpoint of the camera is randomly sampled from the upper hemisphere. We construct a total of 220 scenes where each contains 6 objects and acquire a continuous sequence of images from each scene, resulting in 33k images. More examples of the dataset can be found in the appendix.

Training. We choose 20 keypoints of each object to construct positive-negative pairs, and the temperature τ in intra-object InfoNCE (van den Oord, Li, and Vinyals 2019) loss and inter-object InfoNCE loss are set to 0.07, 0.2 respectively. The data augmentation is composed of color jittering, random gray-scale conversion, gaussian noise, gaussian blur, and random rotation. The δ and N are set to 8 and 16 pixels. We set the trade-off weights of two subparts of descriptor $\lambda_1 = 1$ and $\lambda_2 = 1$.

Our model is implemented in PyTorch (Paszke et al. 2019) with a mini-batch size of 4 and optimized with the Adam (Kingma and Ba 2017) for 20 epochs, and all the input images are cropped to 320×320 . We use a learning rate of 10^{-4} for the first 15 epochs, which is dropped ten times for the remainder.

Testing. To reduce the domain discrepancy between training and test data, we modify the statistics of BN (Ioffe and Szegedy 2015) learned in simulation for adapting the model to real scenes. Strictly speaking, we suppose that the real test data cannot be accessed, so we only use the mean and variance of BN layers from a current real image, i.e., the batch size is set to 1, and do not update the statistics. For comparative evaluations, we record the best results of each baseline which adopts the BN layer, with or without this trick.

Baselines. As a classical keypoint detection and description, we choose the handcrafted method SIFT (Lowe 2004) as the baseline. We also compare against Superpoint (DeTone, Malisiewicz, and Rabinovich 2018), R2D2 (Revaud et al. 2019) and DISK (Tyszkiewicz, Fua, and Trulls 2020) which are data-driven methods.

Image Matching

All the methods are evaluated on the following datasets.

YCB-Video (Xiang et al. 2018) consists of 21 objects and 92 RGB-D video sequences with pose annotations. We use the 2,949 keyframes in 12 videos which are commonly evaluated in other works. In this scene, all the objects are static, and the camera is moving with a slight pose change.

Objctcs	SIFT(128)			Superpoint(256)			R2D2(128)			DISK(128)			Ours(96)		
	Kpts	MMA5	MMA7	Kpts	MMA5	MMA7	Kpts	MMA5	MMA7	Kpts	MMA5	MMA7	Kpts	MMA5	MMA7
cracker_box	72.1	23.4%	27.4%	21.9	25.8%	32.7%	26.1	22.6%	29.0%	41.2	26.2%	32.1%	122.2	37.8%	49.8%
sugar_box	25.2	5.6%	6.8%	9.2	19.1%	25.3%	12.8	9.6%	15.0%	14.4	15.7%	21.0%	64.4	18.9%	29.4%
tomato_soup	21.3	8.7%	11.6%	9.7	47.8%	57.1%	9.6	41.3%	46.0%	11.7	66.8%	73.4%	54.0	60.1%	70.1%
mustard_bottle	27.6	18.8%	21.6%	10.2	25.6%	32.5%	12.8	26.3%	29.9%	19.2	41.2%	51.9%	88.1	43.6%	61.1%
bleach_s	27.9	11.6%	12.0%	10.7	27.6%	35.0%	16.6	19.8%	22.9%	15.9	22.8%	25.4%	72.7	33.6%	42.2%
ALL	34.8	13.6%	15.9%	12.3	29.2%	36.5%	15.6	23.9%	28.5%	20.5	34.5%	40.7%	80.3	38.8%	50.5%

Table 1: Quantitative evaluation for real-real image matching.

YCBInEOAT (Wen et al. 2020) consists of 9 video sequences and each video has one manipulated object from YCB-Video. In this dataset, objects are translated and rotated by different end-effectors while the camera is static. We selected 5 valid videos with a total of 1112 keyframes.

We set two object-centric image matching tasks.

- **Synthetic-real matching.** The test images are 2949 keyframes from YCB-Video. Two adjacent frames are selected where the next frame is adopted as a target (real) image, and the rendered (synthetic) images on the previous pose of each object are used as the references. Pairs are matched and filtered by RANSAC (Fischler and Bolles 1981).
- **Real-real matching.** The keypoints and descriptors from the manipulated object whose mask is known in the initial frame of each video, predicted by keypoint methods, are matched with the subsequent frames (targets) to show the tracking performance on the object keypoints.

In this work, the bounding box or mask of each object in the target frame is not provided. We utilize the nearest neighbor search to find the matched keypoints from two views, i.e., mutual nearest neighbors are considered matches. We adopt the Mean Matching Accuracy (MMA) (Mikolajczyk and Schmid 2005) for matching evaluation, i.e., the average percentage of correct matches per image pair. A correct match represents its reprojection error, which is below a given matching threshold. We record the MMA5 and MMA7 of each object with an error threshold of 5 and 7 pixels, respectively. Each method would detect top-5k keypoints per image. Kpts means the average number of matches per object per image. Dim is the length of descriptors.

Comparison to baselines. In terms of synthetic-real matching, our object-centric method significantly outperforms the scene-centric methods, as shown in Table 2. Our method surpasses all the counterparts by a large margin, i.e., more than 20% in the MMA5 and MMA7. We can extract more matching keypoints with broader distribution on the surface of the objects, while the matches are not affected by other occluded objects, as shown in Figure 6. It should be noted that the scene-level methods can not extract matching keypoints in the junction of object and background, since the different backgrounds between the target and the rendered objects.

In Table 1, we further provide a quantitative comparison with the baselines in the real-real matching track. Compared with the baselines, our method also clearly attains consistent improvements. And it shows that our method outperforms the Superpoint by at least 68.0 Kpts, the R2D2 by 64.7 Kpts

Method	Dim	Kpts	MMA5	MMA7
SIFT	128	16.9	24.2%	30.1%
Superpoint	256	15.7	33.6%	43.8%
R2D2	128	20.0	34.8%	44.6%
DISK	128	15.8	28.2%	35.1%
Ours	96	92.1	50.0%	57.2%

Table 2: Evaluation for synthetic-real image matching. The last three metrics are averaged of each object.

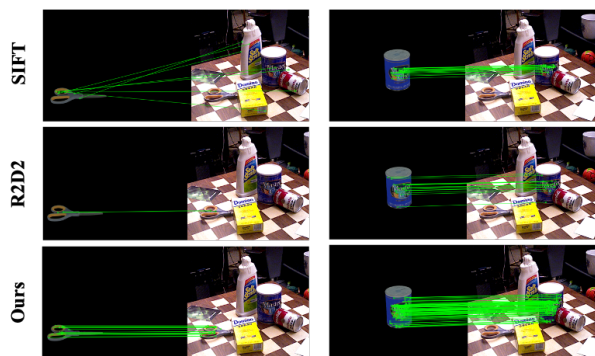


Figure 6: Synthetic-real image matching.

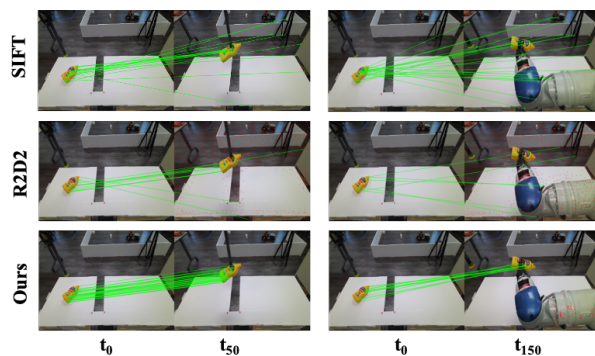


Figure 7: Real-real image matching. $t_k = k$ -th frame.

and the DISK 59.8 Kpts. In the MMA7, our method surpasses the DISK by 9.8% and the Superpoint by 14.0%. Despite the obvious movement of the target in the scene, our method can still detect the matching points. In contrast, the performance of other methods will deteriorate significantly, as shown in Figure 7. It shows the encouraging generalization ability of our method from simulation to reality. More details of matching results can be seen in the appendix.

Objects	Keypoint-Based Method										End2End Method	
	Scene-Centric								Object-Centric		(Supervised)	
	SIFT(128)		S-Point(256)		R2D2(128)		DISK(128)		Ours(96)		PoseCNN	
	ADD	ADDS	ADD	ADDS	ADD	ADDS	ADD	ADDS	ADD	ADDS	ADD	ADDS
master_chef_can	20.1	38.4	48.5	76.8	31.6	56.2	22.2	44.9	48.2	77.5	50.9	84.0
cracker_box	0.0	3.6	55.2	69.0	39.8	57.9	46.5	60.1	35.4	60.3	51.7	76.9
sugar_box	28.7	35.8	65.1	79.0	53.0	61.5	49.2	60.2	74.6	86.3	68.6	84.3
tomato_soup_can	19.4	27.0	42.1	57.9	49.1	60.8	33.1	42.6	56.5	75.3	66.0	80.9
mustard_bottle	9.7	13.1	47.1	52.2	40.4	46.8	31.0	38.3	54.5	72.8	79.9	90.2
tuna_fish_can	7.8	11.4	13.8	21.1	1.6	2.2	0.2	0.7	53.6	73.5	70.4	87.9
pudding_box	2.9	6.8	5.8	8.2	0.5	0.7	5.3	8.8	37.9	49.0	62.9	79.0
gelatin_box	68.2	79.8	55.4	68.0	34.1	39.4	54.9	65.0	55.7	70.8	75.2	87.1
potted_meat_can	8.1	11.4	30.4	38.4	23.3	32.1	24.1	30.4	51.9	70.5	59.6	78.5
banana	0.3	0.6	0.0	0.8	0.3	0.5	0.5	2.2	11.5	28.5	72.3	85.9
pitcher_base	0.0	3.2	3.4	9.8	0.9	3.2	0.2	3.8	16.9	30.9	52.5	76.8
bleach_cleanser	17.6	22.4	43.3	54.4	37.2	49.5	38.7	50.3	39.8	55.2	50.5	71.9
bowl	0.0	0.3	0.2	0.9	0.2	0.7	0.3	3.5	3.7	26.8	6.5	69.7
mug	0.2	0.3	0.2	0.9	0.2	0.2	0.3	1.2	14.3	45.7	57.7	78.0
power_drill	0.9	2.9	55.8	63.3	36.5	43.1	15.2	21.5	61.4	76.5	55.1	72.8
wood_block	0.0	2.2	0.8	4.1	0.0	0.0	0.0	0.4	3.2	22.1	31.8	65.8
scissors	0.0	0.0	1.0	2.1	0.0	0.0	0.0	0.6	20.9	38.3	35.8	56.2
large_marker	22.3	26.9	24.3	34.8	17.2	18.9	0.6	0.7	56.8	68.5	58.0	71.4
large_clamp	0.0	0.4	1.0	3.8	0.0	0.1	0.1	0.5	14.8	36.8	25.0	49.9
ex_large_clamp	0.0	0.6	0.6	4.2	0.1	0.5	0.2	0.5	14.5	45.5	15.8	47.0
foam_brick	0.0	0.0	0.7	1.3	0.0	0.0	0.6	1.2	43.4	69.5	40.4	87.8
ALL	10.6	15.2	30.3	39.9	23.4	30.6	19.0	25.9	42.2	61.9	53.7	75.9

Table 3: Evaluation for 6D pose estimation on YCB-Video dataset.

6D Pose Estimation

Evaluation protocol. Pipeline of pose estimation : (1) Render multiple images in different poses of the test objects as templates. To balance the evaluation speed and accuracy, we rendered 96 templates for each object. (2) Match the templates with the real images one by one, and select the best template according to the number of matched pairs. (3) 6D pose can be solved through the Perspective-n-Point (PnP) and RANSAC algorithms. We report the average recall (%) of ADD(-S) for pose evaluation which is the same as in PoseCNN (Xiang et al. 2018).

Comparison to baselines. In Table 3, our method achieves the state-of-the-art performance among other keypoint-based methods with a large margin (approximately 20% on both ADD and ADD-S). The pose estimation results of each object can be seen in the appendix. By introducing an object-centric mechanism, our method can significantly bootstrap performance on 6D pose estimation. As shown in Figure 8, the keypoint matching result indicates our method can serve a crucial role on downstream 6D pose estimation tasks, even in the occluded and clustered environment. Figure 8 displays some image matching and 6D pose estimation results. It can be seen that ours can detect matching points even when the object is severely occluded or in a large pose difference between a template and target.

Comparison to un/weak-supervised pose estimation. Self6D (Wang et al. 2020) is a sim2real pose estimation method, where the model is fully trained on the synthetic RGB data in a self-supervised way. Further, the model is fine-tuned on an unannotated real RGB-D dataset, called self6D(R). Note that the self6D(R) is a weak-supervised

baseline due to the access to real-world data. In Table 4, our method achieves an overall average recall of 59.4%, which surpasses 10.8% of Self6D(R) and 29.5% of Self6D. And our method brings the encouraging generalization ability from simulation to reality, while the performance of PoseCNN only trained in the simulation would drop dramatically.

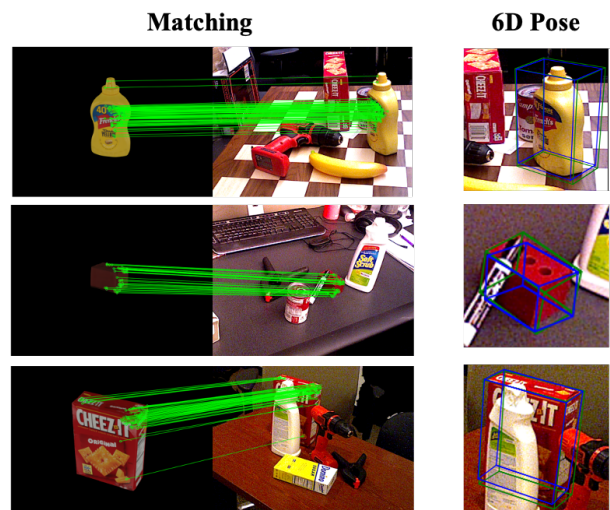


Figure 8: Examples of 6D object pose estimation results on YCB-Video by our method. Left: the keypoints matching between the best template and a real image. Right: the pose estimation results, where the green bounding boxes stands for ground truths while the blue boxes are our predictions.




Figure 9: Visualization of confidence heatmap in the detector. The highlighted regions are selected as keypoints.

Objects	W-Supervised	Sim2Real/Unsupervised		
	Self6D(R)	Self6D	PoseCNN	Ours
mustard_bottle	88.2	73.7	3.7	72.8
tuna_fish_can	69.7	26.6	3.1	73.5
banana	10.3	4.0	0.0	28.5
mug	43.4	23.9	0.0	45.7
power_drill	31.4	21.4	0.0	76.5
ALL	48.6	29.9	1.4	59.4

Table 4: Results for un/weak-supervised on YCB-Video dataset. W-Supervised=Weak-Supervised.

Ablation Study

In this subsection, we will explore the sensitivity of our method in terms of reliability threshold and perform a diverse set of analyses on assessing the impact of each component that contributes to our method.

Sensitivity of confidence threshold. Table 5 illustrates performance of keypoint matching and pose estimation under various confidence threshold r_{thr} , i.e., 0.0, 1.5, 3.0, 4.5. In particular, $r_{thr} = 0$ means that all pixels may be selected as keypoints. Obviously, such performance will be significantly deteriorated. We use confidence threshold r_{thr} being 1.5 as the default setting, as the performance becomes superior in all tasks. We visualize the dense detector heatmap of some objects, i.e., confidence of pixels.

Training Strategy. In Table 6, we compare different training strategies used to learn detector and descriptor. 1) if we randomly select negative samples from all the pixels, like R2D2, it will exhibit extremely poor performance. Object-centric sample strategy will be essential for object keypoint detection. 2) Each module, including the repeatability of keypoint detector (\mathcal{L}_r), the decoupled descriptor (Dec.), and the semantics consistency (Sem.) have positive effects on the final result.

Model Efficiency and Generalization

Efficiency. Our model costs about 0.51s to extract keypoints and descriptors from a 640×480 image, while SIFT, R2D2, Superpoint, and DISK take about 0.04s, 0.10s, 0.19s, and 0.48s, respectively. It indicates that the computation overhead by our method is acceptable, particularly given the remarkable improvement in performance.

Generalization. The generalization ability of our method lies in two aspects. The first one is the sim2real adaptation. Our model is trained on simulation data and tested straightly

r_{thr}	Syn-Real		Real-Real		6D Pose	
	Kpts	MMA5	Kpts	MMA5	ADD	ADDS
0.0	10.0	14.2%	43.7	3.0%	15.8	24.3
1.5	92.1	50.0%	80.3	38.8%	42.2	61.9
3.0	80.4	45.1%	78.9	39.3%	40.2	59.3
4.5	71.1	41.0%	67.7	41.0%	36.1	53.3

Table 5: Sensitivity to reliability threshold r_{thr} .

NEG R. Obj.	\mathcal{L}_r	Dec.	Sem.	Syn-Real	Real-Real	6D Pose	
				MMA5	MMA5	ADD	ADDS
✓				1.1%	2.3%	1.5	2.0
✓				46.9%	29.3%	22.1	33.8
✓	✓			45.5%	35.1%	26.7	42.7
✓	✓	✓		47.8%	44.4%	40.8	58.2
✓	✓		✓	52.4%	30.0%	37.6	54.7
✓	✓	✓	✓	<u>50.0%</u>	<u>38.8%</u>	42.4	61.8

Table 6: Ablation study of training strategy. NEG./R.= negatives randomly sampled from all the pixels; NEG./Obj.= object-centric sampling; \mathcal{L}_r = repeatability of detector; Dec.= decoupled descriptors; Sem.= semantics consistency. We bold the best and underline the second best.

on real images. Such sim2real generalization benefits the field of robotic perception/manipulation where the keypoint annotations are difficult to obtain, but the CAD model of a target object is always given. The second one is the generalization to unseen objects. To illustrate, we provide the matching evaluations on objects with simulated images outside the training set in Table 7. Twenty unseen objects are selected from the OCRTOC dataset (Liu et al. 2021), in which half of them have the same class as YCB-Video objects but with different shapes or textures (seen class), and other objects with novel class have not been seen in training (unseen class). The evaluation protocol is similar to real-real matching. From Table 7, it can be seen that our model also achieves satisfactory matching performance for unseen objects. More details are provided in the appendix.

Method	Seen class		Unseen class	
	Kpts	MMA5	Kpts	MMA5
SIFT	23.3	32.5%	16.8	29.6%
R2D2	21.2	61.3%	16.1	49.3%
Ours	90.5	65.4%	75.9	61.3%

Table 7: Image matching evaluation on unseen objects.

Conclusion

We present for the first time a sim2real contrastive learning framework for object-centric keypoint detection and description, which is only trained from synthetic data. Our experiments demonstrate that (1) our object keypoint detector and descriptor can be robust for both synthetic-to-real and real-to-real image matching tasks, and (2) our method leads to a superior result on unsupervised (sim2real) 6D pose estimation. Future work may explore integrating 2D and 3D inputs to find more repeatable keypoints and distinctive descriptors for texture-less objects.

Acknowledgments

This research was funded by the National Science and Technology Major Project of the Ministry of Science and Technology of China (No.2018AAA0102900). Meanwhile, this work is jointly sponsored by the National Natural Science Foundation of China (Grant No. 62006137) and CAAI-Huawei MindSpore Open Fund. It was also partially supported by the National Science Foundation of China (NSFC) and the German Research Foundation (DFG) in the project Cross Modal Learning, NSFC 61621136008/DFG TRR-169.

References

- Balntas, V.; Johns, E.; Tang, L.; and Mikolajczyk, K. 2016. PN-Net: Conjoined Triple Deep Network for Learning Local Image Descriptors. arXiv:1601.05030.
- Barroso-Laguna, A.; Riba, E.; Ponsa, D.; and Mikolajczyk, K. 2019. Key. net: Keypoint detection by handcrafted and learned cnn filters. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5836–5844.
- Bay, H.; Tuytelaars, T.; and Van Gool, L. 2006. SURF: Speeded Up Robust Features. In Leonardis, A.; Bischof, H.; and Pinz, A., eds., *Computer Vision – ECCV 2006*, 404–417.
- Chai, C.-Y.; Hsu, K.-F.; and Tsao, S.-L. 2019. Multi-step pick-and-place tasks using object-centric dense correspondences. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 4004–4011. IEEE.
- Chan, J.; Addison Lee, J.; and Kema, Q. 2017. BIND: Binary integrated net descriptors for texture-less object recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2068–2076.
- DeTone, D.; Malisiewicz, T.; and Rabinovich, A. 2018. SuperPoint: Self-Supervised Interest Point Detection and Description. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Dusmanu, M.; Rocco, I.; Pajdla, T.; Pollefeys, M.; Sivic, J.; Torii, A.; and Sattler, T. 2019. D2-net: A trainable cnn for joint description and detection of local features. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8092–8101.
- Fischler, M. A.; and Bolles, R. C. 1981. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Commun. ACM*, 24(6): 381–395.
- Florence, P.; Manuelli, L.; and Tedrake, R. 2018. Dense Object Nets: Learning Dense Visual Object Descriptors By and For Robotic Manipulation. *Conference on Robot Learning*.
- Godard, C.; Mac Aodha, O.; Firman, M.; and Brostow, G. J. 2019. Digging into Self-Supervised Monocular Depth Prediction.
- Han, X.; Leung, T.; Jia, Y.; Sukthankar, R.; and Berg, A. C. 2015. MatchNet: Unifying feature and metric learning for patch-based matching. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3279–3286.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Ioffe, S.; and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, 448–456. PMLR.
- J. Lee, J. P. B. H., D. Kim. 2019. SFNet: Learning Object-aware Semantic Flow. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Kendall, A.; and Gal, Y. 2017. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Kingma, D. P.; and Ba, J. 2017. Adam: A Method for Stochastic Optimization. arXiv:1412.6980.
- Kulkarni, T. D.; Gupta, A.; Ionescu, C.; Borgeaud, S.; Reynolds, M.; Zisserman, A.; and Mnih, V. 2019. Unsupervised Learning of Object Keypoints for Perception and Control. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Liu, Z.; Liu, W.; Qin, Y.; Xiang, F.; Gou, M.; Xin, S.; Roa, M. A.; Calli, B.; Su, H.; Sun, Y.; and Tan, P. 2021. OCRTOC: A Cloud-Based Competition and Benchmark for Robotic Grasping and Manipulation. arXiv:2104.11446.
- Lowe, D. G. 2004. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 60: 91–110.
- Mikolajczyk, K.; and Schmid, C. 2005. A performance evaluation of local descriptors. *IEEE transactions on pattern analysis and machine intelligence*, 27(10): 1615–1630.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32, 8024–8035. Curran Associates, Inc.
- Peng, S.; Liu, Y.; Huang, Q.; Zhou, X.; and Bao, H. 2019. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4561–4570.
- Piasco, N.; Sidibé, D.; Demonceaux, C.; and Gouet-Brunet, V. 2018. A survey on Visual-Based Localization: On the benefit of heterogeneous data. *Pattern Recognition*, 74: 90–109.
- Poggi, M.; Aleotti, F.; Tosi, F.; and Mattocchia, S. 2020. On the Uncertainty of Self-Supervised Monocular Depth Estimation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3224–3234.
- Rad, M.; and Lepetit, V. 2017. BB8: A Scalable, Accurate, Robust to Partial Occlusion Method for Predicting the 3D Poses of Challenging Objects without Using Depth. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 3848–3856.
- Revaud, J.; Weinzaepfel, P.; de Souza, C. R.; and Humenberger, M. 2019. R2D2: Repeatable and Reliable Detector and Descriptor. In *Advances in Neural Information Processing Systems*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. arXiv:1505.04597.
- Sadran, E.; Wurm, K. M.; and Burschka, D. 2013. Sparse keypoint models for 6D object pose estimation. In *2013 European Conference on Mobile Robots*, 307–312. IEEE.
- Selvaraju, R. R.; Desai, K.; Johnson, J.; and Naik, N. 2021. CASTing Your Model: Learning To Localize Improves Self-Supervised Representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11058–11067.
- Smith, S. M.; and Brady, J. M. 1997. SUSAN—a new approach to low level image processing. *International journal of computer vision*, 23(1): 45–78.

Strecha, C.; von Hansen, W.; Van Gool, L.; Fua, P.; and Thoennessen, U. 2008. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, 1–8.

Trajković, M.; and Hedley, M. 1998. Fast corner detection. *Image and vision computing*, 16(2): 75–87.

Tyszkiewicz, M. J.; Fua, P.; and Trulls, E. 2020. DISK: Learning local features with policy gradient. *arXiv preprint arXiv:2006.13566*.

van den Oord, A.; Li, Y.; and Vinyals, O. 2019. Representation Learning with Contrastive Predictive Coding. *arXiv:1807.03748*.

Vecerik, M.; Regli, J.-B.; Sushkov, O.; Barker, D.; Pevcevičute, R.; Rothörl, T.; Schuster, C.; Hadsell, R.; Agapito, L.; and Scholz, J. 2020. S3K: Self-Supervised Semantic Keypoints for Robotic Manipulation via Multi-View Consistency. *arXiv preprint arXiv:2009.14711*.

Verdie, Y.; Yi, K.; Fua, P.; and Lepetit, V. 2015. Tilde: A temporally invariant learned detector. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5279–5288.

Wang, G.; Manhardt, F.; Shao, J.; Ji, X.; Navab, N.; and Tombari, F. 2020. Self6d: Self-supervised monocular 6d object pose estimation. In *European Conference on Computer Vision*, 108–125. Springer.

Wang, X.; Zhang, R.; Shen, C.; Kong, T.; and Li, L. 2021. Dense contrastive learning for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3024–3033.

Wang, Z.; Bovik, A.; Sheikh, H.; and Simoncelli, E. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612.

Wen, B.; Mitash, C.; Ren, B.; and Bekris, K. 2020. se(3)-TrackNet: Data-driven 6D Pose Tracking by Calibrating Image Residuals in Synthetic Domains. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

Xiang, Y.; Schmidt, T.; Narayanan, V.; and Fox, D. 2018. PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes. In *Robotics: Science and Systems (RSS)*.

Yang, F.; Li, X.; Cheng, H.; Li, J.; and Chen, L. 2017. Object-Aware Dense Semantic Correspondence. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4151–4159.

Yang, N.; Stumberg, L. v.; Wang, R.; and Cremers, D. 2020. D3VO: Deep Depth, Deep Pose and Deep Uncertainty for Monocular Visual Odometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Yi, K. M.; Trulls, E.; Lepetit, V.; and Fua, P. 2016. Lift: Learned invariant feature transform. In *European conference on computer vision*, 467–483. Springer.

Zhao, W.; Zhang, S.; Guan, Z.; Zhao, W.; Peng, J.; and Fan, J. 2020. Learning Deep Network for Detecting 3D Object Keypoints and 6D Poses. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14122–14130.