

Novel Object Viewpoint Estimation through Reconstruction Alignment

Mohamed El Banani Jason J. Corso David F. Fouhey
University of Michigan

{mbanani, jjcorso, fouhey}@umich.edu

Abstract

The goal of this paper is to estimate the viewpoint for a novel object. Standard viewpoint estimation approaches generally fail on this task due to their reliance on a 3D model for alignment or large amounts of class-specific training data and their corresponding canonical pose. We overcome those limitations by learning a reconstruct and align approach. Our key insight is that although we do not have an explicit 3D model or a predefined canonical pose, we can still learn to estimate the object’s shape in the viewer’s frame and then use an image to provide our reference model or canonical pose. In particular, we propose learning two networks: the first maps images to a 3D geometry-aware feature bottleneck and is trained via an image-to-image translation loss; the second learns whether two instances of features are aligned. At test time, our model finds the relative transformation that best aligns the bottleneck features of our test image to a reference image. We evaluate our method on novel object viewpoint estimation by generalizing across different datasets, analyzing the impact of our different modules, and providing a qualitative analysis of the learned features to identify what representations are being learnt for alignment.

1. Introduction

Consider the two views of the owl in Figure 1. How are they related? As humans, we can easily imagine how to move from view 1 to view 2, even if we have never seen this object before. This problem of understanding object pose and viewpoint has long fascinated researchers in both computer vision (starting with the first PhD thesis on computer vision [21]) and psychology [23, 27]. Within computer vision, many approaches have been proposed to estimate an image’s viewpoint with respect to a given, or assumed, oriented 3D model. Despite the success of those approaches on known classes, they struggle with novel objects since both the 3D model and its canonical pose are unknown. In this work, we learn to predict the relative viewpoint for a novel object using a single reference view.

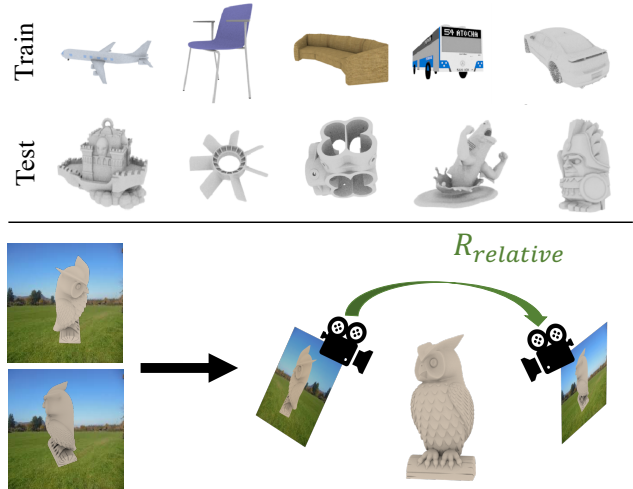


Figure 1. Humans can not help but see the 3D structure in those two views, which makes aligning their viewpoints very easy. This paper proposes a reconstruct-and-align approach to learning viewpoint estimation for novel objects.

Historically, there has been two general approaches to viewpoint estimation. First, 3D model alignment approaches find a transformation that aligns an image to a known 3D model [9, 20, 39]. Despite their efficacy, those approaches were limited to objects with available 3D models. Second, end-to-end discriminative approaches learn to directly estimate the image’s viewpoint with respect to a canonical pose (e.g., head-on for a car) [24, 29, 30, 33]. While those approaches do not rely on explicit 3D models, their predictions are with respect to an implicitly-defined, canonically-oriented model. Due to their reliance on explicitly- or implicitly-defined 3D models, those approaches struggle with novel objects where neither the model nor its *canonical* orientation are known.

How can we estimate the viewpoint of a novel object? Viewpoint is defined with respect to a class-specific coordinate frame or a canonical pose (e.g., head-on for a car) [33]. For small generalization (e.g., from bikes to motorcycles), one can exploit similarity in viewpoint appearance between

categories [15], but this is rarely the case. Motivated by research on mental rotation [23], we observe that an image, as opposed to a 3D model, can serve as a reference pose. Hence, instead of aligning to a 3D model, we find the relative transformation between two views of the same object.

In this work, we combine insights from both 3D alignment and end-to-end learning. We do not use a 3D model, but instead learn to map the input image to a 3D feature grid. Our key insight is that, although we do not have an explicit 3D model, we can still learn a deep network that maps each object instance to a 3D reference frame. This is done by combining learned 2D layers with projection layers constrained to follow projective geometry, similar to [13]. We then learn to align these 3D feature grids to identify the relative transformation that would minimize misalignment in the joint 3D feature grid.

We learn the 3D feature grid via 2D supervision (but not direct 3D supervision via voxels), while incorporating physical constraints inspired by 3D shape carving [14]. We hypothesize that forcing the representation to pass through a geometry-aware bottleneck provides a useful inductive bias. Further, the unlearned projection layers make it far easier for the network to learn an implicit 3D shape representation that can be inspected, where unstructured layers might lead to pure memorization that generalizes worse [28]. We then train an alignment network to estimate whether two views would produce an aligned 3D shape representation by training on randomly misaligned examples similar to [3, 16].

We evaluate novelty via generalization across different datasets. For instance, we show that our model can generalize well from ShapeNet [4] to Thingi10K [45] models, despite them looking very different as shown in Figure 1 and the supplementary material. This sidesteps difficulties involved in using object classes as a proxy for novelty [15, 44, 38]: class distinctions can be arbitrary, with some classes like motorcycles and bicycles sharing a lot of view-specific appearances.

We demonstrate the effectiveness of our proposed approach in Section 5, and show that while standard approaches are challenged by novel objects, our approach generalizes gracefully to new datasets. Specifically, our model achieves a viewpoint estimation accuracy of 40% when generalizing from ShapeNet to Thingi10K, compared to 25% accuracy achieved by standard approaches. We run several analysis experiments to understand the learned representations, and find that our model also shows good across-dataset generalization on view prediction and voxel prediction; achieving an IoU score of 0.43 on ShapeNet airplanes despite being trained on Thingi10K. Overall, our experiments show that a reconstruction and alignment approach generalizes to novel objects better than discriminative approaches, and we hope that our work will spark more interest in this direction.

2. Related Work

The ability to accurately orient an object in 3D space in one of the oldest problems in computer vision [21]. Early work posed the problem as an alignment problem, specifically that of estimating “the transformation needed to map an object model from its inherent coordinate system into agreement with sensory data” [9]. This formulation was powerful as it was task agnostic as long as one had access to the 3D model. Early approaches focused on finding correspondences between images and specific model instances [2, 10, 17]. More recent approaches focused on generalizing to all instances within an object classes by using class-specific 3D key-points [31, 37, 39, 46] or deformable part models [1, 20]. While we are inspired by alignment approaches, we assume no access to a 3D model and generalize to novel object classes by learning to both reconstruct the 3D model and align images to it.

Motivated by the success of large-scale image classification, a recent line of work focused on learning pose estimation with little or no supervision beyond image-pose pairs [8, 18, 24, 29, 30, 33, 40]. Those approaches replaced the alignment approach with fully-learned 2D feature-based models like convolutional neural networks (CNNs). Although this approach has been successful at synthetic-to-real generalization [24, 29] and within-class generalization [8, 33], it has remained unclear how successful it would be at generalizing to novel classes. While we also do not assume access to a 3D model; unlike this work, we train our model to be agnostic to object class, learn a 3D object representation, and generalize to novel classes.

A recent line of work incorporates other losses or tasks with pose estimation such as object classification [15, 19, 20, 24], keypoint detection [31, 33, 46], and object reconstruction [12, 32, 42]. Closest to our approach are [12, 32] who jointly learn 3D reconstruction and pose prediction from unannotated images. While our work stems from the same observation that pose and shape are closely related, our goals are different. While [12, 32] are interested in learning shape and pose for a specific classes from very weak supervision, our focus is to generalize to novel classes of objects by learning from standard viewpoint supervision.

Independent of this work, there has been a rising interest in learning more structured or meaningful intermediate representations through representational bottlenecks [20, 33, 35, 36, 43, 46]. The representational bottleneck is interpreted as providing a good inductive bias to the learning process. More relevant to us are approaches that force a 3D representational bottleneck in the form of a voxel grid [13, 32, 35, 41, 42]. While some approaches use 3D supervision, others leverage the 3D representation bottleneck to learn 3D reconstruction using only 2D supervision. We would like to emphasize that while we are inspired by this line of work, our goals are different; while those approaches

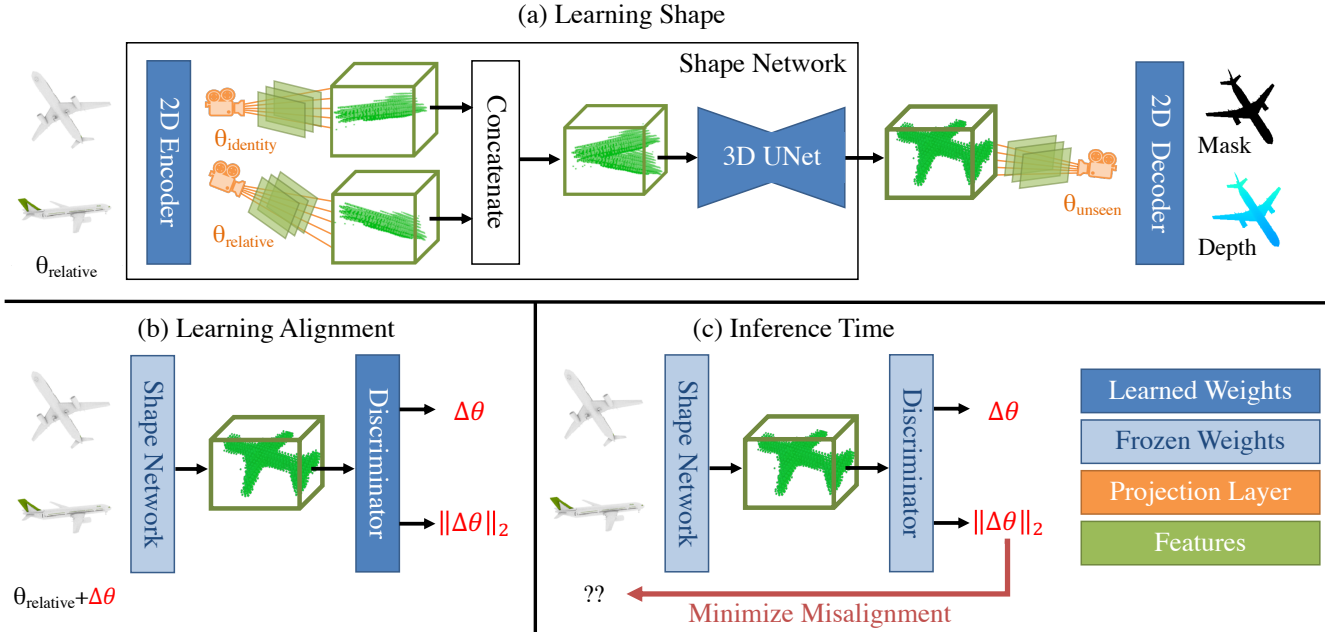


Figure 2. **Approach Overview.** Our approach consists of two learning stages and an inference stage. (a) We first learn to predict an object’s shape from two views of the object and their relative pose. We train our Shape Network to estimate the object’s appearance from a third viewpoint. (b) Using the trained Shape Network, we train a discriminator to predict the degree of misalignment between two views. (c) At inference time, we find the relative pose that would best align the two inputs. Note: Input image backgrounds are removed for clarity.

focused on performing accurate single- or multi-view 3D reconstruction of a known class of objects, our goal is not to reconstruct the object, but rather to use the 3D bottleneck to allow us to estimate the relative viewpoint for novel objects.

Finally, our goal is to generalize to unseen classes. While this problem has been studied extensively in the scope of image classification (see [38] for a survey), it has not received a lot of attention for 3D tasks such as pose estimation. To the best of our knowledge, there has only been two previous approaches that tackled similar problems. Kuznetsova *et al.* [15] proposed a metric learning approach that performs joint classification and pose-estimation, and they leverage the learned space for zero-shot pose estimation. However, their approach only works if the novel object is similar to previously seen objects (*e.g.*, cars to buses and bicycles to motorcycles), which limits the generalizability of their method. Tseng *et al.* [31] proposed a keypoint alignment approach that learns to predict 3D keypoints to align the novel class. However, they expect 3D keypoint annotation for the reference images of the novel object at test time, while we only assume a single, unannotated, reference image.

3. Novel Object Viewpoint Estimation

Our overarching goal is to build a system that understands the *viewpoint* of a previously unseen object using as little information about the object as possible, and ideally

generalizes between completely unrelated objects, such as zebras to forks instead of cars to buses. The difficulty is that viewpoint is defined with respect to a coordinate frame or canonical pose (*e.g.*, head-on for a car): without it, the problem is fundamentally undefined and any viewpoint could be the origin viewpoint. Accordingly, past work has defined the coordinate frame using keypoints [31], or generalized across semantically similar objects [15] (*e.g.*, cars to buses).

Rather than use a predefined coordinate frame via keypoints or semantic similarity, we propose to instead use a single *image* as a reference. Our goal then is, given a reference image I_1 with viewpoint \mathbf{v}_1 and an image I_2 whose viewpoint \mathbf{v}_2 we would like to estimate, we want to predict the relative rotation $\mathbf{R}_{\text{relative}}$ such that $\mathbf{v}_2 = \mathbf{R}_{\text{relative}}\mathbf{v}_1$. This formulation maintains the essence of viewpoint estimation, while circumventing the limitations posed by requiring a canonical pose. It should be noted given the relative pose between two images and the canonical pose of the object in one image, it is trivial to calculate the *canonicalized* pose in the second image.

During training, we assume we have access to image-viewpoint pairs from an arbitrary number of classes. At test time, we are presented with two images from a new class and tasked with predicting the relative viewpoint of the second image with respect to the first image.

4. Approach

The goal of this work is to build a system that can predict the relative viewpoint between two views of an object that it has never seen before. Our approach, shown in Figure 2, takes inspiration from the early formulation of pose estimation as alignment [9]. However, instead of using a 3D model, we learn to predict a 3D representation from each view that can be aligned. We first present a high-level sketch of our approach before explaining each stage in more detail. Architectural details are presented in Appendix A.1. We analyze the effectiveness of our design choices in Section 5.3.

Approach Sketch. In the first stage, we *learn shape* by training our model on a view prediction task. Given two images and the relative viewpoint between them, we train our model to extract 2D features from each image and back-project them into a joint 3D feature tensor. During back-projection, we use the object’s depth from each view to carve out the empty space between the viewer and the object. For an arbitrary third view, we project the 3D feature tensor into 2D, and predict either the object mask or depth in that view. In the second stage, we *learn alignment* by training a discriminator to predict if the extracted 3D feature tensor is well-aligned or not. At test time, we find the relative viewpoint by finding the transformation that minimizes predicted misalignment.

4.1. Learning Shape

Given two images and a relative pose, $\{I_1, I_2, \mathbf{R}_{\text{relative}}\}$, we train a Shape Network to generate a 3D feature grid, $\mathcal{F}_{\text{object}} \in \mathbb{R}^{F \times N \times N \times N}$, that captures the object’s shape. Since we assume that the viewer is looking straight at the object for the first image, the second viewpoint is $\mathbf{R}_{\text{relative}}$.

The Shape Network consists of a 2D CNN encoder, a differentiable back-projection layer, and a 3D UNet [22, 5]. We provide the specific architectural details in the supplementary material. The 2D encoder extracts a 2D feature map for each image that is then back-projected into a 3D voxel grid of dimension N ($N=32$) using a viewpoint and a depth map. The back-projection is based on the differentiable projection layers from Kar *et al.* [13]. We back-project the features by projecting rays for each pixel location in the 2D feature map into the 3D grid using the appropriate viewpoint, the value for each voxel is then set via bi-linear sampling.

Inspired by early work on space carving [14], we leverage known or estimated depth from the object to carve away the empty space in-front and around the object by setting the features in those voxels to 0. After space carving, we have two feature grid tensors, $\mathcal{F}_{v_1}, \mathcal{F}_{v_2}$, that capture the appearance and view of each object in the same frame of reference.

The two feature grid tensors from each view do not see the same parts of the objects and are, essentially, hallucinating the unseen parts.

Therefore, we fuse them with the goal of predicting a consistent 3D representation of the object. We accomplish this by concatenating the two feature grids along the feature dimension to produce a tensor, $\mathbb{R}^{2F \times N \times N \times N}$, and pass it into a 3D UNet. The 3D UNet’s goal is to refine this representation by aligning the seen and hallucinating the unseen.

We train the Shape Network on 2D view prediction of the object’s mask or depth from a third view. This is done by projecting the 3D features to 2D using the relative viewpoint of the third view. Since we do not know the object’s depth, we sample the features at multiple depth values along each ray. We use a 1×1 convolution to aggregate features for all sampled depth values for a given pixel location. This is followed by two convolutional layers and a single up-sampling layer in the middle to match the depth or mask output size. We train the network to minimize binary cross entropy from mask prediction, and L_1 loss for depth prediction.

4.2. Learning Alignment

Once we have trained a Shape Network to estimate the object’s 3D feature representation, our goal for the second stage is to use the 3D representation to predict whether two images of an object are aligned or not. We generate training data by freezing the Shape Network’s weights and perturbing its input relative viewpoint by a random rotation. We then train a discriminator to predict the magnitude of that perturbation. Our intuition is that while determining the relative viewpoint between two objects might be difficult, misalignment is much easier to detect as it produces inconsistent object shapes, with the degree of inconsistency being a measure of how misaligned the images were.

We implement the discriminator as a 3D CNN. Our network architecture consists of two 3D inception modules followed by two fully-connected layers. We adapt the inception module [26] to 3D by expanding the kernels along the depth dimension. We use the inception module as it allows the network to detect features at multiple scales which would be useful when detecting inconsistencies.

We train the discriminator to predict the perturbation direction (Euler angles) and magnitude (geodesic distance). We randomly perturbed the relative viewpoint of two thirds of our instances by at least 10 degrees. We train the network using an L_2 loss for all 4 outputs.

4.3. Relative Viewpoint Estimation

Given a trained Shape Network and a discriminator, the relative viewpoint is the viewpoint that will minimize predicted misalignment. This problem could be solved as an optimization problem or via a greedy search. In this work, we uniformly sample viewpoints and output the viewpoint that minimized the discriminator’s predicted perturbation magnitude.

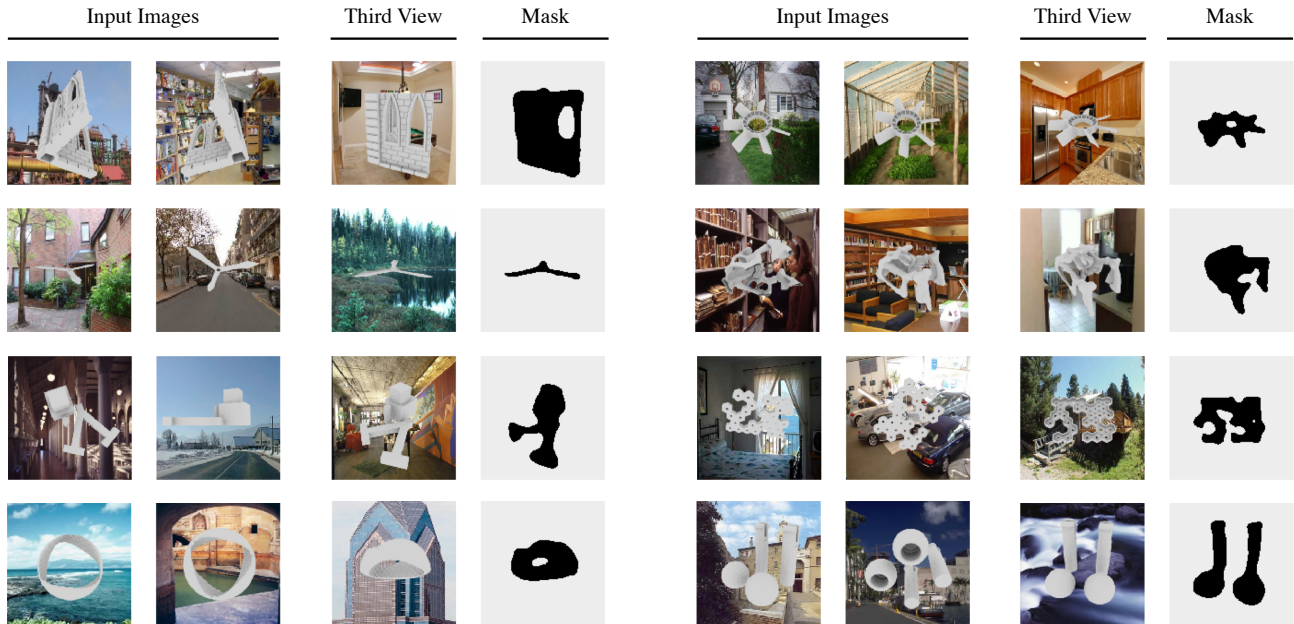


Figure 3. **Mask Prediction on Thing10k.** Given two images, our models learn to predict the mask for a third view. The first two columns present two input views, while the third column is the view for which the mask is being predicted. The masks shown are predicted by a model which is trained on a different dataset; ShapeNet. We can accurately predict the 3D model’s projection from a novel viewpoint despite being trained on a very different set of shapes.

5. Experiments

We now empirically evaluate our model’s performance on relative viewpoint estimation, as well as report additional experiments that analyze the different components of the model. Our experiments aim to answer two questions: (1) how well our proposed approach works, especially in the case of generalization to novel objects; and (2) how the method works and whether it has indeed learned to solve the task by using 3D representations.

We address the first question by evaluating the model’s ability to generalize across datasets for both viewpoint estimation (Section 5.1) and view prediction (Section 5.2). Cross-dataset generalization gives us a good proxy for how the model would perform on novel objects since ShapeNet and Thing10K models look drastically different.

We address the second question by evaluating our model’s ability to learn shape through view prediction (Section 5.2) and voxel prediction (Section 5.4). We also run several ablations to better understand the significance of different model components (Section 5.3).

Datasets. We use three datasets: ShapeNet, Pix3D, and Thing10K. While ShapeNet and Pix3D are fairly uniform in terms of their models’ types and appearances, Thing10K models vary greatly in terms of their size, appearance, and geometric properties. Thing10K [45] consists of 10K models created for 3D printing. The lack of strong patterns or prototypical shapes makes Thing10K very chal-

lenging for both pose estimation and reconstruction methods, and an excellent test dataset for our task. Alternatively, ShapeNet’s large size (55 categories and $\sim 57K$ models) makes it a great training set, while allowing us to evaluate our approach on models from the same class. Finally, Pix3D [25] is a smaller dataset, consisting of 395 3D models of 9 categories of furniture. The smaller number of models and domain-specificity of Pix3D provide an interesting comparison against the larger variance in object categories in ShapeNet and object shapes in Thing10K.

Data Sampling. We choose our poses by uniformly sampling around a viewing sphere. With models centered around the origin, the pose transformation is simply the inverse of the viewpoint transformation. This deviates from previous work which sampled according to consumer photos’ distributions to promote better synthetic to real generalization [24].

After excluding models with missing or corrupted 3D model files, we end up with 55,281 ShapeNet models, 390 Pix3D models, and 9994 Thing10K models. For ShapeNet and Thing10K, we sample 20 different poses per object, while we sample 200 views for Pix3D. We do so to mitigate the large difference between the numbers of models in those datasets. We randomly split our datasets by models to ensure that each of the splits has a disjoint set of models. We use 80%/10%/10% of the data as our training, validation, and test splits.

Table 1. **Relative Viewpoint Estimation Performance.** We evaluate our model on its ability to perform within- and across-dataset generalization. The letters represents the initial letter of each dataset: ShapeNet, Pix3D, and Thingi10k. Each column represents different train/test settings $\mathcal{S} \rightarrow \mathcal{T}$ is training on ShapeNet and testing on Thingi10k. We find that our model is able to generalize across datasets, while previous approaches demonstrate a clear bias towards their training distributions.

	$\mathcal{S} \rightarrow \mathcal{S}$		$\mathcal{P} \rightarrow \mathcal{P}$		$\mathcal{T} \rightarrow \mathcal{T}$		$\mathcal{S} \rightarrow \mathcal{P}$		$\mathcal{S} \rightarrow \mathcal{T}$		$\mathcal{P} \rightarrow \mathcal{T}$	
	$Acc_{\pi/6}$	MedErr	$Acc_{\pi/6}$	MedErr	$Acc_{\pi/6}$	MedErr	$Acc_{\pi/6}$	MedErr	$Acc_{\pi/6}$	MedErr	$Acc_{\pi/6}$	MedErr
Stat. Prior	4.15	114.71	4.13	115.01	4.33	114.41	4.13	115.01	4.33	114.41	4.33	114.41
RPNet	60.23	44.32	68.93	32.61	20.82	91.43	66.42	46.07	25.00	91.85	14.43	111.80
ConvAE	50.49	52.65	51.96	56.36	15.35	100.56	48.63	46.78	8.65	116.11	3.55	129.29
Ours	47.57	66.04	58.08	46.02	41.25	61.96	60.67	43.47	40.12	70.08	33.00	96.82

5.1. Viewpoint Estimation

We first evaluate how well our system performs on viewpoint estimation, our primary task. Our aim here is to evaluate how well the system performs, especially when generalizing from one dataset to another.

Experimental Setup. We evaluate pose estimation performance based on the geodesic distance between the predicted and ground-truth pose as was done in previous work [24, 33]. Given two rotation matrices, the geodesic distance provides the magnitude of the transformation between those two rotation matrices. Following [11], we calculate geodesic distance as $D(R_{gt}, R_{pr}) = \|\log \mathbf{R}_{gt} \mathbf{R}_{pr}^T\|$. Following [24, 33], we report the median geodesic distance and the percentage of viewpoints under threshold $\pi/6$.

Baselines. We compare our models performance to several baselines. For all our baselines, we use binned Euler angles to represent rotation and train the models using a cross-entropy loss.

Statistical Prior. We find the mode pose for the validation and use it on testing. Since we uniformly sampled the viewpoints, we find that this baseline is very weak and corresponds to a random baselines.

RPNet. En *et al.* [7] proposed a Siamese-network to perform relative camera pose estimation. As reported by the authors, we find that for more difficult datasets, it is better to train each network separately on absolute viewpoint and then calculate the relative viewpoint explicitly. We used binned Euler angles instead of quaternions as we found that they perform better. Furthermore, we provide their network with RGB-D input to provide a fair comparison to our method’s use of depth-based space carving.

Convolutional Autoencoder (ConvAE). Given that our model learns both reconstruction and pose estimation, we trained a convolutional autoencoder on both tasks. We initially used UNets [22] due to their success in image translation tasks, but found that removing the skip connections improved the viewpoint estimation performance. Viewpoint is estimated on the encoded features using a 3-layer CNN. Like RPNet, we found that training on absolute viewpoint improved performance.

Results. We find that directly estimating viewpoint scales poorly do datasets like Thingi10K. As shown in Figure 3, Thingi10K models exhibit very high variability which makes them challenging for a model that maps 2D patterns directly to a pose as seen by in Table 1. By contrast, our model is able to leverage this high variability in the data as it allows it to learn better reconstruction features. This is evident by the large performance gain that our model on Thingi10K; both when training and testing on it, or when transferring to it from another dataset.

On the other hand, we find that if there is limited difference between training and testing (e.g., training and testing on ShapeNet), directly estimating viewpoint does well. Interestingly, we find that in this domain, training on more data with higher variance improves performance as seen with training on ShapeNet and evaluating on Pix3D, especially since all of Pix3D’s classes exist in ShapeNet. This pattern is shown both for the direct pose models as well as our approach. Another interesting finding is that joint reconstruction and pose estimation do not explain our models performance as shown by the convolutional autoencoder baseline. It appears that the additional reconstruction task burdens the network as opposed to providing it with an additional learning signal.

We note that RPNet is representative of current end-to-end discriminative viewpoint-estimation models as they generally follow the structure of a CNN backbone and FC layers to predict viewpoint [18, 24, 30, 33, 40]. As a result, the comparison to RPNet represents a more general comparison against problems that frame viewpoint estimation as a discriminative task.

Furthermore, recent approaches proposed jointly learning reconstruction and pose estimation, where the network learns to predict pose in an *emergent* coordinate frame [12, 32]. While those models are trained to minimize a reconstruction error, they are still using a CNN to directly estimate the pose with respect to a class-specific, *emergent* canonical pose. Hence, they face the same challenges as RPNet. Additionally, those approaches can perform better when given ground-truth pose, which we provide to RPNet.

Table 2. **2D Reconstruction Cross-dataset Generalization Experiment.** The rows show the training dataset, while the columns present the test dataset. We find that our model can easily generalize across datasets, with very minor performance drops. This supports our claim that our model is learning a generalizable representation of shape, as opposed to memorizing the shapes available within each dataset.

	ShapeNet				Pix3D				Thing10k			
	Mask		Depth		Mask		Depth		Mask		Depth	
	IoU	F1	LogRMSE	$\delta_{1.25}$	IoU	F1	LogRMSE	$\delta_{1.25}$	IoU	F1	LogRMSE	$\delta_{1.25}$
ShapeNet	0.83	0.90	0.53	0.66	0.80	0.88	0.63	0.60	0.82	0.90	0.35	0.84
Pix3D	0.78	0.87	0.53	0.74	0.79	0.88	0.59	0.74	0.79	0.87	0.34	0.90
Thing10k	0.76	0.85	0.69	0.70	0.73	0.84	0.81	0.63	0.85	0.91	0.30	0.88

5.2. 2D Reconstruction

We now analyze our Shape Network’s ability to represent shape, as well as how well that ability generalizes to novel objects. Specifically, we evaluate the Shape Network’s ability to predict the object’s mask or depth from an unseen viewpoint. Note that our model is trained on either depth or mask prediction. Hence, the depth and mask results are shown for models trained on the respective task. Similar to the previous experiment, we are interested in the model’s ability to generalize to unseen classes. We emphasize that we have a disjoint set of training and testing models, hence, the testing performance on the same dataset evaluates class-specific generalization, while testing across datasets evaluates generalization to unseen classes.

Experimental Setup. We train and test the performance on the 9 pairwise dataset pairs for ShapeNet, Pix3D, and Thing10K. We evaluate mask reconstructions using IoU at a threshold of 0.5 as well as a pixel-level F1 score. Following [6], we evaluate depth estimation using log root mean-square-error and threshold accuracy, $\delta_{1.25}$. For a depth prediction, y , threshold accuracy is calculated as the percentage of pixels, y_i , s.t. $\max(\frac{y_i}{y_i^{gt}}, \frac{y_i^{gt}}{y_i}) < 1.25$.

Results. We find that our model achieves high performance for mask prediction across all the datasets as shown in Table 2. More importantly, we find that there is very little variance between the model’s performance across datasets. This demonstrates our model’s ability to learn to represent 3D shape in a generalizable way. This is in stark contrast to recent findings from reconstruction that find that 3D reconstruction approaches will often memorize some of the training models [28]. We hypothesize that our model is less susceptible to memorization due to (a) the geometry-aware bottleneck and (b) the lack of a shared canonical pose across class instances. The geometry-aware bottle neck forces the model to learn to represent its input in 3D. Furthermore, by removing the canonical pose assumption, the model can no longer rely on specific locations in the tensor having a consistent representation (e.g., airplane wings always being on the side). As a result, the model is forced to represent the specific input, as opposed to just classifying it.

Table 3. **Ablation Study.** We ablate our model by removing two key components in the 3D and analyzing the effect on 2D reconstruction and relative viewpoint estimation.

	$\mathcal{S} \rightarrow \mathcal{S}$			$\mathcal{S} \rightarrow \mathcal{T}$		
	$Acc_{\pi/6}$	MedErr	IoU	$Acc_{\pi/6}$	MedErr	IoU
Ours	47.57	66.04	0.83	40.12	70.08	0.82
- carving	6.38	134.23	0.64	4.21	130.14	0.63
- refine	47.57	59.76	0.63	42.16	54.03	0.64

Of course, like any learned approach, our model is still affected by the quality of the training data. We observe that, in general, the best performing model on a dataset is the one trained on that dataset with two interesting exceptions. The first exception is that for mask prediction, ShapeNet does very well on Pix3D. This is likely due to the degree of class overlap between ShapeNet and Pix3D, as well as the additional variety of models seen in ShapeNet. The second exception is that Pix3D trained model tend to do better on depth prediction. While it is not clear why this would happen, one potential explanation is that the planarity of most furniture might provide a good scaffolding for the networks to learn better depth, while other datasets provide too much variance for the network to learn good representations.

5.3. Ablations

We perform an ablation study to better understand the model’s performance and failure modes. In particular, we ablate two central elements of our model: 3D refinement and space carving. We train our ablated models on ShapeNet and evaluate them on both ShapeNet and Thing10K. We follow the same training procedure as before. The results are shown in Table 3.

Space Carving. We find that removing space carving is extremely detrimental to the model. The model’s 2D reconstruction performance suffers, while the viewpoint performance is reducing to random guessing. Space carving appears to greatly boost our performance by providing a class and model agnostic benefit. Given two projected views, space carving provides a high degree of certainty to the model regarding which voxels can be ignored, and hence, greatly improves the alignment. The uncertainty about the object’s extent appears to stifle its ability to align views.




Figure 4. **Features To Voxels.** Our model can accurately predict occupancy without being trained using voxels or for voxel prediction. We present 3 success cases ($\text{IoU} > 0.7$; in blue) and 3 failure cases of our model ($\text{IoU} < 0.1$; in red), with our prediction on the left and the ground-truth on the right. We observe that our failure cases tend to arise from failure to complete an occluded part or registering the part at an incorrect location, which would be expected from a model that has never seen a plane before.

3D Refinement. We observe that removing the 3D refinement network also greatly deteriorates the reconstruction performance. However, unlike space carving, removing it actually boosts the viewpoint performance. This is likely caused by the network hallucination of features (in order to complete the object), which in turn is mixed with the non-hallucinated features, which may hamper alignment.

5.4. Features to Voxels

Given that the model learns to project to and from a voxel grid, we expect that the features of this grid can capture occupancy despite never having been trained on 3D information. While we do not ever directly extract an occupancy grid during training, we find that we can adapt one of the layers in the network to generate it. Specifically, while training the Shape Network, we add a multi-view consistency loss [32, 34]. In particular, Tulsiani *et al.* [32] propose a loss that samples different points along projected rays to detect ray termination events. We extend their formulation to a feature grid by defining a linear layer that maps the ray-sampled features to a permeability score. Since the permeability layer maps each feature vector to a ray termination probability, it correlates with the occupancy likelihood. Hence, we apply it to the feature tensor output from the Shape Network to generate an occupancy grid. We note that removing this loss does not significantly change our viewpoint estimation or 2D reconstruction performance.

We train our model on Thing10K and evaluate it on the airplane class of ShapeNet. We would like to emphasize that Thing10K does not include an airplane class, so we are evaluating on unseen objects. We use the voxel grids

provided by ShapeNet, but we center and resize them to fit a voxel grid of dimension 32. We evaluate the quality of our rendering by calculating the intersection over union (IoU) of the occupancy grids. We pick the threshold based on the validation set.

Since our embedding is in a viewer-centric frame of reference, we rotate it back to the canonical frame before computing IoU. Our model achieves a mean IoU of 0.43 with 98% of the samples having an IoU at least 0.25. We note that this is calculated on the direct embedding dimension without performing any registration or scaling to better match ShapeNet 3D models. We visualize several success and failure cases in Figure 4.

6. Conclusion

In this paper, we present a reconstruction and alignment approach for novel object viewpoint prediction. We observe that previous viewpoint prediction approaches have either relied on large amounts of class-specific, canonically-oriented data or on having a 3D model for alignment. However, those two assumptions do not apply when dealing with a novel object. Our key insight is that while we may not have access to a 3D model, we can learn a deep network that estimates the object’s shape in the viewer’s frame of reference, and use a reference view to determine the viewpoint with respect to other views of the object. To this end, we propose learning two networks: the first network uses a 3D feature bottleneck to represent objects in 3D; the second network learns whether or not two views are aligned.

We evaluate our approach on several datasets. We find that our approach significantly outperforms standard viewpoint estimation approaches when there is a large domain shift between the train and test models, or when the training models themselves have high variance. Furthermore, we find that we can extract 3D occupancy grids despite not training using 3D supervision. This provides ample evidence that our models can generalize to unseen objects.

Our current models greatly rely on the depth input to accurately reconstruct the object. While the 2.5D input might be needed to understand the 3D structure of a novel object, we would like to explore ways of minimizing this reliance in future work. Furthermore, we would like to explore optimization approaches that would allow us to converge to the best viewpoint instead searching over a set of predetermined viewpoints.

Acknowledgments We would like to thank the reviewers and area chairs for their valuable comments and suggestions, and the members of the UM AI Lab for many helpful discussions. Toyota Research Institute (“TRI”) provided funds to assist the authors with their research but this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity.

References

- [1] Alexander C Berg, Tamara L Berg, and Jitendra Malik. Shape matching and object recognition using low distortion correspondences. In *CVPR*, 2005.
- [2] Rodney A Brooks. Symbolic reasoning among 3-d models and 2-d images. *Artificial intelligence*, 17, 1981.
- [3] Joao Carreira, Pulkit Agrawal, Katerina Fragkiadaki, and Jitendra Malik. Human pose estimation with iterative error feedback. In *CVPR*, 2016.
- [4] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.
- [5] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2016.
- [6] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NeurIPS*, 2014.
- [7] Sovann En, Alexis Lechervy, and Frédéric Jurie. Rpnnet: an end-to-end network for relative camera pose estimation. In *ECCV*, 2018.
- [8] Amir Ghodrati, Marco Pedersoli, and Tinne Tuytelaars. Is 2d information enough for viewpoint estimation? In *BMVC*, 2014.
- [9] William Eric Leifur Grimson and Daniel P Huttenlocher. *Object recognition by computer: the role of geometric constraints*. MIT Press, 1990.
- [10] Daniel P Huttenlocher and Shimon Ullman. Recognizing solid objects by alignment with an image. *International Journal of Computer Vision*, 5(2), 1990.
- [11] Du Q Huynh. Metrics for 3d rotations: Comparison and analysis. *Journal of Mathematical Imaging and Vision*, 35(2):155–164, 2009.
- [12] Eldar Insafutdinov and Alexey Dosovitskiy. Unsupervised learning of shape and pose with differentiable point clouds. In *NeurIPS*, 2018.
- [13] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. In *NeurIPS*, 2017.
- [14] Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. *International Journal of Computer Vision*, 2000.
- [15] Alina Kuznetsova, Sung Ju Hwang, Bodo Rosenhahn, and Leonid Sigal. Exploiting view-specific appearance similarities across classes for zero-shot pose prediction: a metric learning approach. In *AAAI*, 2016.
- [16] Ke Li, Bharath Hariharan, and Jitendra Malik. Iterative instance segmentation. In *CVPR*, 2016.
- [17] David G Lowe. Three-dimensional object recognition from single two-dimensional images. *Artificial intelligence*, 1987.
- [18] Siddharth Mahendran, Haider Ali, and René Vidal. 3d pose regression using convolutional neural networks. In *ICCV*, 2017.
- [19] Xingchao Peng, Baochen Sun, Karim Ali, and Kate Saenko. Learning deep object detectors from 3d models. In *ICCV*, 2015.
- [20] Bojan Pepik, Peter Gehler, Michael Stark, and Bernt Schiele. 3d 2 pm–3d deformable part models. In *ECCV*, 2012.
- [21] Lawrence Roberts. *Machine Perception of 3D Solids*. PhD thesis, Massachusetts Institute of Technology, 1965.
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2015.
- [23] Roger N Shepard and Jacqueline Metzler. Mental rotation of three-dimensional objects. *Science*, 1971.
- [24] Hao Su, Charles R. Qi, Yangyan Li, and Leonidas J. Guibas. Render for cnn: Viewpoint estimation in images using cnns trained with rendered 3d model views. In *ICCV*, 2015.
- [25] Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B Tenenbaum, and William T Freeman. Pix3d: Dataset and methods for single-image 3d shape modeling. In *CVPR*, 2018.
- [26] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [27] Michael J Tarr and Heinrich H Bülthoff. Image-based object recognition in man, monkey and machine. *Cognition*, 67, 1998.
- [28] Maxim Tatarchenko, Stephan R Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In *CVPR*, 2019.
- [29] Jonathan Tremblay, Thang To, and Stan Birchfield. Falling things: A synthetic dataset for 3d object detection and pose estimation. In *CVPR Workshops*, 2018.
- [30] 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 *CoRL*, 2018.
- [31] Hung-Yu Tseng, Shalini De Mello, Jonathan Tremblay, Sifei Liu, Stan Birchfield, Ming-Hsuan Yang, and Jan Kautz. Few-shot viewpoint estimation. In *BMVC*, 2019.
- [32] Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Multi-view consistency as supervisory signal for learning shape and pose prediction. In *CVPR*, 2018.
- [33] Shubham Tulsiani and Jitendra Malik. Viewpoints and keypoints. In *CVPR*, 2015.
- [34] Shubham Tulsiani, Tinghui Zhou, Alexei A. Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *CVPR*, 2017.
- [35] Hsiao-Yu Fish Tung, Ricson Cheng, and Katerina Fragkiadaki. Learning spatial common sense with geometry-aware recurrent networks. In *CVPR*, 2019.
- [36] Jiajun Wu, Yifan Wang, Tianfan Xue, Xingyuan Sun, Bill Freeman, and Josh Tenenbaum. Marnet: 3d shape reconstruction via 2.5d sketches. In *NeurIPS*, 2017.

- [37] Jiajun Wu, Tianfan Xue, Joseph J Lim, Yuandong Tian, Joshua B Tenenbaum, Antonio Torralba, and William T Freeman. Single image 3d interpreter network. In *ECCV*, 2016.
- [38] Yongqin Xian, Bernt Schiele, and Zeynep Akata. Zero-shot learning-the good, the bad and the ugly. In *CVPR*, 2017.
- [39] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *WACV*, 2014.
- [40] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In *RSS*, 2018.
- [41] Xinchun Yan, Jasmined Hsu, Mohammad Khansari, Yunfei Bai, Arkanath Pathak, Abhinav Gupta, James Davidson, and Honglak Lee. Learning 6-dof grasping interaction via deep geometry-aware 3d representations. In *ICRA*, 2018.
- [42] Xinchun Yan, Jimei Yang, Ersin Yumer, Yijie Guo, and Honglak Lee. Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision. In *NeurIPS*, 2016.
- [43] Jimei Yang, Scott E Reed, Ming-Hsuan Yang, and Honglak Lee. Weakly-supervised disentangling with recurrent transformations for 3d view synthesis. In *Advances in Neural Information Processing Systems*, pages 1099–1107, 2015.
- [44] Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Joshua B Tenenbaum, William T Freeman, and Jiajun Wu. Learning to Reconstruct Shapes from Unseen Classes. In *NeurIPS*, 2018.
- [45] Qingnan Zhou and Alec Jacobson. Thingi10k: A dataset of 10,000 3d-printing models. *arXiv preprint arXiv:1605.04797*, 2016.
- [46] Xingyi Zhou, Arjun Karapur, Linjie Luo, and Qixing Huang. Starmap for category-agnostic keypoint and viewpoint estimation. In *ECCV*, 2018.