

Multi-path Learning for Object Pose Estimation Across Domains

Martin Sundermeyer^{1,2}, Maximilian Durner^{1,2}, En Yen Puang¹, Zoltan-Csaba Marton¹,
Narunas Vaskevicius³, Kai O. Arras³, Rudolph Triebel^{1,2}

¹German Aerospace Center (DLR), ²Technical University of Munich (TUM),

³Robert Bosch GmbH

Abstract

We introduce a scalable approach for object pose estimation trained on simulated RGB views of multiple 3D models together. We learn an encoding of object views that does not only describe an implicit orientation of all objects seen during training, but can also relate views of untrained objects. Our single-encoder-multi-decoder network is trained using a technique we denote "multi-path learning": While the encoder is shared by all objects, each decoder only reconstructs views of a single object. Consequently, views of different instances do not have to be separated in the latent space and can share common features. The resulting encoder generalizes well from synthetic to real data and across various instances, categories, model types and datasets. We systematically investigate the learned encodings, their generalization, and iterative refinement strategies on the ModelNet40 and T-LESS dataset. Despite training jointly on multiple objects, our 6D Object Detection pipeline achieves state-of-the-art results on T-LESS at much lower runtimes than competing approaches.¹

1. Introduction

Object pose estimation, i.e. estimating the 3D rotation and translation, is a cornerstone in many perception related applications like augmented reality or robotics. For many years this field was dominated by template- and feature-based approaches. Thereby, a given object model or a set of extracted object features is matched into the underlying sensor data. Depending on the quality of the data and the difficulty of the scene, these approaches are still quite competitive. [22] In recent years however, machine learning, and in particular Deep Learning (DL) techniques have emerged as another key approach []. They are often computationally efficient and can show higher robustness towards sensor noise, clutter and environment changes [45, 35, 41]. On

¹Code can be found here: <https://github.com/DLR-RM/AugmentedAutoencoder/tree/multipath>

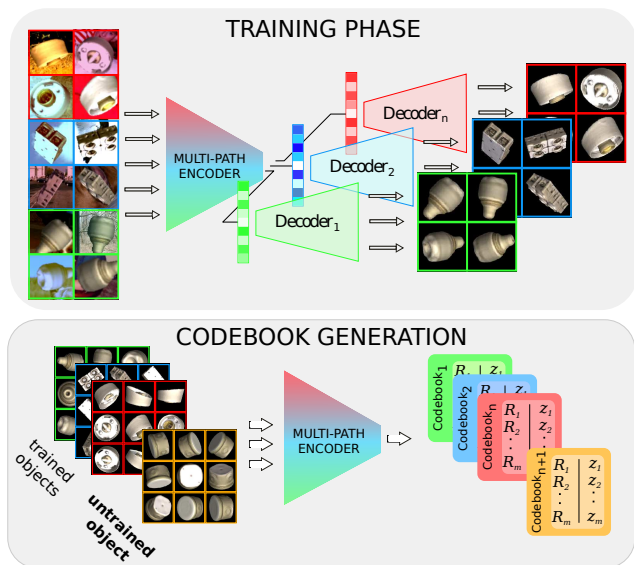


Figure 1: Training (top) and setup phase (bottom) of the MP-Encoder. During training one encoder is shared among all objects, while each decoder reconstructs views of a single object. This turns the encoder into a viewpoint-sensitive feature extractor, that generates expressive encodings for multiple trained and even untrained objects.

the other hand, the need for pose-annotated data as well as lengthy training phases make them less flexible compared to traditional methods. To tackle the lack of annotations, several methods have recently shifted to train on synthetic data rendered from 3D models [23, 34, 38]. However, it is still common to train individual models for every newly encountered instance which is not scalable. Attempts to train on many objects at once often result in deteriorated performance [45]. Furthermore, due to object discriminative training, most of these approaches are not suited to generalize to untrained objects. Since the real world consists of large amounts of object categories and instances, we propose a more adaptive and scalable approach in this paper.

Inspired by Augmented Autoencoders (AAEs) [34] that

extract pose representative features on an instance level, we propose a single-encoder-multi-decoder network for jointly estimating the 3D object orientation of multiple objects. While the encoder and latent space are shared among all objects, the decoders are trained to reconstruct views of specific instances (top of Fig. 1). This *multi-path learning* strategy allows similarly shaped objects to share the same encoding space. After training, we can generate instance-specific codebooks containing encodings of synthetic object views from all over $SO(3)$. So each entry contains a shape and viewpoint dependent descriptor mapped to an explicit 3D object orientation that can be retrieved at test time. (bottom of Fig. 1)

As we show experimentally, the learned encodings generalize well to the real domain and are expressive enough to relate views from all trained and even untrained objects in a viewpoint sensitive way. For a large number of objects the performance of the Multi-Path Encoder (MP-Encoder) does not deteriorate compared to separately trained encoders for each object and even slightly improves encoding quality which leads to new state-of-the-art results in 6-DoF object pose estimation on the T-LESS dataset.

Motivated by this, we also introduce an iterative render-inference scheme based on the learned encodings, which enables relative pose estimation on untrained objects. It resembles the online creation of a local codebook which also helps avoiding $SO(3)$ discretization errors. We apply this method to iteratively refine poses of untrained instances from ModelNet40 and outperform DeepIM [25], a state-of-the-art approach for RGB-based 6D pose refinement.

2. Related Work

Object pose estimation has been widely studied in the literature, mainly from the below aspects.

2.1. Feature-/Template-based approaches

Traditionally, pose estimation methods rely on matching algorithms with local keypoints and descriptors or templates. *E.g.*, both Hinterstoisser *et al.* [17] and Vidal *et al.* [40] apply so-called Point Pair Features (PPF) to match a 3D model into a given scene outperforming other results in the BOP challenge [21] but having high runtimes per target due to the exhaustive search in the large 6D pose space. While these methods require depth information, which can be noisy and sensible to environment changes, Ulrich *et al.* [39] use single RGB images and estimate object poses by matching edges of a rendered 3D model into the image. In general, the benefit of such approaches using explicit matching is that no costly training process is required. However, they tend to require longer execution times. Additionally, usually more expert knowledge is required to set up such pipelines, as feature extraction and matching parameters do not generalize well.

2.2. Learning-based Approaches

More recently, a number of learning-based approaches have been presented with the advantage of reduced computation time and more complex features that can be learned directly from data. One approach is to classify or regress a representation of the object pose directly [23, 28]. Xi-ang *et al.* [44] propose a CNN architecture with multiple stages, where the authors define the pose regression loss as a closest object point measure. Another line of research makes predictions on sparse or dense 2D-3D keypoint locations, *e.g.* the 3D bounding box corners [32, 38, 37] or 3D object coordinates on the surface [26, 45, 31]. The pose is then computed by solving the Perspective-n-Point (PnP) problem. Apart from occlusions, object and view symmetries cause challenges since they introduce one-to-many mappings from object views to poses. While there are strategies to adapt the loss, network or labels for symmetric objects [42, 15, 23], dealing with view-dependent or occlusion-induced symmetries is often intractable and negatively affects the training process. Therefore, a third approach is to learn descriptors of object views that are either conditioned on the pose [1, 42] or completely implicit [34]. The advantage of the implicit approach is that symmetric views share the same target representation during training. Corresponding explicit poses are automatically grouped and can be retrieved at test time. These view-centric representations were also shown to work well for 6D tracking [7].

2.3. Synthetic training data

The demand of learning-based methods for large amounts of labeled training data is particularly unfavorable for object pose estimation. Although there are ways to label object poses semi-automatically [30], it remains a cumbersome, error-prone process. Recent research tackles this by exploiting synthetic training data [8] which solves this problem but creates a new one through the domain gap between real and synthetic data.

To bridge this gap, Tremblay *et al.* [38] use mixed training data consisting of *photo-realistic rendering* and *domain randomized* samples. Since photo-realistic renderings are costly, we follow a less expensive Domain Randomization (DR) strategy [23, 37]: 3D models are rendered in poses randomly sampled from $SO(3)$ using OpenGL with random lighting and superimposed onto real images from data sets such as MS COCO [27] or Pascal VOC [10].

2.4. Generalization to Novel Objects

Most learning-based methods predict the pose of one [31] or few instances [38, 45, 37, 44] and have to be re-trained for every newly encountered object. However, in domains like service and industrial robotics or augmented reality, it would be beneficial to have a general feature extractor that could produce pose-sensitive features for un-

trained objects such that testing on a novel object becomes immediately possible.

When trained on few instances, current pose networks like [38, 37] concurrently classify objects which potentially hinders their ability to generalize to untrained objects. Wohlhart *et al.* [42] and Balntas *et al.* [1] were the first to report qualitative results of deep pose descriptors applied to untrained objects. However, their descriptors are discriminated by both, orientation and object class. So if an untrained object has similar appearance from any viewpoint to one of the trained objects, the corresponding descriptor will get corrupted. Unlike [42, 1], our multi-path training strategy does not separate different object instances in the encoding space and instead allows them to share the same latent features.

Category-level pose estimation [33, 3, 36] can be used to generalize to novel objects from a given category. It assumes that all instances within a category have similar shape and are aligned in a joint coordinate frame. However, these assumptions often do not hold in practice where semantic and geometric similarities often do not coincide. Assigning aligned coordinate frames can be ambiguous because symmetries of instances can vary within a category. Therefore, in this work, we will not explicitly enforce the alignment within semantic categories and instead leave this decision to the self-supervised, appearance-based training.

CNNs trained on large datasets are frequently used to extract low-level features for downstream tasks, e.g. image retrieval [13] or clustering [9]. A naive baseline to predict the 3D orientation of unknown objects would be to compare feature maps of a network trained on a large dataset like ImageNet or COCO. Unsurprisingly this baseline does not work at all because (1) early features are sensitive to translation while the later layers lost geometric information (2) features strongly differ for synthetic and real object views (3) the dimensionality of feature maps is too high to finely discretize $SO(3)$ while reduction techniques like PCA destroy much information.

The pose refinement methods [29, 25] iteratively predict rotation and translation residuals between an estimated and a target view of an object. The former was shown to generalize to untrained objects of the same category, and the latter even generalizes to objects of new categories. These approaches predict an accurate, relative transformation between two object views in a local neighborhood. In contrast, our method is able to yield both, local relative and global 3D orientation estimates.

3. Method

We will first briefly describe the AAE (Sec. 3.1). Building upon these results, we then propose a novel Multi-Path Encoder-Decoder architecture and training strategy (Sec. 3.2). Next, we will investigate the encoder on its ability

to extract pose-sensitive features and examine generalization to novel instances (Sec. 3.3). Different application scenarios depending on the test conditions are discussed (Sec. 3.4). Finally, an iterative render-inference optimization for pose refinements is presented (Sec. 3.5).

3.1. Implicit Object Pose Representations

Sundermeyer *et al.* [34] have shown that implicit pose representations can be learned in a self-supervised way using an encoder-decoder architecture. This so-called AAE allows to encode 3D orientations from arbitrary object views, generalizes from synthetic training data to various test sensors, and inherently handles symmetric object views.

The AAE is trained to reconstruct rendered views of a single object. To encode 3D orientation exclusively, the inputs are randomly translated and scaled while the reconstruction target stays untouched. To encode object views from real images, the background of input views are randomized, occlusions at various locations are generated, and various lighting and color augmentations are produced. As a result of this *domain randomization*, the network learns to represent the objects' orientation of real object views implicitly using the latent code \mathbf{z} .

Concretely, an input sample $\mathbf{x} \in \mathbb{R}^d$ is randomly augmented by $f(\cdot)$ and mapped by an encoder Υ onto a latent code $\mathbf{z} \in \mathbb{R}^m$ where $m \ll d$. The decoder $\Lambda : \mathbb{R}^m \rightarrow \mathbb{R}^d$ is trained to map the code back to the target \mathbf{x} .

$$\hat{\mathbf{x}} = \Lambda(\Upsilon(f(\mathbf{x}))) = \Lambda(\Upsilon(x')) = \Lambda(\mathbf{z}) \quad (1)$$

Both Υ and Λ are neural networks, and their weight parameters are trained such that the ℓ_2 -loss is minimized, *i.e.*

$$\ell_2(\mathcal{B}) = \sum_{i \in \mathcal{B}} \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|_2 = \sum_{i \in \mathcal{B}} \|\mathbf{x}_i - \Lambda(\Upsilon(f(\mathbf{x}_i)))\|_2 \quad (2)$$

where \mathcal{B} contains the indices of input samples of a given batch. After training, the decoder is discarded and latent encodings of object views from all over $SO(3)$ are saved in a codebook together with their corresponding orientations assigned. At test time, a real object crop is encoded and the nearest code(s) in the codebook according to cosine similarity yield the predicted orientation. The rotation can further be corrected for the translational offset as described in [35].

A downside of this formulation is that a new network must be trained for every new object instance. When naively training the original AAE jointly on several objects, they need to be separated in the latent space so that the decoder is able to reconstruct the correct object. Even when conditioning the decoder on the object by concatenating a one-hot vector to the encoding, it can only reconstruct few instances and it diminishes the ability of the encoder to encode object orientations.

3.2. Multi-Path Encoder-Decoder

We propose a simple but effective architecture which, in combination with our multi-path learning strategy, enables the 3D orientation estimation of multiple objects (see Fig. 1). Our architecture consists of a single encoder Υ , an encoding $\mathbf{z} \in \mathbb{R}^{128}$, and n decoders Λ_j with $j = 1, \dots, n$ where n is the number of different object instances. The convolutional encoder is fed with the same augmented inputs as an AAE but with heterogeneous batches $\tilde{\mathcal{B}}$ containing multiple objects. The resulting codes are split and each decoder only receives codes that correspond to a single object instance. The multi-path loss function can be written as:

$$\begin{aligned} \ell_m(\tilde{\mathcal{B}}) &= \sum_{j=1}^b \sum_{k=1}^n \mathbb{I}(s_j = k) \|\mathbf{x}_j - \Lambda_k(\Upsilon(f(\mathbf{x}_j)))\|_2 \\ &= \sum_{j=1}^b \sum_{k=1}^n \mathbb{I}(s_j = k) \|\mathbf{x}_j - \Lambda_k(\mathbf{z}_j)\|_2 \end{aligned} \quad (3)$$

where \mathbb{I} is the indicator function used to select the decoder Λ_k that corresponds to the instance s_j . Note that in this setting only the encoder Υ receives information from the entire mini-batch, while the decoders Λ_j backpropagate a loss ℓ_j from a sub-batch. Since the decoders are only used to learn an efficient encoding, they can be discarded after training, leaving behind a compact encoder model.

In contrast to other approaches [42, 1], where objects are explicitly separated in the descriptor space, our encoder can learn an interleaved encoding where general features can be shared across multiple instances. We consider this ability as the main qualification to obtain encodings that can represent object orientations from novel, untrained instances, even when they belong to untrained categories.

3.3. Principal Component Analysis of Encodings

In order to gain insights into the characteristics of the latent space from trained and untrained objects, we perform an experiment on the ModelNet40 [43] dataset. We first train the multi-path encoder-decoder on 80 CAD instances that originate from the *car* class. A second model is trained on 10 instances from 8 different classes, namely *airplane*, *bed*, *bench*, *car*, *chair*, *piano*, *sink*, *toilet*. Training details can be found in the appendix.

After training, we generate 72 viewpoints along a full revolution of azimuth, elevation and in-plane rotation. We record different objects from these viewpoints and feed them into the encoder. From the encoding space $z_i \in \mathbb{R}^{128}$ of all objects we compute the first three principal components and project all encodings onto these directions. The interpolated results can be seen in Fig. 2. The top row shows the encodings of a car instance from the training set.

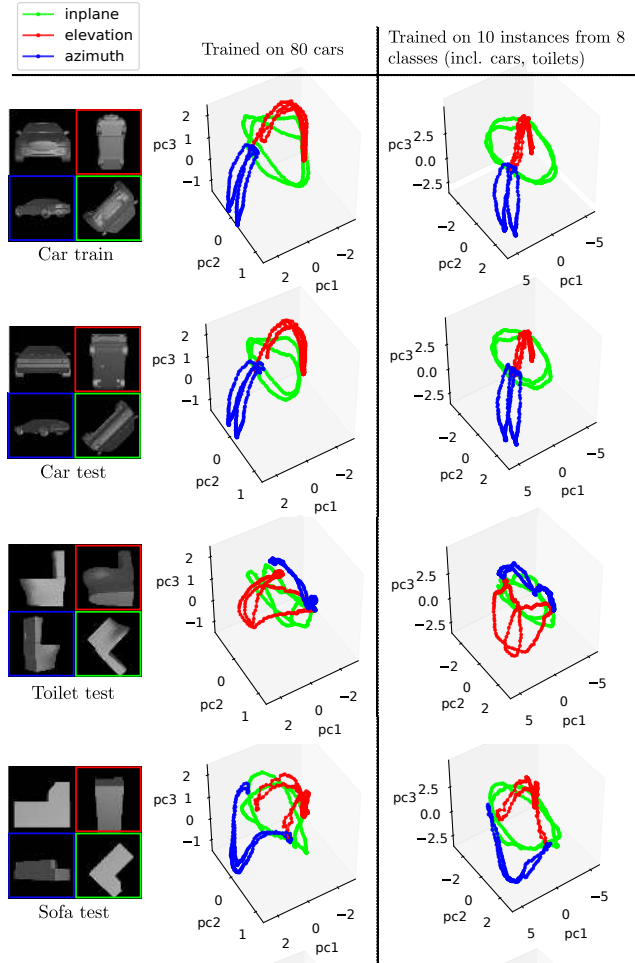


Figure 2: Principal Component Analysis of the learned encodings. Depicted are the encodings of views around elevation (red), azimuth (blue), and in-plane (green). Middle column: Encoder trained only on cars; Right column: Encoder trained on objects from 8 categories

The other rows show instances which are not in the training set of neither model, but different sofa and toilet instances were used to train the second model.

First, it is noticeable that the encodings vary smoothly along each of the rotation axes, even when views of untrained objects are evaluated. It can be seen that the views, while starting at the same point, end up in different sub-manifolds that capture the shape of the encoded objects and their symmetries. For example, the car instance produces closed trajectories with two revolutions because the car shape looks similar at opposite viewpoints.

Furthermore, it can be seen that the car in the training and test set are encoded similarly by each of the models. This indicates that the encoder as well as the codebook could be reused when predicting orientations of novel cars,

Table 1: Different application scenarios for global and iterative pose estimation and the handling of untrained objects

| Prerequisites | Scenario | |
|--|----------|----|
| | I | II |
| Shape category trained | ✗ | ✓ |
| 3D Test Model available | ✓ | ✗ |
| 3D Test Model aligned to Train Model | ✗ | ✗ |
| Available Methods | ↓ | ↓ |
| Reuse codebook from a trained instance | ✗ | ✓ |
| Create a new full / sparse codebook | ✓ | ✗ |
| Iterative relative pose refinement | ✓ | ✗ |

even without access to their 3D models.

The remaining encodings of untrained objects can still be clearly separated and thus a meaningful pose descriptor could be extracted without retraining the encoder but simply by creating a codebook from their 3D model.

Apart from an offset, both models learn quite similar encodings, even for the sofa category which the first model has not seen during training. This implies that low-level features are extracted that generalize across various shapes and categories. However, to fulfill the reconstruction task the features still need to be robust against lighting and color augmentations as well as translations that are applied to the inputs. In the next section, we will explore different approaches to utilize these properties.

3.4. Object Pose Estimation Across Domains

After training, the MP-Encoder can create codebooks for all n training instances for pose estimation. (see Sec. 3.1)

On top of that, it can deal with scenarios including untrained objects which are depicted in Table 1. Here, the trained encoder model is used as a fixed pose-sensitive feature extractor. Available methods depend on the characteristics and given information on the considered test objects.

If a 3D model of the untrained object is available, it is usually preferable to create a new codebook and estimate the 3D orientation from it (I). If no 3D model is available, but we have codebooks of trained instances from a category with coinciding shapes, we could also reuse these codebooks directly for the test instance (II). However, as we do not explicitly align models of a category in a single coordinate frame, the extracted representation will be more dependent on the shape of an object than on any semantics.

Given a 3D model, we can also use our method for iterative pose refinement at test time. This allows us to refine results from sparser codebooks or to perform local, relative pose estimation directly without any codebook.

Algorithm 1: Iterative 6D Pose Refinement

```

Input: encoder  $\Upsilon$ , init pose  $\mathbf{q}_{\text{init}}, \mathbf{t}_{\text{init}}$ , target view  $\mathbf{x}_*$ 
 $\mathbf{z}_* \leftarrow \Upsilon(\mathbf{x}_*)$ 
 $\mathbf{q}_{\text{est}} \leftarrow \mathbf{q}_{\text{init}}$ 
 $\mathbf{t}_{\text{est}} \leftarrow \mathbf{t}_{\text{init}}$ 
for  $k = 0 \dots 2$  do
  for  $j = 0 \dots 3$  do
    for  $i = 0 \dots 40 - 10k$  do
       $\alpha \sim \mathcal{N}(0, \frac{\sigma^2}{j+1})$ 
       $\mathbf{v} \sim \mathcal{N}_3(\mathbf{0}, I)$ 
       $\Delta \mathbf{q} \leftarrow \text{quat}(\frac{\mathbf{v}}{\|\mathbf{v}\|}, \alpha)$ 
       $\mathbf{q}_i \leftarrow \Delta \mathbf{q} \mathbf{q}_{\text{est}}$ 
       $\mathbf{x}_i \leftarrow \text{render}(\mathbf{q}_i, \mathbf{t}_{\text{est}})$ 
       $\mathbf{z}_i \leftarrow \Upsilon(\mathbf{x}_i)$ 
    end
     $k \leftarrow \arg \max_i \frac{\mathbf{z}_i \cdot \mathbf{z}_*}{\|\mathbf{z}_i\| \|\mathbf{z}_*\|}$ 
     $\mathbf{q}_{\text{est}} \leftarrow \mathbf{q}_k$ 
  end
   $\mathbf{x}_{\text{est}} = \text{render}(\mathbf{q}_{\text{est}}, \mathbf{t}_{\text{est}})$ 
   $\Delta \mathbf{t} = \text{multiScaleEdgeMatching}(\mathbf{x}_*, \mathbf{x}_{\text{est}})$ 
   $\mathbf{t}_{\text{est}} = \mathbf{t}_{\text{est}} + \Delta \mathbf{t}$ 
end
Result:  $\mathbf{q}_{\text{est}}, \mathbf{t}_{\text{est}}$ 

```

3.5. Iterative Refinement of Latent Codes

Our method for iterative pose refinement is outlined in Alg. 1. We start with an initial pose estimate and a target view. Next, we render a batch of 3D model views at small rotational perturbations from the initial pose and insert the whole batch into the encoder. The code with the highest cosine similarity to the target view encoding determines the new rotation estimate. We sample again with a smaller perturbation radius. The inner loop of this random optimization scheme, consisting of rendering and inference, can be efficiently parallelized.

Rotation and translation are alternately optimized three times in our experiments. As the MP-Encoder is trained to be invariant against translations, we can first optimize for the more challenging out-of-plane rotation neglecting translation. Afterwards, the rotation is usually roughly aligned to the target and we use a simple edge-based, multi-scale template matching method based on OpenCV [2] to determine the translational offset. The advantage of sampling in SO(3) vs. sampling in the latent space is that (1) the SO(3) space has lower dimensionality and (2) we only search valid orientations.

Apart from relative pose estimation with a novel 3D model without a codebook, the refinement also allows to trade-off between set-up time, inference time, memory resources and accuracy. The full 92232×128 codebook creation takes ~ 5 minutes per object on a modern GPU and 45MB space. Inference only takes ~ 6 ms on a modern GPU while the 6D pose refinement in the presented configura-

Table 2: Pose estimation performance of the MP-Encoder trained on objects 1-18 on the complete T-LESS primesense test set (RGB-only) compared to 30 single object AAEs [34] trained on all objects. We measure recall under the Visible Surface Discrepancy (err_{vsd}) metric. The right column shows the performance of a model trained only on 30 instances of the ModelNet40 dataset. We use ground truth bounding boxes (top) and ground truth masks (bottom) in this experiment. It can be observed that a single MP-Encoder can reach similar performance on unknown objects if segmentation masks are given.

| | Mean | 30 separate AAE encoders [34] | Single Multi-Path Encoder trained on | |
|----------|---------------|-------------------------------|--------------------------------------|-------------------------|
| | VSD recall | | T-Less Objects 1-18 | 30 Instances ModelNet40 |
| +gt bbox | Objects 1-18 | 35.60 | 35.25 | 27.64 |
| | Objects 19-30 | 42.45 | 33.17 | 34.57 |
| | Total | 38.34 | 34.42 | 30.41 |
| +gt mask | Objects 1-18 | 38.98 | 43.17 | 35.61 |
| | Objects 19-30 | 45.33 | 43.33 | 42.59 |
| | Total | 41.52 | 43.24 | 38.40 |

tion takes approximately 1 second. To further speed this up, the random search could be replaced by more sophisticated black-box optimization tools such as CMA-ES [12]. Depending on the application, i.e. number of objects, available resources and constraints on the set-up time, global initialization with sparse codebooks combined with local pose refinement could be a viable option.

3.6. 6DoF Object Detection Pipeline

Our full RGB-based 6DoF object detection pipeline consists of a MaskRCNN with ResNet50 backbone [14], the MP-Encoder for 3D orientation estimation and a projective distance estimate [34] for translation estimation. To compare against other approaches that use depth data, we also report results with a point-to-plane ICP [5, 46] refinement step. Especially in the presence of severe occlusions the RGB-based projective distance estimate does not produce distances accurate enough to fulfill the strict VSD metric. On the other hand, an accurate initialization is crucial for ICP refinement to work well.

For the MaskRCNN we generate training data by pasting object recordings from the T-LESS training set at random translation, scale and in-plane rotation on random background images. Thereby, we produce 80000 images with MS COCO [27] background, 40000 images with black background and 40000 images with random texture backgrounds [6]. We apply several standard online color augmentations like hue, brightness, blur and contrast. Our MaskRCNN reaches an mAP@0.5 performance of 0.68 (bbox detection) and 0.67 (segmentation). Both, MaskRCNN and the MP-Encoder can process multiple objects at once and thus achieve a runtime that stays nearly constant for a large number of objects.

4. Evaluation

We focus our analysis on two benchmarks: ModelNet40 [43] and the T-LESS dataset [19].²

4.1. Metrics

In ModelNet40 we use the absolute angular error

$$e_R = \arccos \left(\text{tr}(\hat{R}^T R - I) / 2 \right) \quad (4)$$

as well as the ADD metric [16] at an average distance threshold of $0.1 \times$ object diameter for the model points \mathcal{M}

$$ADD = \frac{1}{m} \sum_{x \in \mathcal{M}} \|(Rx + t) - (\hat{R}x + \hat{t})\| \quad (5)$$

and the 2D projection metric at an average 5px threshold

$$Proj2D = \frac{1}{m} \sum_{x \in \mathcal{M}} \|K(Rx + t) - K(\hat{R}x + \hat{t})\| \quad (6)$$

In the T-LESS experiments we use the err_{vsd} [20] since here the above metrics are quite meaningless as a result of various object and view symmetries. The err_{vsd} is an ambiguity-invariant pose error metric that is computed from the distance between the estimated and ground truth visible object surfaces. As in the SIXD challenge 2017 [18] and the BOP benchmark 2018 [22], we report the recall of 6D object poses at $err_{vsd} < 0.3$ with tolerance $\tau = 20mm$ and $> 10\%$ object visibility.

4.2. Generalization Capabilities of the MP-Encoder

We first investigate the joint encoding performance and generalization capabilities of the MP-Encoder in isolation, i.e. with ground truth detections and masks. Therefore, we compare one MP-Encoder against 30 separately trained

²Further results can be expected at the BOP challenge 2020 [22]

Table 3: Evaluation of our full 6D Object Detection pipeline with MaskRCNN + Multi-Path Encoder + optional ICP. We report the mean VSD recall following the SIXD challenge / BOP benchmark [21] on the T-LESS Primesense test set. See the appendix for object-wise results. A single MP-encoder model outperforms the result of 30 instance-specific AEs.

| | Template matching | PPF based | | | Learning-based | | | | |
|----------|-------------------|------------------------|-------------------|------------------------------|-----------------------|------------------------|--------------------------|----------------------------|--------------|
| | Hodan-15 Depth | Vidal-18 Depth +ICP | Drost-10 Depth | Drost-10-edge Depth + RGB | Brachmann-16 RGB-D | Kehl-16 RGB-D + ICP | OURS RGB + ICP | Sundermeyer-18 RGB only | OURS |
| Average | 63.18 | 66.51 | 56.81 | 67.5 | 17.84 | 24.6 | 69.53 | 19.26 | 20.53 |
| Time (s) | 13.5 | 4.7 | 2.3 | 21.5 | 4.4 | 1.8 | 0.8 | 0.1 | 0.2 |

AAE models on all scenes of the T-LESS Primesense test set (Table 2). Equivalent encoder and decoder architectures are used. Furthermore, the MP-Encoder is only trained on the first 18 3D object reconstructions of the T-LESS dataset to show the generalization capabilities on untrained objects 19-30. On objects 1-18 which the MP-Encoder has seen during training, the results using ground truth bounding boxes are close to the AAE results. Looking at the next row, the performance on the untrained objects 19-30 is significantly worse compared to the single AAEs. We suppose that the MP-Encoder has difficulties to extract unknown objects from the background. This hypothesis is strongly supported by the results with ground truth masks (bottom) where the MP-Encoder even outperforms the AAEs. Even for the untrained objects 19-30 the performance gap to the AAEs that were trained on these objects is quite small. One possible explanation is that a feature extractor trained on multiple objects learns more versatile and thus more robust features.

The last column of Table 2 depicts the surprisingly good generalization from ModelNet40 to the T-LESS dataset. Here, the MP-Encoder is specifically trained on 30 texture-free CAD models from the 8 categories *airplane*, *bed*, *bench*, *car*, *chair*, *piano*, *sink*, *toilet*. Codebooks are created for all T-LESS objects and it is tested on the same real sensor recordings. These results underline that with the multi-path training combined with an input randomization strategy, we can learn to extract orientation variant features that generalize well across greatly differing domains.

4.3. 6D Object Detection results

Next, we evaluate our full 6D pose estimation pipeline on the T-LESS dataset which is a particularly challenging 6D object detection benchmark containing texture-less, symmetric objects as well as clutter and severe occlusions. Table 3 presents results using the strict vsd metric (Sec. 3.6). We achieve state-of-the-art results on T-LESS at much lower runtimes than previous methods, both in the RGB domain and also in the depth domain when initializing an ICP with our RGB-based pose estimate. Although the gains are marginal, the results are significant because we only train a single encoder on the whole dataset and still obtain state-of-the-art results. This makes our method scalable, partic-



Figure 3: Left: Qualitative 6D Object Detection results of the RGB-based pipeline on T-LESS Primesense scenes. Only 18 of 30 objects are used to train the Multi-Path encoder; Middle: ICP-refined results; Right: Relative pose refinement of instances from the untrained categories guitar and bathtub. Red is the initial pose, green depicts the refined pose.




Figure 4: Failure case: MaskRCNN predicts the wrong class 20 instead of 19 (the object below). Since the shape is quite similar (except scale) the codebook of object 20 still gives a reasonable pose estimate.

ularly considering that no real pose-annotated data is used for training. Fig. 3 (left) shows qualitative examples of the full 6D Object Detection pipeline. Previously, pose estimation of texture-less objects has been dominated by pure depth-based geometric matching approaches that often possess high run-times and do not scale very well with the number of considered objects. Fig. 4 shows a failure case which underlines that strict instance-wise pose estimation may not be suitable for real world applications where objects often

Table 4: Relative Pose Refinement from up to 45° and $\Delta t = (10, 10, 50)[mm]$ perturbations on untrained instances of seen (top) and unseen (bottom) object categories of the ModelNet40 dataset. We measure the recall at the $5cm, 5^\circ$ threshold, the ADD at 0.1d (object diameter) metric and the Proj2D at 5px threshold.

| metric | $(5^\circ, 5cm)$ | | | | 6D Pose (ADD) | | | | Proj2D (5px) | | | | |
|----------------|------------------|---------|------|---------|---------------|---------|-------------|---------|--------------|---------|-------------|---------|-------------|
| | DeepIm [25] | | Ours | | DeepIm [25] | | Ours | | DeepIm [25] | | Ours | | |
| method | init | refined | init | refined | init | refined | init | refined | init | refined | init | refined | |
| novel instance | airplane | 0.8 | 68.9 | 0.9 | 96.9 | 25.7 | 94.7 | 33.4 | 97.9 | 0.4 | 87.3 | 0.1 | 97.4 |
| | car | 1.0 | 81.5 | 0.4 | 96.4 | 10.8 | 90.7 | 13.4 | 98.5 | 0.2 | 83.9 | 0.1 | 94.0 |
| | chair | 1.0 | 87.6 | 0.3 | 96.4 | 14.6 | 97.4 | 16.3 | 98.3 | 1.5 | 88.6 | 0.0 | 94.6 |
| | Mean | 0.9 | 79.3 | 0.5 | 96.6 | 17.0 | 94.3 | 21.0 | 98.2 | 0.7 | 86.6 | 0.1 | 95.3 |
| novel category | bathtub | 0.9 | 71.6 | 0.7 | 85.5 | 11.9 | 88.6 | 15.4 | 91.5 | 0.2 | 73.4 | 0.1 | 80.6 |
| | bookshelf | 1.2 | 39.2 | 0.7 | 81.9 | 9.2 | 76.4 | 13.7 | 85.1 | 0.1 | 51.3 | 0.0 | 76.3 |
| | guitar | 1.2 | 50.4 | 0.5 | 69.2 | 9.6 | 69.6 | 13.1 | 80.5 | 0.2 | 77.1 | 0.3 | 80.1 |
| | range_hood | 1.0 | 69.8 | 0.5 | 91.0 | 11.2 | 89.6 | 14.1 | 95.0 | 0.0 | 70.6 | 0.0 | 83.9 |
| | sofa | 1.2 | 82.7 | 0.6 | 91.3 | 9.0 | 89.5 | 12.2 | 95.8 | 0.1 | 94.2 | 0.0 | 86.5 |
| | wardrobe | 1.4 | 62.7 | 0.7 | 88.7 | 12.5 | 79.4 | 14.8 | 92.1 | 0.2 | 70.0 | 0.0 | 81.1 |
| | tv_stand | 1.2 | 73.6 | 0.6 | 85.9 | 8.8 | 92.1 | 10.5 | 90.9 | 0.2 | 76.6 | 0.1 | 82.5 |
| | Mean | 1.2 | 64.3 | 0.6 | 84.8 | 10.3 | 83.6 | 13.4 | 90.1 | 0.1 | 73.3 | 0.1 | 81.6 |

differ only in details.

4.4. Iterative Refinement on Untrained Objects

In our final experiment, we evaluate the MP-Encoder on the task of iterative refinement of poses on untrained instances from seen and unseen categories from ModelNet40.

We follow the evaluation protocol of DeepIm [25] where the relative pose between two object views of an untrained instance is sought. The target view is rendered at constant translation $t = (0, 0, 500)$ (mm) and random rotation $R \sim SO(3)$. Then we render another object view with that pose perturbed by the angles $\beta_{x/y/z} \sim \mathcal{N}(0, (15^\circ)^2)$ sampled for each rotation axis and a translational offset $\Delta x \sim \mathcal{N}(0, 10^2)$, $\Delta y \sim \mathcal{N}(0, 10^2)$, $\Delta z \sim \mathcal{N}(0, 50^2)$ (mm). If the total angular perturbation is more than 45° , it is resampled.

We train category specific MP-Encoders on 80 instances of the airplane, car and chair class and predict the relative pose on novel instances. We also train another MP-Encoder on 80 instances from 8 different categories³ to obtain a general viewpoint-sensitive feature extractor which we test on instances from novel categories. To retain the relative pose to the target view we use our random optimization scheme described in Alg. 1 with an initial $\sigma = 45^\circ$.

We compare our approach against DeepIm [25] that also minimizes the relative pose between an initial and a target view through an iterative rendering inference cycle. The results in Table 4 demonstrate superior performance of our approach on both, seen and unseen categories. Figure 3 on

the right shows qualitative results of the refinement on unseen categories.

5. Conclusion

In this paper we presented a novel approach for estimating poses of multiple trained and untrained objects from a single encoder model. In contrast to other methods, training on multiple objects does not degrade performance and instead leads to state-of-the-art results on the textureless, symmetric objects of the T-LESS dataset. The same MP-Encoder architecture is also used for iterative pose refinement on untrained objects outperforming previous approaches on ModelNet40.

The ability of this pose estimator to generalize across objects from different categories, datasets and image domains indicates that low-level viewpoint-sensitive features are shared across various domains. Higher-level features for discriminating semantics require viewpoint-invariance and usually generalize less well. Therefore, our results suggest that these two tasks should be learned separately.

We believe that this is a step towards coping with the large number of instances in industrial settings where 3D models are often available and service robotics where categories often appear repeatedly and interactions with novel objects should be immediately possible.

³airplane, bed, bench, car, chair, piano, sink, toilet

References

- [1] Vassileios Balntas, Andreas Doumanoglou, Caner Sahin, Juil Sock, Rigas Kouskouridas, and Tae-Kyun Kim. Pose Guided RGB-D Feature Learning for 3D Object Pose Estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3856–3864, 2017. 2, 3, 4
- [2] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000. 5
- [3] Esteves Carlos, Sud Avneesh, Luo Zhengyi, Daniilidis Kostas, and Makadia Ameesh. Cross-domain 3d equivariant image embeddings. *arXiv preprint arXiv:1812.02716*, 2018. 3
- [4] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 801–818, 2018.
- [5] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Image and vision computing*, 10(3):145–155, 1992. 6
- [6] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014. 6
- [7] Xinke Deng, Arsalan Mousavian, Yu Xiang, Fei Xia, Timothy Bretl, and Dieter Fox. Poserbpf: A rao-blackwellized particle filter for 6d object pose tracking. *Robotics: Science and Systems (RSS)*, 2019. 2
- [8] Maximilian Denninger, Martin Sundermeyer, Dominik Winkelbauer, Youssef Zidan, Dmitry Olefir, Mohamad Elbadrawy, Ahsan Lodhi, and Harinandan Katam. Blenderproc. *arXiv preprint arXiv:1911.01911*, 2019. 2
- [9] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning*, pages 647–655, 2014. 3
- [10] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 2
- [11] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>, 2012.
- [12] Nikolaus Hansen, Youhei Akimoto, and Petr Baudis. CMA-ES/pycma on Github. Zenodo, DOI:10.5281/zenodo.2559634, Feb. 2019. 6
- [13] Jiedong Hao, Jing Dong, Wei Wang, and Tieniu Tan. What is the best practice for cnns applied to visual instance retrieval? *arXiv preprint arXiv:1611.01640*, 2016. 3
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. *arXiv preprint arXiv:1703.06870*, 2017. 6
- [15] Stefan Hinterstoisser, Cedric Cagniard, Slobodan Ilic, Peter Sturm, Nassir Navab, Pascal Fua, and Vincent Lepetit. Gradient response maps for real-time detection of textureless objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(5):876–888, 2012. 2
- [16] Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary Bradski, Kurt Konolige, and Nassir Navab. Model based training, detection and pose estimation of texture-less 3D objects in heavily cluttered scenes. In *Asian conference on computer vision*, pages 548–562. Springer, 2012. 6
- [17] Stefan Hinterstoisser, Vincent Lepetit, Naresh Rajkumar, and Kurt Konolige. Going further with Point Pair Features. In *European Conference on Computer Vision*, pages 834–848. Springer, 2016. 2
- [18] Tomas Hodan. SIXD Challenge 2017, http://cmp.felk.cvut.cz/sixd/challenge_2017/, 2017. 6
- [19] Tomáš Hodaň, Pavel Haluza, Štěpán Obdržálek, Jiří Matas, Manolis Lourakis, and Xenophon Zabulis. T-LESS: An RGB-D Dataset for 6D Pose Estimation of Texture-less Objects. *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2017. 6
- [20] Tomáš Hodaň, Jiří Matas, and Štěpán Obdržálek. On evaluation of 6D object pose estimation. In *European Conference on Computer Vision*, pages 606–619. Springer, 2016. 6
- [21] Tomas Hodan, Frank Michel, Eric Brachmann, Wadim Kehl, Anders GlentBuch, Dirk Kraft, Bertram Drost, Joel Vidal, Stephan Ihrke, Xenophon Zabulis, et al. Bop: benchmark for 6d object pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 19–34, 2018. 2, 7
- [22] Tomas Hodan, Frank Michel, Eric Brachmann, Wadim Kehl, Anders GlentBuch, Dirk Kraft, Bertram Drost, Joel Vidal, Stephan Ihrke, Xenophon Zabulis, Caner Sahin, Fabian Manhardt, Federico Tombari, Tae-Kyun Kim, Jiri Matas, and Carsten Rother. Bop: Benchmark for 6d object pose estimation. In *The European Conference on Computer Vision (ECCV)*, September 2018. 1, 6
- [23] 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 Conference on Computer Vision and Pattern Recognition*, pages 1521–1529, 2017. 1, 2
- [24] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [25] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. Deepim: Deep iterative matching for 6d pose estimation. *arXiv preprint arXiv:1804.00175*, 2018. 2, 3, 8
- [26] Zhigang Li, Gu Wang, and Xiangyang Ji. Cdpn: Coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2
- [27] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2, 6

- [28] Siddharth Mahendran, Haider Ali, and René Vidal. 3d pose regression using convolutional neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2174–2182, 2017. [2](#)
- [29] Fabian Manhardt, Wadim Kehl, Nassir Navab, and Federico Tombari. Deep model-based 6d pose refinement in rgb. In *The European Conference on Computer Vision (ECCV)*, September 2018. [3](#)
- [30] Pat Marion, Peter R Florence, Lucas Manuelli, and Russ Tedrake. Label fusion: A pipeline for generating ground truth labels for real rgbd data of cluttered scenes. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3325–3242. IEEE, 2018. [2](#)
- [31] Kiru Park, Timothy Patten, and Markus Vincze. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. [2](#)
- [32] 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 *IEEE International Conference on Computer Vision (ICCV)*, 2017. [2](#)
- [33] Caner Sahin and Tae-Kyun Kim. Category-level 6d object pose recovery in depth images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018. [3](#)
- [34] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3d orientation learning for 6d object detection from rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 699–715, 2018. [1](#), [2](#), [3](#), [6](#)
- [35] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, and Rudolph Triebel. Augmented autoencoders: Implicit 3d orientation learning for 6d object detection. *International Journal of Computer Vision*, Oct 2019. [1](#), [3](#)
- [36] Supasorn Suwajanakorn, Noah Snavely, Jonathan J Tompson, and Mohammad Norouzi. Discovery of latent 3d keypoints via end-to-end geometric reasoning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 2059–2070. Curran Associates, Inc., 2018. [3](#)
- [37] Bugra Tekin, Sudipta N. Sinha, and Pascal Fua. Real-Time Seamless Single Shot 6D Object Pose Prediction. In *CVPR*, 2018. [2](#), [3](#)
- [38] Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. In *Conference on Robot Learning*, pages 306–316, 2018. [1](#), [2](#), [3](#)
- [39] Markus Ulrich, Christian Wiedemann, and Carsten Steger. CAD-based recognition of 3D objects in monocular images. In *ICRA*, volume 9, pages 1191–1198, 2009. [2](#)
- [40] Joel Vidal, Chyi-Yeu Lin, and Robert Martí. 6D Pose Estimation using an Improved Method based on Point Pair Features. *arXiv preprint arXiv:1802.08516*, 2018. [2](#)
- [41] Chen Wang, Danfei Xu, Yuke Zhu, Roberto Martín-Martín, Cewu Lu, Li Fei-Fei, and Silvio Savarese. Densefusion: 6d object pose estimation by iterative dense fusion. *CoRR*, abs/1901.04780, 2019. [1](#)
- [42] Paul Wohlhart and Vincent Lepetit. Learning descriptors for object recognition and 3D pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3109–3118, 2015. [2](#), [3](#), [4](#)
- [43] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. [4](#), [6](#)
- [44] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017. [2](#)
- [45] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Dpod: Dense 6d pose object detector in rgb images. *arXiv preprint arXiv:1902.11020*, 2019. [1](#), [2](#)
- [46] Zhengyou Zhang. Iterative point matching for registration of free-form curves and surfaces. *International journal of computer vision*, 13(2):119–152, 1994. [6](#)