

ShapeCodes: Self-Supervised Feature Learning by Lifting Views to Viewgrids

Dinesh Jayaraman^{1,2}, Ruohan Gao², and Kristen Grauman^{3,2}

¹ UC Berkeley

² UT Austin

³ Facebook AI Research

Abstract. We introduce an unsupervised feature learning approach that embeds 3D shape information into a single-view image representation. The main idea is a self-supervised training objective that, given only a single 2D image, requires all unseen views of the object to be predictable from learned features. We implement this idea as an encoder-decoder convolutional neural network. The network maps an input image of an unknown category and unknown viewpoint to a latent space, from which a deconvolutional decoder can best “lift” the image to its complete viewgrid showing the object from all viewing angles. Our class-agnostic training procedure encourages the representation to capture fundamental shape primitives and semantic regularities in a data-driven manner—without manual semantic labels. Our results on two widely-used shape datasets show 1) our approach successfully learns to perform “mental rotation” even for objects unseen during training, and 2) the learned latent space is a powerful representation for object recognition, outperforming several existing unsupervised feature learning methods.

1 Introduction

The field has made tremendous progress on object recognition by learning image features from supervised image datasets labeled by object categories [20, 23, 37]. Methods tackling today’s challenging recognition benchmarks like ImageNet or COCO [11, 42] capitalize on the 2D regularity of Web photos to discover useful appearance patterns and textures that distinguish many object categories. However, there are limits to this formula: manual supervision is notoriously expensive, not all objects are well-defined by their texture, and (implicitly) learning viewpoint-specific category models is cumbersome if not unscalable.

Restricting learned representations to a 2D domain presents a fundamental handicap. While visual perception relies largely on 2D observations, objects in the world are inherently three-dimensional entities. In fact, cognitive psychologists find strong evidence that inferring 3D geometry from 2D views is a useful skill in human perception. For example, in their seminal “mental rotation” work, Shepard and colleagues observed that people attempting to determine if two views portray the same abstract 3D shape spend time that is linearly




Fig. 1. Learning *ShapeCodes* by lifting views to “viewgrids”. Given one 2D view of an unseen object (possibly from an unseen category), our deep network learns to produce the remaining views in the viewgrid. This self-supervised learning induces a feature space for recognition. It embeds valuable cues about 3D shape regularities that transcend object category boundaries.

proportional to the 3D angular rotation between those views [55]. Such findings suggest that humans may explicitly construct mental representations of 3D shape from individual 2D views, and further, the act of mentally rotating such representations is integral to registering object views, and by extension, to object recognition.

Inspired by this premise, we propose an unsupervised approach to image feature learning that aims to “lift” individual 2D views to their 3D shapes. More concretely, we pose the feature learning task as one-shot viewgrid prediction from a single input view. A viewgrid—an array of views indexed by viewpoints—serves as an implicit image-based model of 3D shape. We implement our idea as an encoder-decoder deep convolutional neural network (CNN). Given one 2D view of any object from an arbitrary viewpoint, the proposed training objective learns *a latent space from which images of the object after arbitrary rotations are predictable*. See Fig 1. Our approach extracts this latent “ShapeCode” encoding to generate image features for recognition. Importantly, our approach is *class-agnostic*: it learns a single model to accommodate all objects seen during training, thereby encouraging the representation to capture basic shape primitives, semantic regularities, and shading cues. Furthermore, the approach is *self-supervised*: it aims to learn a representation generically useful to object perception, but without manual semantic labels.

We envision the proposed training phase as if an embodied visual agent were learning visual perception from scratch by inspecting a large number of objects. It can move its camera to arbitrary viewpoints around each object to acquire its own supervision. At test time, it must be able to observe only one view and hallucinate the effects of all camera displacements from that view. In doing so, it secures a representation of objects that, in a departure from purely 2D view-based recognition, is inherently shape-aware. We discuss the advantages of viewgrids over explicit voxels/point clouds in Sec. 3.1.

Our work relates to a growing body of work on self-supervised representation learning [1, 13, 18, 27, 29, 40, 46, 47, 60, 69], where “pretext” tasks requiring no manual labels are posed to the feature learner to instill visual concepts useful for recognition. We pursue a new dimension in this area—instilling 3D reasoning about shape. The idea of inferring novel viewpoints also relates to work in CNN-based view synthesis and 3D reconstruction [8, 14, 19, 36, 52, 59, 62, 67, 70]. However, our intent and assumptions are quite different. Unlike our approach, prior work targets reconstruction as the end task itself (not recognition), builds

category-specific models (e.g., one model for chairs, another for cars), and relies on networks pretrained with heavy manual supervision.

Our experiments on two widely used object/shape datasets validate that (1) our approach successfully learns class-agnostic image-based shape reconstruction, generalizing even to categories that were not seen during training and (2) the representations learned in the process transfer well to object recognition, outperforming several popular unsupervised feature learning approaches. Our results establish the promise of explicitly targeting 3D understanding as a means to learn useful image representations.

2 Related Work

Unsupervised representation learning While supervised “pretraining” of CNNs with large labeled datasets is useful [20], it comes at a high supervision cost and there are limits to its transferability to tasks unlike the original labeled categories. As an increasingly competitive approach, researchers investigate *unsupervised* feature learning [1, 4, 9, 13, 18, 27, 29, 40, 46, 47, 60, 61, 69]. An emerging theme is to identify “pretext” tasks, where the network targets an objective for which supervision is inherently free. In particular, features tasked with being predictive of contextual layout [13, 46, 47], camera egomotion [1, 27, 48], stereo disparities [17], colorization [40], or temporal slowness [18, 29, 60] simultaneously embed basic visual concepts useful for recognition. Our approach shares this self-supervised spirit and can be seen as a new way to force the visual learner to pick up on fundamental cues. In particular, our method expands this family of methods to multi-view 3D data, addressing the question: does learning to infer 3D from 2D help perform object recognition? While prior work considers 3D egomotion [1, 27], it is restricted to impoverished glimpses of the 3D world through “unattached” neighboring view pairs from video sequences. Instead, our approach leverages viewgrid representations of full 3D object shape. Our experiments comparing to egomotion self-supervision show our method’s advantage.

Recognition of 3D objects Though 2D object models dominate recognition in recent years (e.g., as evidenced in challenges like PASCAL, COCO, ImageNet), there is growing interest in grounding object models in their 3D structures and shapes. Recent contributions in large-scale data collections are fostering such progress [5, 64, 65], and researchers are developing models to integrate volumetric and multiview approaches effectively [49, 58], as well as new ideas for relating 3D properties (pose, occlusion boundaries) to 2D recognition schemes [66]. Active recognition methods reason about the information value of unseen views of an object [28, 30, 34, 51, 53, 64].

Geometric view synthesis For many years, new view synthesis was solved with geometry. In image-based rendering, rather than explicitly constructing a 3D model, new views are generated directly from multiple 2D views [35], with methods that establish correspondence and warp pixels according to projective or

multi-view geometry [2, 54]. Image-based models for object shape (implicitly) intersect silhouette images to carve a visual hull [39, 45].

Learning 2D-3D relationships More recently, there is interest in instead *learning* the connection between a view and its underlying 3D shape. The problem is tackled on two main fronts: image-based and volumetric. Image-based methods infer the new view as a function of a specified viewpoint. Given two 2D views, they learn to predict intermediate views [12, 16, 24, 33]. Given only a single view, they learn to render the observed object from new camera poses, e.g., via disentanglement with deep inverse graphic networks [38], tensor completion [7], recurrent encoder-decoder nets [31, 68], appearance flow [70], or converting partial RGBD to panoramas [57]. Access to synthetic object models is especially valuable to train a generative CNN [14]. Volumetric approaches instead map a view(s) directly to a 3D representation of the object, such as a voxel occupancy grid or a point cloud, e.g., with 3D recurrent networks [8], direct prediction of 3D points [15], or generative embeddings [19]. While most efforts study synthetic 3D object models (e.g., CAD datasets), recent work also ventures into real-world natural imagery [36]. Beyond voxels, inferring depth maps [59] or keypoint skeletons [62] offer valuable representations of 3D structure.

Our work builds on such advances in learning 2D-3D ties, and our particular convolutional autoencoder (CAE)-based pipeline (Sec. 3.2) echoes the de facto standard architecture for pixel-output tasks [38, 44, 59, 67, 69]. However, our goal differs from any of the above. Whereas existing methods develop category-specific models (e.g., chairs, cars, faces) and seek high-quality images/voxels as the end product, we train a class-agnostic model and seek a transferable image representation for recognition.

3 Approach

Our goal is to learn a representation that lifts a single image from an arbitrary (unknown) viewpoint and arbitrary class to a space where the object’s 3D shape is predictable—its *ShapeCode*. This task of “mentally rotating” an object from its observed viewpoint to arbitrary relative poses requires 3D understanding from single 2D views, which is valuable for recognition. By training on a one-shot shape reconstruction task, our approach aims to learn an image representation that embeds this 3D understanding and applies the resulting embedding for single-view recognition tasks.

3.1 Task setup: One-shot viewgrid prediction

During training, we first evenly sample views from the viewing sphere around each object. To do this, we begin by selecting a set S_{az} of M camera azimuths $S_{az} = \{360^\circ/M, 720^\circ/M, \dots, 360^\circ\}$ centered around the object. Then we select a set S_{el} of N camera elevations $S_{el} = \{0^\circ, \pm 180^\circ/(N-1), \pm 360^\circ/(N-1), \dots, \pm 90^\circ\}$. We now sample all $M \times N$ views of every object corresponding to the cartesian

product $S = S_{az} \times S_{el}$ of azimuth and elevation positions: $\{\mathbf{y}(\boldsymbol{\theta}_i) : \boldsymbol{\theta}_i \in S\}$.⁴ Note that each $\boldsymbol{\theta}_i$ is an elevation-azimuth pair, and represents one position in the viewing grid S .

Now, with these evenly sampled views, the one-shot viewgrid prediction task can be formulated as follows. Suppose the observed view is at an unknown camera position $\boldsymbol{\theta}$ sampled from our viewing grid set S of camera positions. The system must learn to predict the views $\mathbf{y}(\boldsymbol{\theta}')$ at position $\boldsymbol{\theta}' = \boldsymbol{\theta} + \boldsymbol{\delta}_i$ for all $\boldsymbol{\delta}_i \in S$. Because of the even sampling over the full viewing sphere, $\boldsymbol{\theta}'$ is itself also in our original viewpoint set S , so we have already acquired supervision for all views that our system must learn to predict.

Why viewgrids? The viewgrid representation has advantages over other more explicit 3D representations such as point clouds [15] and voxel grids [8]. First, viewgrid images can be directly acquired by an embodied agent through object manipulation or inspection, whereas voxel grids and point clouds require noisy 3D inference from large image collections. While our experiments leverage realistic 3D object CAD models to render viewpoints on demand (Sec. 4), it is actually less dependent on CAD data than prior work requiring voxel supervision. Ultimately we envision training taking place in a physical scenario, where an embodied agent builds up its visual representation by examining various objects. By moving to arbitrary viewpoints around an object, it acquires self-supervision to understand the 3D shape. Finally, viewgrids facilitate the representation of missing data—if some ground truth views are unavailable for a particular object, the only change required in our training loss (below in Eq 1) would be to drop the terms corresponding to unseen views.

3.2 Network architecture and training

To tackle the one-shot viewgrid prediction task, we employ a deep feed-forward neural network. Our network architecture naturally splits into four modular sub-networks with different functions: elevation sensor, image sensor, fusion, and finally, a decoder. Together, the elevation sensor, image sensor, and fusion modules process the observation and proprioceptive camera elevation information to produce a single feature vector that encodes the full object model. That vector space constitutes the learned *ShapeCode* representation. During training *only*, the decoder module processes this code through a series of learned deconvolutions to produce the desired image-based viewgrid reconstruction at its output.

Encoder: First, the image sensor module embeds the observed view through a series of convolutions and a fully connected layer into a vector. In parallel, the camera elevation angle is processed through the elevation sensor module. Note that the object pose is not fully known—while camera elevation can be determined from gravity cues, there is no way to determine the azimuth.

The outputs of the image and elevation sensor modules are concatenated and passed through a fusion module which jointly processes their information to

⁴ omitting object indices throughout to simplify notation.




Fig. 2. Architecture of our system. A single view of an object (top left) and the corresponding camera elevation (bottom left) are processed independently in “image sensor” and “elevation sensor” neural net modules, before fusion to produce the *ShapeCode* representation of the input, which embeds the 3D object shape aligned to the observed view. This is now processed in a deconvolutional decoder. During training, the output is a sequence of images representing systematically shifted viewpoints relative to the observed view. During testing, novel 2D images are lifted into the *ShapeCode* representation to perform recognition.

produce a $D = 256$ -dimensional output “code” vector, which embeds knowledge of object shape. In short, the function of the encoder is to lift a 2D view to a single vector representation of the full 3D object shape.

Decoder: To learn a representation with this property, the output of the encoder is processed through another fully connected layer to increase its dimensionality before reshaping into a sequence of small 4×4 feature maps. These maps are then iteratively upsampled through a series of learned deconvolutional layers. The final output of the decoder module is a sequence of MN output maps $\{\hat{y}_i : i = 1, \dots, M \times N\}$ of same height and width as the input image. Together these MN maps represent the system’s output viewgrid on which the training loss is computed.

The complete architecture, together with more detailed specifications, is visualized in Fig 2. Our convolutional encoder-decoder [44] neural network architecture is similar to [31, 59, 67, 70]. As discussed above, however, the primary focus of our work is very different. We consider one-shot reconstruction as a path to useful image representations that lift 2D views to 3D, whereas existing work addresses the image/voxel generation task itself, and accordingly builds category-specific models.

By design, our approach must not exploit artificial knowledge about the absolute orientation of an object it inspects, either during training or testing. Thus, there is an important issue to deal with—what is the correspondence between the individual views in the target viewgrid and the system’s output maps? Recall that at test time, the system will be presented with a single view of a novel object, from an unknown viewpoint (elevation known, azimuth unknown). How then can it know the correct viewpoint coordinates for the viewgrid it must produce? It instead produces viewgrids aligned with the *observed* viewpoint at the azimuthal coordinate origin, similar to [31]. Azimuthal rotations of a given viewgrid all form an equivalence class. In other words, circularly shifting the

7×12 viewgrid in Fig 2 by one column will produce a different, but entirely valid viewgrid representation of the same airplane object.

To optimize the entire pipeline, we regress to the target viewgrid \mathbf{y} , which is available for each training object. Since our output viewgrids are aligned by the observed view, we must accordingly shift the target viewgrid before performing regression. This leads to the following minimization objective:

$$\mathcal{L} = \sum_{i=1}^{M \times N} \|\hat{\mathbf{y}}_i - \mathbf{y}(\boldsymbol{\theta} + \boldsymbol{\delta}_i)\|^2, \quad (1)$$

where we omit the summation over the training set to keep the notation simple. Each output map $\hat{\mathbf{y}}_i$ is thus penalized for deviation from a specific *relative* circular shift $\boldsymbol{\delta}_i$ from the observed viewpoint $\boldsymbol{\theta}$. This one-shot reconstruction task enforces that the encoder must capture the *full* 3D object shape from observing just one 2D view. A similar reconstruction loss was proposed in [31] in a different context; they learned exploratory action policies by training them to select a sequence of views best suited to reconstruct entire objects and scenes.

Recent work targeting image synthesis has benefited from using adversarial (GAN) losses [21]. GAN losses help achieve correct low-level statistics in image patches, improving photorealism [26]. Rather than targeting realistic image synthesis as an end in itself, we target shape reconstruction for *feature* learning, so we use the standard ℓ_2 loss. During feature transfer, we discard the decoder entirely (see Sec. 3.3). See Supp for optimization details.

3.3 ShapeCode features for object recognition

During training, the objective is to minimize viewgrid error, to learn the latent space from which unseen views are predictable. Then, to apply our network to novel examples, the representation of interest is that same latent space output by the fusion module of the encoder—the ShapeCode. In the spirit of self-supervised representation learning, we hypothesize that features trained in this manner will facilitate high-level visual recognition tasks. This is motivated by the fact that in order to solve the reconstruction task effectively, the network must implicitly learn to lift 2D views of objects to inferred 3D shapes. A full 3D shape representation has many attractive properties for generic visual tasks. For instance, pose invariance is desirable for recognition; while difficult in 2D views, it becomes trivial in a 3D representation, since different poses correspond to simple transformations in the 3D space. Furthermore, the ShapeCode provides a representation that is *equivariant* to egomotion transformations, which is known to benefit recognition and emerges naturally in supervised networks [10, 24, 32, 41].

Suppose that the visual agent has learned a model as above for viewgrid prediction by inspecting 3D shapes. Now it is presented with a new recognition task, as encapsulated by a dataset of class-labeled training images *from a disjoint set of object categories*. We aim to transfer the 3D knowledge acquired in the one-view reconstruction task to this new task. Specifically, for each new

class-labeled image, we directly represent it in the feature space represented by an intermediate fusion layer in the network trained for reconstruction. These features are then input to a generic machine learning pipeline that is trained for the categorization task.

Recall that the output of the fusion module in Fig 2, which is the fc3 feature vector, is trained to encode 3D shape. In our experiments, we test the usefulness of features from fc3 and its two immediate preceding layers, fc2, and fc1, for solving object classification and retrieval tasks.

4 Experiments

First, we quantify performance for class-agnostic viewgrid completion (Sec 4.2). Second, we evaluate the learned features for object recognition (Sec 4.3).

4.1 Datasets

In principle, our self-supervised learning approach can leverage viewpoint-calibrated viewgrids acquired by an agent systematically inspecting objects in its environment. In our experiments, we generate such viewgrids from datasets of synthetic object shapes. We test our method on two such publicly available datasets: ModelNet [65] and ShapeNet [6]. Both of these datasets provide a large number of manually generated 3D models, with class labels. For each object model, we render 32×32 grayscale views from a grid of viewpoints that is evenly sampled over the viewing sphere centered on the object.

ModelNet [65] has 3D CAD models downloaded from the Web, and then manually aligned and categorized. ModelNet comes with two standard subsets: ModelNet-10 and ModelNet-40, with 10 and 40 object classes respectively. The 40 classes in ModelNet-40 include the 10 classes in ModelNet-10. We use the 10 ModelNet-10 classes as our *unseen classes*, and the other 30 ModelNet-40 classes as *seen classes*. We use the standard train-test split, and set aside 20% of seen-class test set models as validation data. ModelNet is the most widely used dataset in recent 3D object categorization work [28, 30, 34, 50, 58, 63, 65].

ShapeNet [6] contains a large number of models organized into semantic categories under the WordNet taxonomy. All models are consistently aligned to fixed canonical viewpoints. We use the standard ShapeNetCore-v2 subset which contains 55 diverse categories. Of these, we select the 30 largest categories as seen categories, and the remaining 25 are unseen. We use the standard train-test split. Further, since different categories have highly varying numbers of object instances, we limit each category in our seen-class training set to 500 models, to prevent training from being dominated by models of a few very common categories. Table 1 (left) shows more details for both datasets.

4.2 Class-agnostic one-shot viewgrid prediction

First, we train and test our method for viewgrid prediction. For both datasets, the system is trained on the seen-classes training set. The trained model is subsequently tested on both seen and unseen class test sets.

The evaluation metric is the per-pixel mean squared deviation of the inferred viewgrid vs. the ground truth viewgrid. We compare to several baselines:

- **Avg view:** This baseline simply predicts, at each viewpoint in the viewgrid, the average of all views observed in the training set *over all viewpoints*.
- **Avg viewgrid:** Both ModelNet and ShapeNet have consistently aligned models, so there are significant biases that can be exploited by a method that has access to this canonical alignment information. This baseline aims to exploit this bias by predicting, at each viewpoint in the viewgrid, the average of all views observed in the training set *at that viewpoint*. Note that our system does **not** have access to this alignment information, so it cannot exploit this bias.
- **GT class avg view:** This baseline represents a model with perfect object classification. Given an arbitrary object from some ground truth category, this baseline predicts, at each viewpoint, the average of all views observed in the training set for that category.
- **GT class avg viewgrid:** This baseline is the same as GT category avg view, but has knowledge of canonical alignments too, so it produces the average of views observed at each viewpoint over all models in that category in the training set.
- **Ours w. CA:** This baseline is our approach but trained with the (unrealistic) addition of knowledge about canonical alignment (“CA”) of viewgrids. It replaces Eq (1) to instead optimize the loss: $\mathcal{L} = \sum_{i=1}^{M \times N} (\hat{\mathbf{y}}_i - \mathbf{y}(\delta_i))^2$, so that each output map $\hat{\mathbf{y}}_i$ of the system is now assigned to a specific coordinate in the canonical viewgrid axes.

A key question in these experiments is whether the class-agnostic model we train can generalize to predict unseen views of objects from *classes* never seen during training.⁵

Table 1 (right) shows the results. “Avg viewgrid” and “GT category avg viewgrid” improve by large margins over “Avg view” and “GT category avg view”, respectively. This shows that viewgrid alignment biases can be useful for reconstruction in ModelNet and ShapeNet. Recall however that while these baselines can exploit the (unrealistic) bias, our approach cannot; it knows only the elevation as sensed from gravity. Our approach is trained to produce views at various *relative* displacements from the observed view, i.e., its target is a viewgrid that has the current view at its origin. So it cannot learn to memorize and align an average viewgrid.

Despite this, our approach outperforms the baselines by large margins. It even outperforms its variant “Ours w. CA” that uses alignment biases in the data. Why is “Ours w. CA” weaker? It is trained to produce viewgrids with

⁵ For this very reason, it is not clear how to map existing view synthesis models that are category-specific (e.g., chairs in [14], chairs/cars in [70]) to our class-agnostic setting so that they could compete fairly.

Dataset→	ModelNet		ShapeNet		Dataset→	ModelNet		ShapeNet	
Methods↓/ Data→	seen	unseen	seen	unseen	Methods↓/ Data→	seen	unseen	seen	unseen
Cam elevations	0,±30°,±60°,±90°		0,±30°,±60°		Avg view	13.514	15.956	14.793	16.394
Cam azimuths	0,30°,60°,... ,330°		0,45°,±315°		Avg vgrid	12.954	15.725	14.334	15.942
View size	32×32		32×32		GT class avg view	11.006	-	12.279	-
Categories	30	10	30	25	GT class avg vgrid	8.891	-	9.374	-
Train models	5,852	-	11,532	-	Ours w. CA	4.689	10.440	5.879	9.021
Val models	312	-	1,681	-	Ours	3.718	7.005	4.656	6.811
Test models	1,248	726	3,569	641					

Table 1. Left: Dataset statistics. Right: Quantitative results for viewgrid completion. Results are reported as $\text{MSE} \times 1000$ on images normalized to lie in $[0, 1]$, for seen and unseen classes. Lower is better. Per-category results are shown in Supp. “GT category” methods are not applicable to unseen classes since they rely on category knowledge.

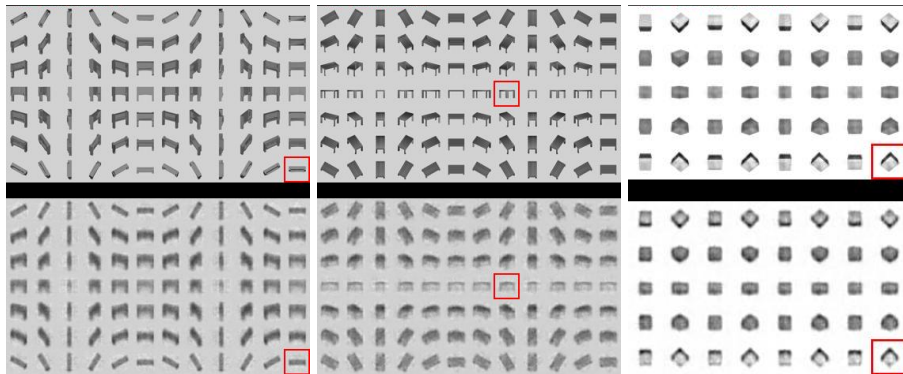


Fig. 3. Shape reconstructions from a single view (rightmost example from ShapeNet, other two from ModelNet). In each panel, ground truth viewgrids are shown at the top, the observed view is marked with a red box, and our method’s reconstructions are shown at the bottom. (Best seen in pdf at high resolution.) See text in Sec 4.2 for description of these examples.

canonical alignments (CA). CA’s are loosely manually defined, usually class-specific conventions in the dataset (e.g., 0° azimuth and elevation for all cars might be “head on” views). “Ours w. CA” naturally has no notion of CA’s for unseen categories, where it performs particularly poorly. On seen classes with strong alignment biases, CA’s make it easier to capture category-wide information (e.g., if the category is recognizable, produce its corresponding average aligned training viewgrid as output). However, it is harder to capture instance-specific details, since the network must not only mentally rotate the input view but also infer its canonical pose correctly. “Ours” does better by aligning outputs to the observed viewpoint.

Fig 3 shows example viewgrids generated by our method. In the leftmost panel, it reconstructs an object shape from a challenging viewpoint, effectively exploiting the semantic structure in ModelNet. In the center panel, the system observes an ambiguous viewpoint that could be any one of four different views at the same azimuth. In response to this ambiguity, it attempts to play it safe




Fig. 4. ModelNet reconstruction MSE for three classes, conditioned on observed view (best viewed in pdf at high resolution). Yellower colors correspond to high MSE (bad) and bluer colors correspond to low MSE (good). See text in Sec 4.2.

to minimize MSE loss by averaging over possible outcomes, producing blurry views. In the rightmost panel, our method shows the ability to infer shape from shading cues for simple objects.

Fig 4 examines which views are informative for one-shot viewgrid prediction. For each of the three classes shown, the heatmap of MSE is overlaid on the average viewgrid for that class. The yellowish (high error) horizontal and vertical stripes correspond to angles that only reveal a small number of faces of the object. Top and bottom views are consistently uninformative, since very different shapes can have very similar overhead projections. The middle row (0° elev.) stands out as particularly bad for “airplane” because of the narrow linear projection, which presents very little information. See Supp. for more. These trends agree with intuitive notions of which views are most informative for 3D understanding, and serve as evidence that our method learns meaningful cues to infer unseen views.

Overall, the reconstruction results demonstrate that our approach successfully learns one *single unified category-agnostic* viewgrid reconstruction model that handles not only objects from the large number of generic categories that are represented in its training set, but also objects from unseen categories.

4.3 ShapeCode features for object recognition

We now validate our key claim: the lifted features—though learned without manual labels—are a useful visual representation for recognition.

First, as described in Sec 3.3, we extract features from various layers in the network (fc1, fc2, fc3 in Fig 2) and use them as inputs to a classifier trained for categorization of individual object views. Though any classifier is possible, following [18, 32, 60], we employ a simple k -nearest neighbor classifier, which makes the power of the underlying representation most transparent. We run this experiment on both the seen and unseen class subsets on both ModelNet and ShapeNet. In each case, we use 1000 samples per class in the training set, and set $k = 5$. We compare our features against a variety of baselines:

- **DrLIM [22]:** A commonly used unsupervised feature learning approach. During training, DrLIM learns an invariant feature space by mapping features of views of the same training object close to each other, and pushing features of views of different training objects far apart from one another.

- **Autoencoder** [3, 25, 44]: A network is trained to observe an input view from an arbitrary viewpoint and produce exactly that same view as its output (compared to our method which produces the full viewgrid, including views from *other* viewpoints too). For this method, we use an architecture identical to ours except at the very last deconvolutional layer, where, rather than producing $N \times M$ output maps, it predicts just one map corresponding to the observed view itself.
- **Context** [47]: Representative of the popular paradigm of exploiting spatial context to learn features [13, 46, 47]. The network is trained to “inpaint” randomly removed masks of arbitrary shape covering up to $\frac{1}{4}$ of the 32×32 object views, thus learning spatial configurations of object parts. We adapt the public code of [47].
- **Egomotion** [1]: Like our method, this baseline also exploits camera motion to learn unsupervised representations. While our method is trained to predict all rotated views given a starting view, [1] trains to predict the camera rotation between a given pair of images. We train the model to predict 8 classes of rotations, i.e., the immediately adjacent viewpoints in the viewgrid for a given view (3×3 neighborhood).
- **PointSetNet** [15]: This method reconstructs object shape point clouds from a single image, plus the ground truth segmentation mask. We extract features from their provided encoder network trained on ShapeNet. Since segmentation masks are unavailable in the feature evaluation setting, we set them to the whole image.
- **3D-R2N2** [8]: This method constructs a voxel grid from a single view. We extract features from their provided encoder network trained on ShapeNet.
- **VGG** [56]: While our focus is unsupervised feature learning, this baseline represents current standard *supervised* features, trained on millions of manually labeled images. We use the VGG-16 architecture [56] trained on ImageNet, and extract fc6 features from 224×224 images.
- **Shape classifier**: To provide a supervised baseline trained on in-domain data, we train a network for single-view 3D shape categorization using 1k labeled images per seen class. The architecture is kept identical to the encoder of our method, and features are extracted from the same layers.
- **Pixels**: For this baseline the 32×32 image is vectorized and used directly as a feature vector.
- **Random weights**: A network with identical architecture to ours and initialized with the same scheme is used to extract features with no training.

The “Random weights”, “DrLIM”, “Egomotion”, and “Autoencoder” methods use identical architectures to ours until fc3 (see Supp). For “Context”, we stay close to the authors’ architecture and retrain on our 3D datasets. For “VGG”, “PointSetNet”, and “3D-R2N2”, we use author-provided model weights.

Recall that our model is trained to observe camera elevations together with views, as shown in Fig 2. While this is plausible in a real world setting where an agent may know its camera elevation angle from gravity cues, for fair comparison with our baselines, we omit proprioception inputs when evaluating our unsupervised features. Instead, we feed in camera elevation 0° for all views.

Datasets→ Methods↓ / Classes→	ModelNet		ShapeNet	
	seen	unseen	seen	unseen
Chance	3.3	10.0	3.3	4.0
VGG [56] (supervised)	66.0	64.9	55.9	53.7
Shape classifier (supervised)	62.6	68.1	52.8	49.9
Pixels	52.5	60.7	43.1	44.9
Random weights	49.6	59.4	39.6	39.7
DrLIM [22]	57.4	64.9	47.5	47.2
Autoencoder [3, 25, 44]	52.5	60.8	44.3	46.0
Context [47]	52.6	60.5	46.2	46.5
Egomotion [1]	56.1	65.0	49.0	49.7
PointSetNet [15]	35.5	38.8	28.6	32.2
3D-R2N2 [8]	49.4	55.5	39.0	41.2
Ours w. CA	64.0	69.6	56.9	54.5
Ours	65.2	71.2	57.7	54.8

Table 2. Single view recognition accuracy (%) with features from our model vs. baselines on the ModelNet and ShapeNet datasets. For each method, we report its best accuracy across layers (fc1, fc2, fc3). Results are broken out according to classes seen and unseen during representation learning. Our approach consistently outperforms other unsupervised representations, and even competes well with off-the-shelf VGG features pretrained with 1M ImageNet labels.

Table 2 shows the results for both datasets. Since trends across fc1, fc2, and fc3 were all very similar, for each method, we report the accuracy from its best performing layer (see Supp for per-layer results). “Ours” strongly outperforms all prior approaches. Critically, our advantage holds whether or not the objects to be recognized are seen during training of the viewgrid prediction network.

Among the baselines, all unsupervised learning methods outperform “Pixels” and “Random weights”, as expected. The two strongest unsupervised baselines are “Egomotion” [1] and “DrLIM” [22]. Recall that “Egomotion” is especially relevant to our approach as it also has access to relative camera motion information. However, while their method only sees neighboring view pairs sampled from the viewgrid at training time, our approach learns to infer full viewgrids for each instance, thus exploiting this information more effectively. “Autoencoder” features perform very poorly. “Context” features [47] are also poor, suggesting that spatial context within views, though successful for 2D object appearance [47], is a weak learning signal for 3D shapes.

The fact that our method outperforms PointSetNet [15] and 3D-R2N2 [8] suggests that it is preferable to train for generating implicit 3D viewgrids rather than explicit 3D voxels or point clouds. While we include them as baselines for this very purpose, we stress that the goal in those papers [8, 15] is reconstruction—not recognition—so this experiment expands those methods’ application beyond the authors’ original intent. We believe their poor performance is partly due to domain shift: ModelNet performance is weaker for both methods since they were trained on ShapeNet. Further, the authors train PointSetNet only on objects with specific poses (elevation 20°) and it exploits ground truth segmentation masks which are unavailable in this setting. Both methods were also trained on object views rendered with small differences from ours (see Supp).

Finally, Table 2 shows our *self-supervised* representation outperforms the ImageNet-pretrained *supervised* VGG features in most cases. Note that recognition tasks on synthetic shape datasets are commonly performed with ImageNet-pretrained neural networks [28, 30, 34, 58]. However, the domain gap between Im-




Fig. 5. ShapeCodes embedding for data from unseen ModelNet10 classes, visualized with t-SNE in two ways: (left) categories are color-coded, and (right) instances are shown at their t-SNE embedding. Best viewed in pdf with zoom.

ageNet and ModelNet/ShapeNet is not the sole reason for this; our method also beats supervised features trained on in-domain data (“Shape classifier”). These results illustrate that our geometry-aware self-supervised pretraining has potential to supersede traditional ImageNet pretraining, and even domain-specific supervised pretraining.

Fig 5 visualizes a t-SNE [43] embedding of unseen ModelNet-10 class images using ShapeCode features. As can be seen, categories tend to cluster, including those with diverse appearance like ‘chair’. ‘Chair’ and ‘toilet’ are close, as are ‘dresser’ and ‘night stand’, showing the emergence of high-level semantics.

In Supp, we test ShapeCode’s transferability to recognition tasks under many varying experimental conditions, including varying training dataset sizes for the k-NN classifier, and performing object category retrieval rather than categorization. A unified and consistent picture emerges: ShapeCode is a significantly better feature representation than the baselines for high-level recognition tasks.

5 Conclusions

We proposed ShapeCodes, self-supervised features trained for a mental rotation task so that they embed generic 3D shape priors useful to recognition. ShapeCodes outperform an array of state-of-the-art approaches for unsupervised feature learning, establishing the promise of explicitly targeting 3D understanding to learn useful image representations.

While we test our approach on synthetic object models, we are investigating whether features trained on synthetic objects could generalize to real images. Further, an embodied agent could in principle inspect physical objects to acquire viewgrids to allow training on real objects. Future work will explore extensions to permit sequential accumulation of observed views of real objects. We will also investigate reconstruction losses expressed at a more abstract level than pixels, e.g., in terms of a feature content loss.

Acknowledgements: This research is supported in part by DARPA Lifelong Learning Machines, ONR PECASE N00014-15-1-2291, an IBM Open Collaborative Research Award, and Berkeley DeepDrive.

References

1. Agrawal, P., Carreira, J., Malik, J.: Learning to see by moving. In: ICCV (2015)
2. Avidan, S., Shashua, A.: Novel view synthesis in tensor space. In: CVPR (1997)
3. Bengio, Y., et al.: Learning deep architectures for ai. *Foundations and trends® in Machine Learning* **2**(1), 1–127 (2009)
4. Bojanowski, P., Joulin, A.: Unsupervised learning by predicting noise. ICML (2017)
5. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: ShapeNet: An Information-Rich 3D Model Repository. Tech. Rep. arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago (2015)
6. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
7. Chen, C.Y., Grauman, K.: Inferring unseen views of people. In: CVPR (2014)
8. Choy, C., Xu, D., Gwak, J., Chen, K., Savarese, S.: 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In: ECCV (2016)
9. Coates, A., Ng, A., Lee, H.: An analysis of single-layer networks in unsupervised feature learning. In: Proceedings of the fourteenth international conference on artificial intelligence and statistics. pp. 215–223 (2011)
10. Cohen, T.S., Welling, M.: Transformation properties of learned visual representations. ICLR (2014)
11. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
12. Ding, W., Taylor, G.: “mental rotation” by optimizing transforming distance. In: NIPS Workshop (2014)
13. Doersch, C., Gupta, A., Efros, A.: Unsupervised visual representation learning by context prediction. In: ICCV (2015)
14. Dosovitskiy, A., Springenberg, J., Brox, T.: Learning to generate chairs with convolutional neural networks. In: CVPR (2015)
15. Fan, H., Su, H., Guibas, L.: A point set generation network for 3d object reconstruction from a single image. In: Conference on Computer Vision and Pattern Recognition (CVPR). vol. 38 (2017)
16. Flynn, J., Neulander, I., Philbin, J., Snavely, N.: Deepstereo: Learning to predict new views from the world’s imagery. In: CVPR (2015)
17. Gan, C., Gong, B., Liu, K., Su, H., Guibas, L.J.: Geometry guided convolutional neural networks for self-supervised video representation learning. In: CVPR (2018)
18. Gao, R., Jayaraman, D., Grauman, K.: Object-centric representation learning from unlabeled videos. In: ACCV (2016)
19. Girdhar, R., Fouhey, D., Rodriguez, M., Gupta, A.: Learning a predictable and generative vector representation for objects. In: ECCV (2016)
20. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2014)
21. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS (2014)
22. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality Reduction by Learning an Invariant Mapping. CVPR (2006)

23. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
24. Hinton, G., Krizhevsky, A., Wang, S.: Transforming auto-encoders. In: ICANN (2011)
25. Hinton, G.E., Salakhutdinov, R.R.: Reducing the dimensionality of data with neural networks. *science* **313**(5786), 504–507 (2006)
26. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017)
27. Jayaraman, D., Grauman, K.: Learning image representations tied to ego-motion. In: ICCV (2015)
28. Jayaraman, D., Grauman, K.: Look-ahead before you leap: End-to-end active recognition by forecasting the effect of motion. In: ECCV (2016)
29. Jayaraman, D., Grauman, K.: Slow and steady feature analysis: Higher order temporal coherence in video. In: CVPR (2016)
30. Jayaraman, D., Grauman, K.: End-to-end policy learning for active visual categorization. In: TPAMI (2018)
31. Jayaraman, D., Grauman, K.: Learning to look around: Intelligently exploring unseen environments for unknown tasks. CVPR (2018)
32. Jayaraman, D., Grauman, K.: Learning image representations tied to egomotion from unlabeled video. *International Journal of Computer Vision* (2017)
33. Ji, D., Kwon, J., McFarland, M., Savarese, S.: Deep view morphing. In: CVPR (2017)
34. Johns, E., Leutenegger, S., Davison, A.: Pairwise decomposition of image sequences for active multi-view recognition. In: CVPR (2016)
35. Kang, S.B.: A survey of image-based rendering techniques. *Videometrics SPIE Intl Symp on Elec Imag: Science and Technology* (1999)
36. Kar, A., Tulsiani, S., Carreira, J., Malik, J.: Category-specific object reconstruction from a single image. In: CVPR (2015)
37. Krizhevsky, A., Sutskever, I., Hinton, G.: Imagenet classification with deep convolutional neural networks. NIPS (2012)
38. Kulkarni, T., Whitney, W., Kohli, P., Tenenbaum, J.: Deep convolutional inverse graphics network. In: NIPS (2015)
39. Kutulakos, K., Seitz, S.: A theory of shape by space carving. *IJCV* (2000)
40. Larsson, G., Maire, M., Shakhnarovich, G.: Colorization as a proxy task for visual understanding. In: CVPR (2017)
41. Lenc, K., Vedaldi, A.: Understanding image representations by measuring their equivariance and equivalence. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 991–999 (2015)
42. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
43. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *JMLR* (2008)
44. Masci, J., Meier, U., Cireşan, D., Schmidhuber, J.: Stacked convolutional auto-encoders for hierarchical feature extraction. In: *International Conference on Artificial Neural Networks*. pp. 52–59. Springer (2011)
45. Matusik, W., Buehler, C., Raskar, R., Gortler, S., McMillan, L.: Image-based visual hulls. In: SIGGRAPH (2000)
46. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: ECCV (2016)
47. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: CVPR (2016)

48. Poier, G., Schinagl, D., Bischof, H.: Learning pose specific representations by predicting different views. In: CVPR (2018)
49. Qi, C., Su, H., Niessner, M., Dai, A., Yan, M., Guibas, L.: Volumetric and multi-view cnns for object classification on 3d data. In: CVPR (2016)
50. Qi, C.R., Su, H., Nießner, M., Dai, A., Yan, M., Guibas, L.J.: Volumetric and multi-view cnns for object classification on 3d data. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5648–5656 (2016)
51. Ramanathan, V., Pinz, A.: Active object categorization on a humanoid robot. In: VISAPP (2011)
52. Rezende, D., Eslami, S., Mohamed, S., Battaglia, P., Jaderberg, M., Heess, N.: Unsupervised learning of 3d structure from images. In: NIPS (2016)
53. Schiele, B., Crowley, J.: Transinformation for active object recognition. In: ICCV (1998)
54. Seitz, S., Dyer, C.: View morphing. In: SIGGRAPH (1996)
55. Shepard, R.N., Metzler, J.: Mental rotation of three-dimensional objects. *Science* **171**, 701–703 (1971)
56. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. ICLR (2015)
57. Song, S., Zeng, A., Chang, A., Savva, M., Savarese, S., Funkhouser, T.: Im2pano3d: Extrapolating 360 structure and semantics beyond the field of view. In: CVPR (2018)
58. Su, H., Maji, S., Kalogerakis, E., Learned-Miller, E.: Multiview convolutional neural networks for 3d shape recognition. In: ICCV (2015)
59. Tatarchenko, M., Dosovitskiy, A., Brox, T.: Multi-view 3d models from single images with a convolutional network. In: ECCV (2016)
60. Wang, X., Gupta, A.: Unsupervised learning of visual representations using videos. In: ICCV (2015)
61. Wang, X., He, K., Gupta, A.: Transitive invariance for self-supervised visual representation learning. ICCV (2017)
62. Wu, J., Xue, T., Lim, J., Tian, Y., Tenenbaum, J., Torralba, A., Freeman, W.: Single image 3d interpreter network. In: ECCV (2016)
63. Wu, J., Zhang, C., Xue, T., Freeman, B., Tenenbaum, J.: Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In: NIPS (2016)
64. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: CVPR (2015)
65. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: CVPR (2015)
66. Xiang, Y., Choi, W., Lin, Y., Savarese, S.: Data-driven 3d voxel patterns for object category recognition. In: CVPR (2015)
67. Yan, X., Yang, J., Yumer, E., Guo, Y., Lee, H.: Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision. In: NIPS (2016)
68. Yang, J., Reed, S., Yang, M.H., Lee, H.: Weakly supervised disentangling with recurrent transformations in 3d view synthesis. In: NIPS (2015)
69. Zhang, R., Isola, P., Efros, A.: Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In: CVPR (2017)
70. Zhou, T., Tulsiani, S., Sun, W., Malik, J., Efros, A.: View synthesis by appearance flow. In: ECCV (2016)