

ZePhyR: Zero-shot Pose Hypothesis Rating

Brian Okorn*, Qiao Gu*, Martial Hebert and David Held

Abstract—Pose estimation is a basic module in many robot manipulation pipelines. Estimating the pose of objects in the environment can be useful for grasping, motion planning, or manipulation. However, current state-of-the-art methods for pose estimation either rely on large annotated training sets or simulated data. Further, the long training times for these methods prohibit quick interaction with novel objects. To address these issues, we introduce a novel method for zero-shot object pose estimation in clutter. Our approach uses a hypothesis generation and scoring framework, with a focus on learning a scoring function that generalizes to objects not used for training. We achieve zero-shot generalization by rating hypotheses as a function of unordered point differences. We evaluate our method on challenging datasets with both textured and untextured objects in cluttered scenes and demonstrate that our method significantly outperforms previous methods on this task. We also demonstrate how our system can be used by quickly scanning and building a model of a novel object, which can immediately be used by our method for pose estimation. Our work allows users to estimate the pose of novel objects without requiring any retraining. Additional information can be found on our website <https://bokorn.github.io/zephyr/>

I. INTRODUCTION

6D pose describes the position and orientation of an object, defined in a reference frame relative to a predefined model of the object. An object’s 6D pose fully describes the state of a static rigid object and, as such, is commonly used as a representation for planning [1], [2]. A robot can use an estimate of an object’s pose to perform complex manipulation interactions with the object [3], [4], [5], [6].

Current state-of-the-art methods for object pose estimation train a new model for each object they are being evaluated on [7], [8], [9]. This requires a large amount of annotated training data, either produced by capturing and annotating large datasets or through rendering the object in synthetically generated scenes. For example, the YCB-Video dataset [7] contains 133,827 human-annotated images with roughly 25,000 images per object. Although this dataset has enabled the training of powerful deep learning methods [7], [8], curating such a human-labeled dataset (including both capturing a diverse dataset and labeling the data) for each new object that a robot must interact with is cumbersome. Methods that rely on purely simulated data [10], [11], [12] avoid this limitation but must instead contend with the sim2real gap between the synthetic data and real sensor observations. Improved rendering [13] and domain randomization techniques [14] have been suggested to alleviate this gap, but ensuring




Fig. 1: Pose hypotheses scored using Zero-shot Pose Hypothesis Rating on novel drill object, reconstructed at test time. The highest scoring pose is rendered in color. Poses are outlined in color corresponding to score, with highly rated poses in red transitioning to lower ones in blue.

that the simulated data accurately represents the variations observed in the real world continues to be an open problem.

Regardless of how this data is obtained, training new networks has a time and space cost. This training can take many hours, which prevents robots using such systems from quickly being able to interact with new objects. Additionally, new network weights are trained for each new object, which presents a difficulty for memory-constrained robot systems. These constraints do not scale well in cases where robots need to interact with many different types of objects.

One approach to mitigate these issues is to use a non-learned geometry-based method [15], [16]. These methods, however, do not typically capture visual texture well, and they rely on hard-coded, rather than learned, invariances, which limits the potential accuracy of the system (based on our experiments in Section V-A). A few recent learning-based approaches have attempted to perform zero-shot object pose estimation [17], [18] but these methods require instance segmentation masks to be provided as input, which limits their use in a “zero-shot” system, as such masks are typically trained per-object.

We seek to remove these limitations by developing a novel learning-based method for zero-shot object pose estimation that can handle both textured and untextured objects in cluttered scenes and does not require object masks as input. Our method uses the paradigm of pose hypothesis generation and evaluation: given a scene, a large number of candidate poses are generated that are consistent with the observation. The fitness of each hypothesis is then evaluated and the best-fit candidate is selected. Such an approach requires

* indicates equal contribution.

B. Okorn, Q. Gu, M. Hebert and D. Held are with the Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213. (bokorn@andrew.cmu.edu, qiaog@andrew.cmu.edu, hebert@cs.cmu.edu, dheld@andrew.cmu.edu)

the hypothesis rating function to give appropriate weight to the features that most correlate with the correct pose. The variation between sensor data and the object model, caused by sensor noise or lighting changes, as well as partial occlusions, can make designing this scoring function non-trivial. Past approaches to hypothesis scoring have used voting over hypotheses or feature matching [19], [20], [15]; in contrast, this paper proposes a scoring function that learns to compare the observed images and rendered model points. Our learned scoring function demonstrates a significant improvement on zero-shot object pose estimation over a wide set of objects and environmental variations.

The key insight of our method is to use a learned scoring function that compares the sensor observation to a sparse rendering of each candidate pose hypothesis. This scoring function receives as input *an unordered set of point differences*, shown in Fig. 2, which we show is crucial to perform zero-shot generalization to novel objects not seen in the training set. Our method is trained over a disparate set of objects and then evaluated on novel objects not included in the training set.

We demonstrate that our method works on objects in clutter without requiring object masks as input, unlike past zero-shot methods [18], [17]. Our method handles both untextured objects as well as objects with significant visual texture, not seen at training time. Our method thus achieves the goal of zero-shot object pose estimation mentioned earlier:

- We require no new human annotations or large scale synthetic data generation to interact with novel objects.
- We require no retraining for novel objects.
- Our method uses only a single set of network weights, rather than requiring new weights for each unique object, reducing the memory constraints.

We evaluate our method on YCB-Video and Occluded-Linmod, two challenging pose estimation datasets. Our method reduces the pose error by a large margin and achieves state-of-the-art results over previous zero-shot pose estimation methods.

II. RELATED WORK

A. Non-learned Zero-shot Pose Estimation

Zero-shot pose estimation is the task of estimating the pose of objects not seen at training time. Non-learning based approaches are inherently zero-shot, leveraging robust features and the available object model at test time. Point Pair Features (PPF) [16], [15], [21], [22], [21], [23] use pairs of oriented points to generate geometrically consistent pose hypotheses and select the best hypothesis using voting and clustering. These are the top-performing zero-shot methods on the BOP leader board [24], when averaged over all datasets, but struggle to compete with deep learned methods on the highly textured YCB dataset due to the methods being exclusively based on depth.

B. Learned Zero-shot Object Pose Estimation

Several learned methods solve the zero-shot pose estimation problem using class-based pose estimation [25], [26] as

opposed to instance-based pose estimation. These methods learn a pose estimator capable of generalizing among objects in the given class, but such methods are not intended to generalize to novel classes. While this is a step in the direction of zero-shot pose estimation, it still requires training a new network for each class.

Pose refinement methods like DeepIM [27] learn to estimate the residual pose between the observed data and a rendered viewpoint and have shown to generalize well to unseen classes of objects. These methods, however, require the initial rendered pose to be relatively close to the observation to produce accurate results, and as such is primarily used to refine a coarse pose prediction. Our method requires no such close initialization.

A few recent zero-shot methods use a learned representation of the object in their pose estimation pipeline [28], [18], [17]. While these methods have been shown to generalize across objects, they require a bounding box for the target object, which is obtained using an object-specific learned detector (and hence not a zero-shot system) or the ground-truth bounding box. This requirement is avoided in the MOPED dataset [17], as there is only a single object in the scene, which greatly simplifies the task of estimating the object mask [29]. For the occluded LineMOD dataset, ground truth object masks are used [17]. Our method does not require such bounding boxes or masks as input, making it truly zero-shot.

C. Pose Scoring

There has been some study of learned fitness functions. Differentiable RANSAC (DSAC) [20] explores learning a fully differentiable RANSAC algorithm. Specifically, they study the use of a REINFORCE style loss for scoring candidate hypotheses. We take inspiration from this work; however, their method focuses on a different task of camera localization rather than object pose estimation; as a result, many important details of our method, such as the input featurization and network architecture, are significantly different from their approach. Pose Proposal Critic [30] learns to regress to the reprojection error between a rendered pose and the observation. They numerically differentiate this error function as a means of pose refinement. However, they only evaluate this approach as a pose refinement technique, with a close initial pose estimate; in contrast, our focus is on evaluating a large set of pose hypotheses that span the entire observation space.

III. METHOD

A. Overview

The primary objective of this work is zero-shot object pose estimation in clutter. To achieve this, we train our pose estimation method on one set of objects and then evaluate on a set of novel objects, without requiring any re-training. This differentiates our work from previous work that requires real or synthetic training data of the test objects [7], [8]. Our work additionally differentiates from other zero-shot pose estimation work [17], [18], [16], in that it operates well

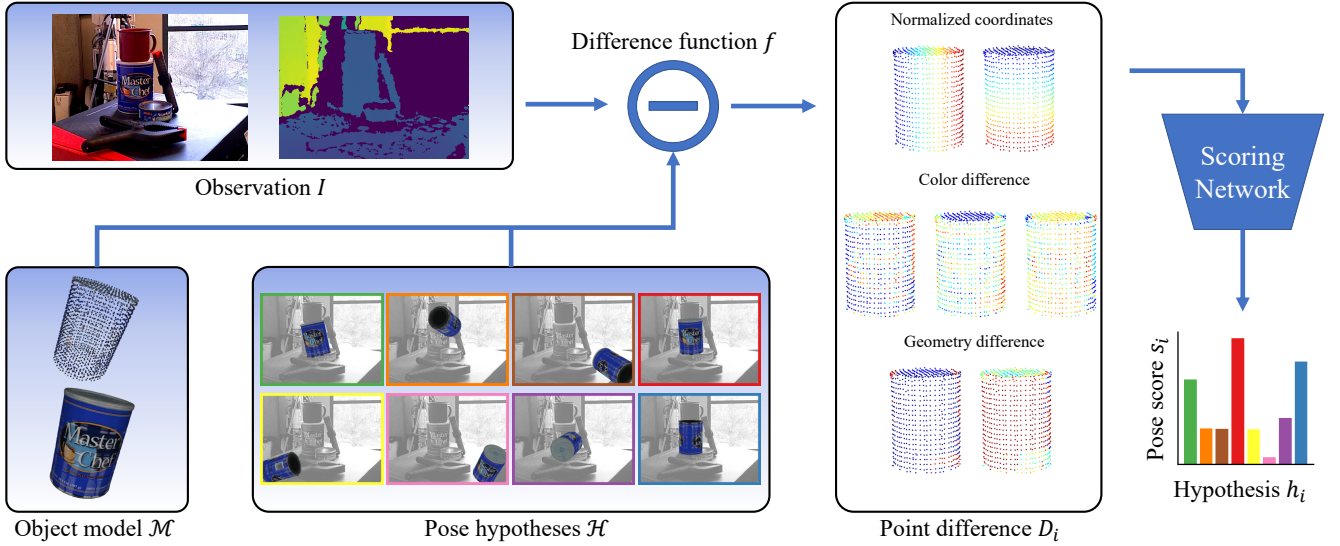


Fig. 2: System Pipeline. Our method first projects the sampled model points \mathcal{M} onto the observation I according to a pose hypothesis h_i . Then D_i are extracted as the point-wise differences between the observation and the projected model points, describing the alignment of the pose hypothesis at each projected point. A network takes in D_i and regresses to a score s_i for each pose h_i which evaluates how well the pose matches the observation.

in cluttered scenes, requires no object masks as input, and produces accurate poses for both textured and untextured objects.

An overview of our method is shown in Figure 2. Given a set of 6D pose hypotheses, we first project each hypothesis into the scene. Our method learns to score each hypothesis by comparing differences in the projected object model point cloud to the RGB-D observation. For each projected model point, we extract the color and geometry information from both the model and the observation and compute the local differences of the extracted information. This yields a set of *point-differences*, one for each projected model point. Each element in this set encodes the local alignment between the model and the observation with respect to color and geometry. We adopt a point-based network [31], [32] to analyze this unordered set of point-differences and regress to an overall score for each pose hypothesis. Focusing on differences as well as adopting a point-based neighborhood structure helps us avoid overfitting to object-specific properties from the training set and allows us to generalize to unseen objects at test time.

In this work, our primary focus is the learned scoring function and we use existing methods to generate our initial pose hypothesis set. While many algorithms could be used to generate these potential object pose hypotheses [20], [33], [34], we use a combination of Point Pair Features [15] and SIFT features [35].

B. Learned Scoring Function

The main goal of our method is to score pose hypotheses by projecting them into the observed scene and learning to compare their local geometric and color differences. Suppose that we have a set of 6D pose hypotheses $\mathcal{H} = \{h_i\}_{i=1}^m$ that we wish to evaluate. We represent the object as a point cloud $\mathcal{M} = \{x_j\}_{j=1}^n$, sampled from the provided object

mesh model, or obtained from a 3D reconstruction pipeline. Each point contains both geometric (depth and normal) and color information drawn from its local region on the object. Similarly the observation image I contains geometric and color values from the observation. To evaluate hypothesis h_i , we project each object point x_j onto the observation’s image plane, using the known camera parameters. This projection gives a point at image coordinates y_{ij} with transformed point values \tilde{x}_{ij} (the point depth and normal vector are transformed; the color of the projected point is unchanged). For each pose hypothesis, the difference between the projected values, \tilde{x}_{ij} , and their corresponding image values, $I(y_{ij})$, is computed according to a simple distance function, $d_{ij} = f(\tilde{x}_{ij}, I(y_{ij}))$ (see Section III-D.2 for details).

The set $D_i = \{d_{ij}\}_{j=1}^m$ represents an unordered set of point differences for pose hypothesis h_i , each of which is associated with a given point x_j in the model and a location y_{ij} in the observation image. We train a deep neural network $g_\theta(D_i)$ with parameters θ to analyze this difference set and regress to a pose fitness score, s_i . While one might assume that a simple hand-designed function for g would be sufficient, in practice, however, occlusions, lighting differences and other confounding factors make such simple methods ineffective. Our learned function is able to intelligently combine point differences on multiple parts of the object to robustly estimate the most likely pose hypothesis.

C. Loss Function

To train this hypothesis scoring function, we adopt the probabilistic selection loss proposed by DSAC [20], as it directly optimizes the expected pose error when hypotheses are sampled according to the predicted scores. For each pose hypotheses h_i with corresponding true pose error ϵ_i , we compute the expected pose error of sampling ac-

cording to the softmax distribution induced by s_i , $\mathcal{L} = \sum_{i=1}^m \text{softmax}(s_i) \epsilon_i$.

In our experiment, ϵ_i is defined as the log of the average point distance metric (ADD) for non-symmetric objects and its symmetric analog (ADD-S) for symmetric objects [7]. Empirically, we find that using the log of this error better dampens the effects of outliers. More discussion can be found in Section III-D.2. At test time, the highest-scoring pose hypothesis is selected. The full test time pipeline is described in Algorithm 1.

Algorithm 1: Hypothesis Scoring Pose Estimation

```

Compute initial pose hypothesis set  $\mathcal{H} = \{h_i\}_{i=1}^m$ ;
foreach  $h_i$  in  $\mathcal{H}$  do
    Project all model points according to  $h_i$  onto the
    image plane to get projected model points  $\tilde{x}_{ij}$  at
    projected image coordinates  $y_{ij}$ ;
    Get observation points  $I(y_{ij})$ ;
    Compute point differences  $d_i = f(\tilde{x}_{ij}, I(y_{ij}))$ ;
    Score point-differences  $s_i = g_\theta(\{d_{ij}\}_{j=1}^m)$ ;
end
Return hypothesis  $h_{i^*}$ , where  $i^* = \arg \max_i s_i$ ;

```

D. Implementation details

1) *Hypothesis Generation*: We generate the initial hypotheses set using the commercially available Point Pair Feature software, HALCON 20.05 Progress software [36], which implements the PPF algorithm described in Drost *et al.* [15]. For each observation, we use the top 100 pose hypotheses generated by PPF. For detecting objects with high visual texture (e.g. for all objects in YCB-V), we augment these hypotheses using Dense SIFT feature matching. We obtain pose hypotheses from these features by aligning the surface normals and SIFT orientations of pairs of matched SIFT features; aligning the SIFT orientations and normals enables a single pair of matched SIFT features to define a 6D pose hypothesis.

2) *Network Input*: As input to the hypothesis scoring function, we use very simple geometric and color information for both the model and observation data. For each point on the model, we compute its 3D location, surface normal, and color in HSV space. When projecting each point into the observation frame, we transform both the normals and 3D coordinates to compute the depth and normal with respect to the camera. The color data is unaffected by the projection. Similarly, we compute local surface normals from the observation, and thus we obtain depth, normal and HSV color information at each pixel of the observation image.

To create the network inputs d_{ij} , we compute the signed difference between the projected and observed points for both depth and color. For surface normals, we use the cosine of the angle difference between the projected and observed normals. Additionally, we concatenate the projected image coordinates, y_{ij} , of the associated image point, normalized to zero mean and unit variance, as an additional input to the network, to provide some structural neighborhood information.

3) *Network Structure*: Our network takes in the set of point-differences $D_i = \{d_{ij}\}_{j=1}^m$ and outputs a single score, s_i , that estimates how well the pose hypothesis h_i matches the observation. Because D_i is an unordered set of point-differences, we use a network architecture designed to handle unordered sets of points; specifically, we use PointNet++ [32]. Our experiments show that the loose neighborhood structure of this architecture enables zero-shot generalization to unseen objects.

To define the spatial neighborhood for grouping points in PointNet++’s point set abstraction layers, we use the normalized image coordinates. We explore the effect of networks with different neighborhood structures in Section V-C. See the supplementary material on our website for more experimental details and hyperparameters.

IV. EXPERIMENTS

A. Datasets

We evaluated our method on two of the most popular datasets in the BOP Challenge [24], the YCB-Video (YCB-V) dataset [7] and the LineMOD-Occlusion (LM-O) dataset [9]. In these experiments, we follow the evaluation protocol set up by the BOP Challenge, with the additional constraint that our method is not trained on the objects it is tested on. This allows us to test our ability to perform zero-shot generalization to novel objects.

YCB-Video dataset (YCB-V) [7] contains 92 RGB-D video sequences of 21 YCB objects [39] of varying shape and texture, annotated with 6D poses. This is a particularly challenging dataset for object pose estimation due to its varying lighting conditions, occlusions, and sensor noise. We follow the dataset split in [7], and for the evaluation, we adopt the BOP testing set [24], where 75 images with higher-quality ground-truth poses from each of 12 test videos are used. To demonstrate the generalization ability of our method, one half of the objects are used for training, and the other half are used for testing. To accommodate the full dataset, a second network is trained with train and test objects exchanged, such that each network only sees half of the objects during training, and no network is trained on the objects that it will be tested on. Note that we train our network on the training (seen) objects in the YCB-V training split and test on the testing (unseen) objects in the testing split, so not a single test image or object is seen during training. When evaluating on YCB-V, we use hypotheses generated from both PPF and SIFT matching to handle the high degree of visual texture. We also adopt a ICP refinement step [40] for post-processing.

LineMOD-Occlusion dataset (LM-O) [9] adopted a single scene from the test set of the larger LineMOD (LM) dataset [41] and provides ground-truth 6D pose annotations for 8 low-textured objects. For training, we used the PBR-BlenderProc4BOP [42] training images provided by the BOP challenge. This dataset contains photo-realistic synthetic images of LM objects dropped onto a table, with randomized background texture and object materials. Our model is only trained on synthetic images of the 7 objects

	Zero-Shot Methods				Object Specific Methods	
	Drost [15]	Vidal [16]	Multipath [18]	ZePHYR + Drost (Ours)	CosyPose [37]	Pix2Pose [38]
YCB-V	0.344	0.450	0.289	0.516	0.861	0.675
LM-O	0.527	0.581	0.217	0.598	0.714	0.588

TABLE I: AR scores for methods of zero-shot and object specific pose estimation on object pose datasets (higher is better).

that are in the LM dataset but *not* in the LM-O dataset; we then evaluate on the LM-O objects, which were not seen at training time. When evaluating on LM-O, we only use hypotheses generated by PPF; we find that SIFT hypotheses are ineffective on this dataset since the objects do not contain much visual texture.

B. Metrics

As suggested by the BOP challenge, we report the average recall (AR) scores as the average of the following three average-recall pose error metrics: Visible Surface Discrepancy (VSD), Maximum Symmetry-Aware Surface Distance (MSSD), and Maximum Symmetry-Aware Projection Distance (MSPD). For a detailed formulation of each metric, please refer to [24].

C. Baselines

We compare our method to both zero-shot and object-specific pose estimation methods. As we are most concerned with our performance as compared to other zero-shot methods, we compare to two variants of Point Pair Features, Drost [15] and Vidal [16]. An implementation of Drost’s PPF is used as the hypothesis generation algorithm in our work. Vidal had until recently been the top-performing method in the BOP challenge, and demonstrates the peak performance of PPF-only systems (although their code is not available). Other recent papers have proposed learning-based methods for zero-shot pose estimation, namely Multipath Augmented Autoencoders [18], which we compare against. While this method has been shown to generalize to unseen objects, the reported results that we include are a product of training a single model on the test objects; further, their method utilizes an object-specific detection network (also trained on the test objects) [43]. In addition to the zero-shot baselines, we report the current state of the art in object-specific methods as CosyPose [37] and Pix2Pose [38]. Both of these methods train a network on annotated instances of the test objects and have weights specifically associated with each object. While we are not attempting to match the performance of these systems, we report their results to illustrate the still remaining gap between zero-shot and object-specific methods.

V. RESULTS

A. Zero-shot Pose Estimation

In Table I, we find that our method outperforms all zero-shot methods, significantly improving over our initial pose hypotheses produced by Drost and outperforming the best PPF-only solution in Vidal [16]. We see the largest improvement on the YCB dataset, where PPF is unable to fully resolve the pose of the geometrically symmetric but textually asymmetric objects, seen in failure to match the cylindrical

objects in Figure 3. Our method is able to leverage both color and geometry, selecting the most accurate pose hypothesis. Additionally, we find our method to be comparable to the object-specific results produced by Pix2Pose [38].

B. Evaluating Generalization

As we stated previously, in order to ensure our network is not trained on the test objects we split the objects in YCB-V into two halves, training a network on each set of objects. We select via index parity, as it separates the dataset into splits with roughly equal numbers of symmetric and asymmetric objects, with “Object Set 1” and “Object Set 2” representing the set of objects with even and odd object IDs respectively. To evaluate how well our network generalizes, we compare our results on unseen objects to the objects each network was trained on. The full breakdown of each network’s scores are shown in Table II. Although there is some performance drop on unseen objects, the gap is relatively small, showing the generalization abilities of our method. The “Zero-Shot” column of shows the zero-shot performance of each model on the objects it does *not* see during training.

Tested on	Our method trained on		Zero-Shot
	Set 1	Set 2	
Object Set 1	0.624	0.543	0.543
Object Set 2	0.488	0.496	0.488
All Object	0.557	0.520	0.516

TABLE II: AR scores on YCB-V object subsets.

C. Neighborhood Structure

We explore the effects of different neighborhood structures on the accuracy and generalization of our method. Our method uses a PointNet++ [31] architecture that uses a hierarchical neighborhood structure; we compare this to a CNN architecture that uses a strict neighborhood structure and a PointNet-based architecture [31] that uses a global structure. For the CNN approach, we generate a sparse difference image using the projected point differences before passing it to a ResNet18 network [44]. Our PointNet++ approach uses normalized image coordinates for neighborhood grouping. The PointNet approach contains the normalized image coordinates but it does not perform explicit neighborhood grouping. In Table III, we see that the loose local neighborhood structure found in PointNet++ outperforms the global structure of PointNet as well as the strict structure used in image convolutions. This implies that some neighborhood structure is important for evaluating these sparse point differences, but a too strict neighborhood hampers both performance and generalization.



Input Image



Original PPF Results



Our Improved Results

Fig. 3: Qualitative results on image from YCB-V dataset showing the improved accuracy of our method.

On YCB-V dataset	PointNet++ (Hierarchical)	PointNet (Global)	CNN (Strict)
Seen (Training) Objects	0.624	0.477	0.533
Unseen (Test) Objects	0.488	0.355	0.386
Total	0.557	0.416	0.459

TABLE III: Comparison of the performance of the different neighborhood structure through network architectures.

D. Input Ablations

To determine the relative importance of each of our input channels, we retrain our networks without each dimension. We show results on YCB in Table IV, training on the “Object Set 1” and testing on “Object Set 2”. Additionally, this table shows the effects of concatenating observation and model inputs (“Model without Diff”), as opposed to computing their difference (as in our method). As can be seen, using concatenation instead of differencing gives little change in performance for seen objects, whereas it gives worse performance for unseen objects. Unsurprisingly, the color information has the greatest effect on the accuracy of our system, as it is the most orthogonal to the information used by our PPF hypotheses.

	Model without				
	Color	Depth	Normal	Coords	Diff
Unseen Objects (Zero-shot)	-18%	-15%	-7.1%	-8.9%	-6.3%
Seen Objects (Training)	-24%	-4.2%	0.8%	1.1%	2.1%

TABLE IV: Percent change in AR scores on YCB Video dataset caused by removal of each input to our method.

E. Timing analysis

We analyze the inference speed of our method in Table V. We separate our method into 5 stages, including generating pose hypotheses from SIFT feature matching (“SIFT”), generating pose hypotheses from PPF (“PPF”), computing, transforming and comparing the observation and model values for all hypotheses (“Projection”) and inference with our scoring network (“Scoring”). Note that we only use 100 PPF hypotheses for LM-O, whereas we use an additional 1000 SIFT hypotheses for YCB-V. We found that the LM-O dataset required more accurate initial pose hypotheses, requiring significantly more processing time. To compensate for this, we evaluate the time-performance trade-off of different sets of PPF parameters on the LM-O dataset, shown in blue on Figure 4. Since the LM-O




Fig. 4: Speed accuracy analysis of our method (blue) over various PPF hypothesis generation hyperparameters on LM-O. Base PPF accuracy shown in orange.

dataset is challenging due to strong occlusions and limited scales of objects in the scene, PPF methods [15], [16] need a high sampling rate to produce reasonable pose estimates. Therefore, increased speed comes at the cost of performance, but our method consistently improves the accuracy of the initial hypotheses, shown in red, at all stages of the curve.

	SIFT	PPF	Projection	Scoring	Total
YCB-V	0.142	0.291	0.051	0.135	0.619
LM-O	0	2.900	0.014	0.034	2.949

TABLE V: Test time spent (sec) in each stage of our pipeline.

F. Reconstructed Model Results

To show the effectiveness of our method in robotic scenarios, we test our pipeline on newly generated object model reconstructions. Using fiducial markers [45] and TSDF based surface reconstruction [46], we build textured mesh model of a novel drill object. As shown in Figure 1, we are able to estimate the pose of the target object while in human hands and while being manipulated by the robot; because our method is zero-shot, we do not require any retraining to estimate the pose of new objects such as this one.

VI. CONCLUSION

We propose a method for zero-shot object pose estimation, focusing on pose hypothesis scoring. By extracting point differences between the projected object points and the observation and imposing a loose neighborhood structure on these points, we learn a pose scoring function that generalizes well to novel objects. On the challenging YCB-Video and LineMOD-Occlusion datasets, our method achieves state-of-the-art performance for zero-shot object pose estimation in clutter, evaluated on both textured and untextured objects. We hope that our method paves the way for roboticists to obtain accurate pose estimates for novel objects without needing additional training or data annotation.

REFERENCES

- [1] N. T. Dantam, Z. K. Kingston, S. Chaudhuri, and L. E. Kavraki, "Incremental task and motion planning: A constraint-based approach," in *RSS*, 2016.
- [2] M. Zucker, N. Ratliff, A. D. Dragan, M. Pivtoraiko, M. Klingensmith, C. M. Dellin, J. A. Bagnell, and S. S. Srinivasa, "Chomp: Covariant hamiltonian optimization for motion planning," *IJRR*, 2013.
- [3] S.-K. Kim and M. Likhachev, "Planning for grasp selection of partially occluded objects," in *ICRA*, 2016.
- [4] G. Thomas, M. Chien, A. Tamar, J. A. Ojea, and P. Abbeel, "Learning robotic assembly from cad," in *ICRA*, 2018.
- [5] C. Goldfeder, M. Ciocarlie, H. Dang, and P. K. Allen, "The columbia grasp database," in *ICRA*, 2009.
- [6] M. Ciocarlie, K. Hsiao, E. G. Jones, S. Chitta, R. B. Rusu, and I. A. Sucan, "Towards reliable grasping and manipulation in household environments," in *Experimental Robotics*, 2014.
- [7] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes," *RSS*, 2018.
- [8] C. Wang, D. Xu, Y. Zhu, R. Martín-Martín, C. Lu, L. Fei-Fei, and S. Savarese, "Densefusion: 6d object pose estimation by iterative dense fusion," in *CVPR*, 2019, pp. 3343–3352.
- [9] E. Brachmann, A. Krull, F. Michel, S. Gumhold, J. Shotton, and C. Rother, "Learning 6d object pose estimation using 3d object coordinates," in *ECCV*, Cham, 2014, pp. 536–551.
- [10] J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, and S. Birchfield, "Deep object pose estimation for semantic robotic grasping of household objects," *arXiv preprint arXiv:1809.10790*, 2018.
- [11] X. Deng, A. Mousavian, Y. Xiang, F. Xia, T. Bretl, and D. Fox, "Poserbpf: A rao-blackwellized particle filter for 6d object pose estimation," in *RSS*, 2019.
- [12] M. Sundermeyer, Z.-C. Marton, M. Durner, and R. Triebel, "Augmented autoencoders: Implicit 3d orientation learning for 6d object detection," *IJCV*, vol. 128, no. 3, pp. 714–729, 2020.
- [13] J. Tremblay, T. To, and S. Birchfield, "Falling things: A synthetic dataset for 3d object detection and pose estimation," in *CVPRW*, 2018, pp. 2038–2041.
- [14] D. Dwibedi, I. Misra, and M. Hebert, "Cut, paste and learn: Surprisingly easy synthesis for instance detection," in *ICCV*, 2017, pp. 1301–1310.
- [15] B. Drost, M. Ulrich, N. Navab, and S. Ilic, "Model globally, match locally: Efficient and robust 3d object recognition," in *CVPR*, 2010, pp. 998–1005.
- [16] J. Vidal, C.-Y. Lin, X. Lladó, and R. Martí, "A method for 6d pose estimation of free-form rigid objects using point pair features on range data," *Sensors*, vol. 18, no. 8, p. 2678, 2018.
- [17] K. Park, A. Mousavian, Y. Xiang, and D. Fox, "Latentfusion: End-to-end differentiable reconstruction and rendering for unseen object pose estimation," in *CVPR*, 2020.
- [18] M. Sundermeyer, M. Durner, E. Y. Puang, Z.-C. Marton, N. Vaskevicius, K. O. Arras, and R. Triebel, "Multi-path learning for object pose estimation across domains," in *CVPR*, June 2020.
- [19] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [20] E. Brachmann, A. Krull, S. Nowozin, J. Shotton, F. Michel, S. Gumhold, and C. Rother, "Dsac-differentiable ransac for camera localization," in *CVPR*, 2017, pp. 6684–6692.
- [21] B. Drost and S. Ilic, "3d object detection and localization using multimodal point pair features," in *2012 Second International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission*. IEEE, 2012, pp. 9–16.
- [22] E. Kim and G. Medioni, "3d object recognition in range images using visibility context," in *IROS*, 2011, pp. 3800–3807.
- [23] S. Hinterstoisser, V. Lepetit, N. Rajkumar, and K. Konolige, "Going further with point pair features," in *ECCV*, 2016, pp. 834–848.
- [24] T. Hodan, F. Michel, E. Brachmann, W. Kehl, A. Glent Buch, D. Kraft, B. Drost, J. Vidal, S. Ihke, X. Zabulis, C. Sahin, F. Manhardt, F. Tombari, T.-K. Kim, J. Matas, and C. Rother, "BOP: Benchmark for 6D object pose estimation," *ECCV*, 2018.
- [25] L. Manuelli, W. Gao, P. Florence, and R. Tedrake, "kpam: Keypoint affordances for category-level robotic manipulation," *arXiv preprint arXiv:1903.06684*, 2019.
- [26] H. Wang, S. Sridhar, J. Huang, J. Valentin, S. Song, and L. J. Guibas, "Normalized object coordinate space for category-level 6d object pose and size estimation," in *CVPR*, 2019, pp. 2642–2651.
- [27] Y. Li, G. Wang, X. Ji, Y. Xiang, and D. Fox, "Deepim: Deep iterative matching for 6d pose estimation," in *ECCV*, 2018.
- [28] Y. Xiao, X. Qiu, P.-A. Langlois, M. Aubry, and R. Marlet, "Pose from shape: Deep pose estimation for arbitrary 3d objects," *arXiv preprint arXiv:1906.05105*, 2019.
- [29] C. Xie, Y. Xiang, A. Mousavian, and D. Fox, "The best of both modes: Separately leveraging rgb and depth for unseen object instance segmentation," in *CoRL*, 2020, pp. 1369–1378.
- [30] L. Brynte and F. Kahl, "Pose proposal critic: Robust pose refinement by learning reprojection errors," *arXiv preprint arXiv:2005.06262*, 2020.
- [31] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *CVPR*, July 2017.
- [32] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," in *NeurIPS*, 2017, pp. 5099–5108.
- [33] V. Narayanan and M. Likhachev, "Perch: Perception via search for multi-object recognition and localization," in *ICRA*, 2016, pp. 5052–5059.
- [34] A. Collet, M. Martinez, and S. S. Srinivasa, "The moped framework: Object recognition and pose estimation for manipulation," *IJRR*, vol. 30, no. 10, pp. 1284–1306, 2011.
- [35] D. G. Lowe, "Object recognition from local scale-invariant features," in *ICCV*, vol. 2. Ieee, 1999, pp. 1150–1157.
- [36] MVTec Software GmbH, "Halcon." [Online]. Available: <https://www.mvtec.com/products/halcon/documentation/release-notes-1911-0/>
- [37] Y. Labbe, J. Carpentier, M. Aubry, and J. Sivic, "Cosypose: Consistent multi-view multi-object 6d pose estimation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [38] K. Park, T. Patten, and M. Vincze, "Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation," in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2019.
- [39] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The ycb object and model set: Towards common benchmarks for manipulation research," in *2015 International Conference on Advanced Robotics (ICAR)*, 2015, pp. 510–517.
- [40] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *TPAMI*, vol. 14, no. 2, pp. 239–256, 1992.
- [41] S. Hinterstoisser, V. Lepetit, S. Ilic, S. Holzer, G. Bradski, K. Konolige, and N. Navab, "Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes," in *ACCV*, Berlin, Heidelberg, 2013, pp. 548–562.
- [42] T. Hodan, V. Vineet, R. Gal, E. Shalev, J. Hanzelka, T. Connell, P. Urbina, S. Sinha, and B. Guenter, "Photorealistic image synthesis for object instance detection," *ICIP*, 2019.
- [43] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *ICCV*, 2017, pp. 2961–2969.
- [44] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016, pp. 770–778.
- [45] S. Garrido-Jurado, R. Muñoz-Salinas, F. J. Madrid-Cuevas, and M. J. Marín-Jiménez, "Automatic generation and detection of highly reliable fiducial markers under occlusion," *Pattern Recognition*, vol. 47, no. 6, pp. 2280–2292, 2014.
- [46] Q.-Y. Zhou and V. Koltun, "Dense scene reconstruction with points of interest," *ACM Transactions on Graphics (ToG)*, vol. 32, no. 4, pp. 1–8, 2013.