

Self-Supervised Learning of Point Clouds via Orientation Estimation

Omid Poursaeed^{*1}

Tianxing Jiang^{*1}

Han Qiao^{*1}

Nayun Xu¹

Vladimir G. Kim²

¹Cornell University

²Adobe Research

Abstract

Point clouds provide a compact and efficient representation of 3D shapes. While deep neural networks have achieved impressive results on point cloud learning tasks, they require massive amounts of manually labeled data, which can be costly and time-consuming to collect. In this paper, we leverage 3D self-supervision for learning downstream tasks on point clouds with fewer labels. A point cloud can be rotated in infinitely many ways, which provides a rich label-free source for self-supervision. We consider the auxiliary task of predicting rotations that in turn leads to useful features for other tasks such as shape classification and 3D keypoint prediction. Using experiments on ShapeNet and ModelNet, we demonstrate that our approach outperforms the state-of-the-art. Moreover, features learned by our model are complementary to other self-supervised methods and combining them leads to further performance improvement.

1. Introduction

As a concise and expressive form of representing 3D shapes, point clouds are ubiquitous in various 3D vision tasks. Massive point cloud datasets have become available with the development of 3D scanning devices and techniques. ShapeNet [9], ModelNet [86] and ScanNet [11] are instances of such datasets. Models such as PointNet [58] have achieved impressive results on these datasets. However, manually annotating these large amounts of data is costly and impractical to scale. This raises the need for methods that can learn from unlabelled data to reduce the number of annotated samples required for learning. Recently, self-supervised learning have shown promising performance on several computer vision tasks. Self-supervision is a learning framework in which a supervised signal for a pretext task is created automatically, in an effort to learn representations that are useful for solving real-world downstream tasks.

A simple form of self-supervision is predicting rotations [20]. This approach has achieved state-of-the-art results on image classification [35]. In this paper we propose the first self-supervised approach based on rotations of 3D data. Among various 3D representations, we focus on point clouds as an efficient representation in which rotations can be easily applied via matrix multiplication. Manually annotating point clouds to provide information such as category labels or keypoints is cumbersome and costly. Notably, most public 3D models have a *canonical* orientation, or at least a known *up* axis. We can leverage this orientation information to provide a label-free supervision signal for downstream tasks. More specifically, we can arbitrarily rotate an input point cloud and train a model to recover the rotation angle. The auxiliary task of predicting these rotations requires high-level understanding of 3D shapes, since canonical orientation of the object implicitly encodes important information about its functionality [16].

We demonstrate that a model trained for this proxy task generates semantic features useful for target tasks such as classification and keypoint prediction. We show that simple models such as SVMs can be trained on these features and achieve strong performance even when trained on a small subset of the training data. Hence, our approach facilitates learning with fewer labels, and obviates the need for manually annotating massive datasets. Finally, we discuss how our approach can be combined with other self-supervised models to achieve further performance gain, indicating that our learned features provide complementary information to other methods.

2. Related Work

2.1. Deep Learning on Geometry

There has been a recent surge of interest on learning-based approaches on geometric data [6, 10, 64, 84, 86]. These models are able to encode rich prior information about the space of 3D shapes which helps to resolve ambiguities in the input. Typical convolutional architectures require highly regular input data formats such as image grids or 3D voxels. Unlike images, geometry usually does

* Indicates equal contribution

not have an underlying grid structure, requiring new building blocks replacing convolution and pooling or adaptation to a grid. As a simple way to overcome this issue, view-based [72, 83, 57] and volumetric representations [34, 43, 74, 86] or their combination [59] place geometric data onto a grid. Voxel representations are a straightforward generalization of pixels to the 3D case. However, the memory footprint of these representations grows cubically with resolution. Data adaptive representations such as octrees [66, 74] mitigate this issue, yet they lead to complex implementations and are still limited to relatively small voxel grids.

Point clouds provide a more efficient representation by only focusing on surface points, yielding a tensor with a fixed number of 3D coordinates [1, 15]. Point clouds as an input modality present a unique set of challenges when building a network architecture. The grid structure is not available in raw point clouds as unordered sets of points. To address this issue, Qi et al. propose PointNet [58] which respects permutation invariance of point clouds by using a max-pooling operation to form a single feature vector representing the global context. PointNet++ [60] extends PointNet by recursively applying it on a nested partitioning of the input point set. Kernel correlation and graph pooling are proposed for improving PointNet-like methods in [70]. PointCNN [39] leverages spatially-local correlations using a χ -Conv operator that weights and permutes input points and features before they are processed by a typical convolution. PointConv [85] and Relation-Shape CNN [41] also focus on local structures of point clouds and further improve the quality of captured features. Recently, graph convolutional networks have shown promising results on learning tasks on point clouds. Dynamic Graph CNN [81] applies a graph convolution to edges of the k-nearest neighbor graph of the point clouds, which is dynamically recomputed in the feature space after each layer. While the aforementioned models achieve great performance on various tasks, they require large number of annotated samples for training. Hence, there is a need for efficient methods with less supervision.

Learning-based approaches have been applied to several other geometric representations such as meshes [30, 62, 78, 56], implicit functions [44, 52, 87, 55] and structured models [38, 97, 96]. We refer the interested reader to [7] for a survey on geometric deep learning.

2.2. Unsupervised and Self-Supervised Learning

The goal of unsupervised learning is to learn representations from raw data without the use of labels such that they are easily adaptable for other tasks. A common approach in unsupervised learning is leveraging generative models such as Variational AutoEncoders (VAEs) [32, 33, 65], Generative Adversarial Networks (GANs) [21, 26, 61, 31] and

autoregressive models [22, 76, 77]. Another prominent paradigm is self-supervision, which defines an annotation-free pretext task in order to provide a surrogate supervision signal for feature learning. Example pretext tasks include colorizing gray scale images [36, 92, 93], image inpainting [54], noise-as-targets [5], generation [29, 63], geometry [14], counting [48] and clustering of images [8, 40, 88]. Other methods use additional information such as temporal data to learn features. Typical proxy tasks are based on temporal-context [45, 82, 68], spatio-temporal cues [27, 19, 80], foreground-background segmentation via video segmentation [53], optical-flow [18, 42], future-frame synthesis [71], audio prediction from video [12, 51], audio-video alignment [3], ego-motion estimation [28], slow feature analysis with higher order temporal coherence [28], transformation between frames [2] and patch tracking in videos [79].

Several papers consider predicting the spatial relation between two randomly chosen image patches [13, 46, 47, 49]. An extension of this idea to 3D representations is presented in [67]. Our work bears a resemblance to [20] which randomly rotates an image by one of four possible angles and let the model predict the rotation. [50] proposes to predict future patches in representation space via autoregressive predictive coding. An overview and comparison of various 2D self-supervised learning algorithms is provided in [35, 4].

Recently, various methods have been proposed for unsupervised learning of point clouds. Autoencoder-based approaches learn features through reconstructing the point cloud data [1, 17, 89, 94]. Generative models learn representations in the process of generating plausible point clouds [37, 73, 75, 84]. Concurrent to our work, various pretext tasks have been proposed for self-supervision on point clouds. Sauder et al. propose to learn features by reconstructing point clouds whose parts have been randomly rearranged [67]. Zhang et al. train deep graph neural networks to solve two proxy tasks, part contrasting and object clustering [91]. A multi-task learning framework is proposed in [24] which learns features by optimizing three different tasks including clustering, prediction, and reconstruction. Our approach requires more holistic understanding of shapes, and is relatively easy to implement. As we demonstrate in experiments, our model outperforms several competing methods, and can also be combined with them to achieve further performance gain.

3. Method

Inspired by the success of self-supervised learning on images, we propose a self-supervised method to learn representations of point clouds. We consider the proxy task of rotation prediction. A set of K rotations $R = \{R_1, R_2, \dots, R_K\}$ is used, where $R_i(X)$ rotates upward di-




Figure 1: Model architecture. A rotation angle is randomly sampled from the set of possible angles and applied to the input point cloud. The rotation prediction model is trained to infer the rotation angle from the rotated point cloud. Features from the rotation prediction model are passed to a linear SVM which is trained to infer the shape category.

rotation of the input point cloud X to the i^{th} angle. We train a classifier $F(\cdot)$ to predict the rotation applied to the input point cloud. The model takes a transformed point cloud $R_i(X)$, with R_i randomly selected from R , and outputs a probability distribution over all possible rotations. We train the model using cross-entropy loss in a supervised manner.

Figure 1 illustrates our architecture. Our models are based on PointNet [58] and DGCNN [81] in which the final layers are modified for rotation classification or regression. The classifier outputs K probabilities and the regression model outputs four real numbers (three for the rotation axis and one for the rotation angle). The features from the rotation prediction model are then used for several target tasks including classification and keypoint prediction. ShapeNet [9], a large scale dataset of 3D shapes, is used to train the rotation estimation model. 3D models in ShapeNet are annotated with upright and front orientations, allowing us to align them to a canonical orientation.

Number of angles K is a hyper-parameter. Using a small K gives a high accuracy on rotation prediction while might not yield the best features for downstream tasks since predicting a small number of rotations is a simple task. On the other hand, a very large K gives a low accuracy on rotation prediction as angles become closer to one other. Table 1 shows the rotation classification accuracy for different values of K . As we observe, the accuracy decreases for larger K values.

For a given choice of K we need to select the angles as uniformly as possible to avoid bias in learning. Ideally, we can consider a regular polyhedron with K vertices. However, regular polyhedra do not exist for all values of K , hence we resort to approximate solutions for custom K

Number of rotation angles K	Prediction accuracy
6	99.8%
18	89.0%
32	90.3%
54	67.8%
100	1.6%

Table 1: Rotation classification accuracy on ShapeNet based on the number of angles K .

values. We consider a range of small to large values for K including 6, 18, 32, 54 and 100. For six rotation angles, we choose the six directions of the $\pm x, \pm y, \pm z$ axes. For $K = 18$, we take directions of the six axes and the twelve angle bisectors for each two consecutive axes. For $K = 32$, a regular icosahedron is centered at the origin, and directions towards its 12 vertices and 20 face centers are chosen. For $K = 54, 100$, we use the golden spiral sunflower distribution, where the golden ratio is $\frac{1+\sqrt{5}}{2}$. Figure 2 illustrates the distribution of angles for various K values.

Alternatively, we can train a model to regress the angles instead of classifying them. We consider two rotation representations for the regression task: axis-angle and a continuous 6D representation presented in [95]. For the axis-angle representation, we first randomly select an axis by sampling coordinates $x, y, z \sim \mathcal{N}(0, 1)$ and normalizing the resulting vector $(x, y, z)^T$ to have a unit norm. It can be shown that the resulting vectors are uniformly distributed on the unit sphere. We then sample a random angle, and rotate the input point cloud using the sampled axis and angle. A re-




Figure 2: Illustration of the rotation angles. The object is placed at the origin, and its upward direction is rotated to one of the K marked points.

gression model is trained to regress the axis of rotation and the rotation angle with respect to this axis. We train the model with L_2 loss using a linear combination of axis and angle prediction losses with equal weights.

We also consider the representation proposed by Zhou et al. [95]. They demonstrate that for 3D rotations all representations are discontinuous in the real Euclidean spaces of four or fewer dimensions. They present continuous representations in 5D and 6D which are more suitable for learning. We use the 6D representation and consider a similar setup as their experiment on pose estimation for 3D point clouds. The network consists of a simplified PointNet structure for feature extraction and an MLP to produce the 6 dimensional rotation representation. A mapping function transforms this representation to $SO(3)$. Note that while they estimate the relative rotation between two point clouds, we predict the rotation of a single point cloud with respect to a canonical orientation.

Finally, we use the trained rotation estimation model to extract features for downstream tasks. These features are

passed to a linear SVM for shape classification. In the keypoint prediction task, we modify final layers of the angle prediction model and fine-tune it for keypoint regression. We train the rotation estimation and target models on ShapeNet and ModelNet respectively.

Note that our method assumes that the shapes have been consistently aligned, but this is a common assumption in online stock repositories. We can extend our approach to predict the relative rotation between point clouds either in a supervised or unsupervised manner. In the supervised case we use ground truth orientations, and it is quite similar to our current setup. In the unsupervised case we can align a point cloud to a rotated version of itself. We can estimate the relative rotation, apply it to one of the point clouds and use a chamfer loss to penalize misalignment. The latter case enables learning even in the absence of canonical orientations. However, the learned features might not be as transferable as current features. We will further explore this direction in future work.




Figure 3: Linear SVM’s classification accuracy on ModelNet-40 for different sizes of training set.

4. Experiments

We demonstrate that features learned by the rotation prediction model are useful for downstream tasks. We consider two tasks: shape classification and 3D keypoint prediction.

4.1. Shape Classification

We consider the task of shape classification on ModelNet [86]. Using the rotation prediction model trained on ShapeNet, we generate features for ModelNet-40’s train split. We use the global activations after the max pooling layers of PointNet and DGCNN as we found them to yield the best performance. A linear SVM classifier is trained on these activations for the task of predicting the shape category. This experiment evaluates the learned features in a transfer learning task, demonstrating their generalizability. Figure 1 illustrates the architecture. Note that the rotation estimation model is trained on the ShapeNet dataset while the linear SVM is trained on features from ModelNet.

Table 2 shows accuracy of the linear SVM on the test set of ModelNet-40. It also compares our model’s performance with various competing methods. As we observe, our approach outperforms competing methods. Note that the accuracy depends on the number of angles, with $K = 18, 32$ achieving the best performance. Using too large or too small values for K leads to inferior performance as the rotation prediction task becomes too difficult or too easy.

We can also combine our approach with other self-supervised learning methods. More specifically, we concatenate the features learned by our model with features extracted from the context prediction model [67] and train the SVM classifier on the combined features. As shown in Table 2, this leads to improved performance compared to indi-

Previous work	
VConv-DAE [69]	75.50%
3D-GAN [84]	83.30%
Latent-GAN [1]	85.70%
FoldingNet [89]	88.40%
VIP-GAN [23]	90.19%
Context Prediction (DGCNN) [67]	90.64%
Context Prediction (PointNet) [67]	87.31%
Ours (DGCNN)	
6 angles	90.06%
18 angles	90.75%
32 angles	89.41%
Ours (PointNet)	
6 angles	87.5%
18 angles	88.5%
32 angles	88.6%
Ours + Context Prediction	91.84%

Table 2: Test accuracy on ModelNet-40 shape classification using the pretext task of rotation classification on ShapeNet.

vidual models, demonstrating that our approach learns complementary features to other self-supervised methods. Intuitively, our rotation estimation model learns global orientation information while [67] learns part-level information required for the task of context prediction. Both of these information sources are useful for the shape classification task, therefore combining them leads to improvements over each individual model.

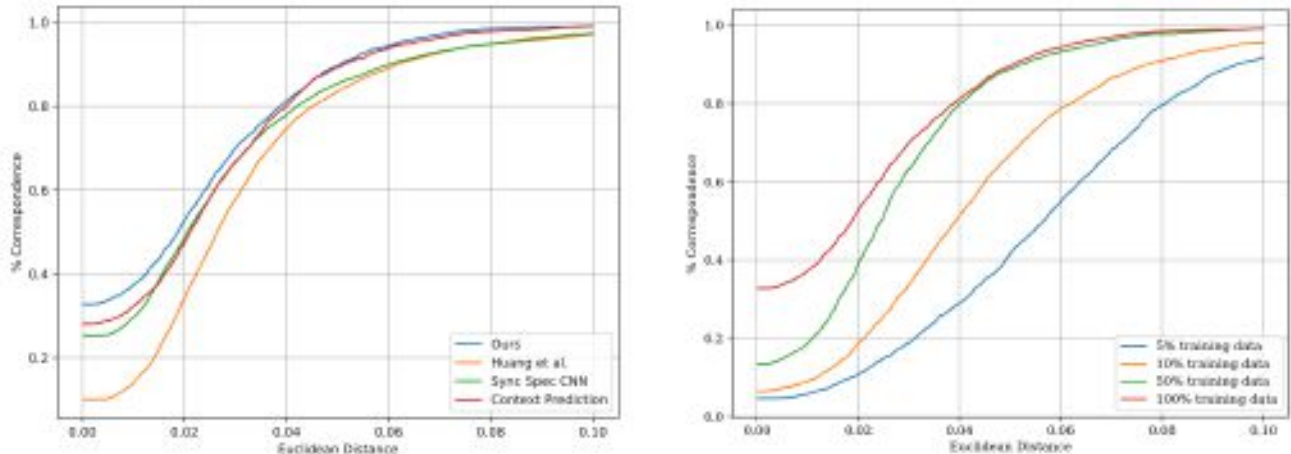


Figure 4: PCK curves for the keypoint prediction task. Left: our approach outperforms competing methods for all error thresholds. Right: PCK curves for various sizes of the training set.

Continuous 6D [95]	86.16%
Axis-angle (PointNet)	85.51%
Axis-angle (DGCNN)	83.12%

Table 3: Test accuracy on ModelNet-40 shape classification using the pretext task of rotation regression on ShapeNet. Different rotation representations are considered.

We can also estimate rotations via regression instead of classification. As described in Section 3, in this approach we aim to regress the exact rotation angle rather than a discretized approximation of it via K pre-selected angles. We consider the axis-angle representation with PointNet and DGCNN backbone architectures, as well as the 6D representation in [95] with a simplified PointNet structure. We train the rotation regression model on ShapeNet and use it to extract features from ModelNet samples. A linear SVM classifier is trained on these features. Table 3 shows the accuracy on ModelNet-40’s test set. We observe that features learned from the regression model achieve inferior performance compared with the classification model. Regressing the exact rotation angle is a challenging task. It resembles our classification setting with $K \rightarrow \infty$ which does not yield the best performance. The learned features become too task-specific and not as transferable to a shape categorization task.

We also study dependence of our model on the size of the training set. We vary ModelNet-40’s train set size for the linear SVM. A certain percentage of data is randomly sampled and used to train the classifier. We then evaluate

the model’s accuracy on the full test set of ModelNet-40. Figure 3 depicts the linear SVM’s classification accuracy as size of the training set varies from 1% to 100%. The results indicate that we still have a relatively strong performance even when a small number of annotated samples are available. The features learned from the rotation prediction model contain information useful for the classification task. Therefore, a simple classifier trained with moderate supervision suffices to achieve high accuracy.

4.2. 3D Keypoint Prediction

We also consider the task of 3D keypoint prediction. We use the training data provided by [90] which includes 10 keypoints for each object in ShapeNet’s chair category. We pre-train a model for rotation prediction and fine-tune it for keypoint regression. More specifically, we first train a rotation prediction model (on all categories) and then modify its final layers to regress 10 keypoints for each input shape. The regression network is initialized with weights of the rotation prediction model and then fine-tuned with chamfer loss between predicted and ground truth keypoints. At inference time, we map each regressed point to its nearest neighbor point in the input point cloud.

PCK¹ curves for our approach are shown in Figure 4, and are compared with [90] and [25]. We observe that our method outperforms competing methods for all error thresholds. We also plot PCK curves for various sizes of the training set and see that our model achieves a good performance even when trained on a subset of training data. Note that compared with the shape classification task, keypoint prediction relies more on local and part-level information.

¹Percentage of Correct Correspondences



Figure 5: Visualization of predicted (red) and ground truth (black) keypoints. Each row represents a separate object from different viewpoints. Zoom in for details.

Our results indicate that our self-supervised approach can transfer useful local and global information. We visualize the keypoints in Figure 5. We observe that predicted and ground truth keypoints are close to each other, and in many cases they overlap.

5. Conclusion and Future Work

We present a self-supervised approach for learning representations of 3D point clouds. A supervised signal for the pretext task of rotation prediction is created to learn representations that are useful for solving target tasks. We devise models for learning proxy and downstream tasks, and demonstrate results on shape classification and 3D keypoint prediction. The results indicate that our approach achieves superior performance to competing methods. Each self-supervised learning method learns features tailored for its own proxy task. We demonstrate that features from different models can complement one another and outperform each individual model. Finally, we show that only a small subset of the training data suffices to achieve high accuracy on shape classification, eliminating the need for annotating massive 3D datasets.

While our work focuses on point clouds, self-supervised learning can also be applied to other 3D representations such as meshes, voxel grids and implicit functions. Each representation has its own merits and disadvantages. Studying self-supervision for different representations is relatively under-explored and devising effective proxy tasks for various 3D representations is a promising venue for future work. We can also consider combining self-supervised techniques for different representations. Finally, evaluating our model on other downstream tasks is another direction for future research.

References

- [1] P. Achlioptas, O. Diamanti, I. Mitliagkas, and L. Guibas. Learning representations and generative models for 3d point clouds. In *International conference on machine learning*, pages 40–49, 2018.
- [2] P. Agrawal, J. Carreira, and J. Malik. Learning to see by moving. In *Proceedings of the IEEE international conference on computer vision*, pages 37–45, 2015.
- [3] R. Arandjelovic and A. Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 609–617, 2017.
- [4] Y. M. Asano, C. Rupprecht, and A. Vedaldi. Acritical analysis of self-supervision, or what we can learn from a single image. *arXiv preprint arXiv:1904.13132*, 2019.
- [5] P. Bojanowski and A. Joulin. Unsupervised learning by predicting noise. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 517–526. JMLR. org, 2017.
- [6] A. Brock, T. Lim, J. M. Ritchie, and N. Weston. Generative and discriminative voxel modeling with convolutional neural networks. *arXiv preprint arXiv:1608.04236*, 2016.
- [7] W. Cao, Z. Yan, Z. He, and Z. He. A comprehensive survey on geometric deep learning. *IEEE Access*, 8:35929–35949, 2020.
- [8] M. Caron, P. Bojanowski, A. Joulin, and M. Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 132–149, 2018.
- [9] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [10] C. B. Choy, D. Xu, J. Gwak, K. Chen, and S. Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *European conference on computer vision*, pages 628–644. Springer, 2016.
- [11] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5828–5839, 2017.
- [12] V. R. de Sa. Learning classification with unlabeled data. In *Advances in neural information processing systems*, pages 112–119, 1994.
- [13] C. Doersch, A. Gupta, and A. A. Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1422–1430, 2015.
- [14] A. Dosovitskiy, P. Fischer, J. T. Springenberg, M. Riedmiller, and T. Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(9):1734–1747, 2015.
- [15] H. Fan, H. Su, and L. J. Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017.
- [16] H. Fu, D. Cohen-Or, G. Dror, and A. Sheffer. Upright orientation of man-made objects. In *SIGGRAPH*, 2008.
- [17] M. Gadelha, R. Wang, and S. Maji. Multiresolution tree networks for 3d point cloud processing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 103–118, 2018.
- [18] C. Gan, B. Gong, K. Liu, H. Su, and L. J. Guibas. Geometry guided convolutional neural networks for self-supervised video representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5589–5597, 2018.
- [19] R. Gao, D. Jayaraman, and K. Grauman. Object-centric representation learning from unlabeled videos. In *Asian Conference on Computer Vision*, pages 248–263. Springer, 2016.
- [20] S. Gidaris, P. Singh, and N. Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR 2018*, 2018.

- [21] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [22] K. Gregor, I. Danihelka, A. Mnih, C. Blundell, and D. Wierstra. Deep autoregressive networks. In *International Conference on Machine Learning*, pages 1242–1250, 2014.
- [23] Z. Han, M. Shang, Y.-S. Liu, and M. Zwicker. View inter-prediction gan: Unsupervised representation learning for 3d shapes by learning global shape memories to support local view predictions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8376–8384, 2019.
- [24] K. Hassani and M. Haley. Unsupervised multi-task feature learning on point clouds. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8160–8171, 2019.
- [25] Q.-X. Huang, H. Su, and L. Guibas. Fine-grained semi-supervised labeling of large shape collections. *ACM Transactions on Graphics (TOG)*, 32(6):1–10, 2013.
- [26] X. Huang, Y. Li, O. Poursaeed, J. Hopcroft, and S. Belongie. Stacked generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5077–5086, 2017.
- [27] P. Isola, D. Zoran, D. Krishnan, and E. H. Adelson. Learning visual groups from co-occurrences in space and time. *arXiv preprint arXiv:1511.06811*, 2015.
- [28] D. Jayaraman and K. Grauman. Slow and steady feature analysis: higher order temporal coherence in video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3852–3861, 2016.
- [29] S. Jenni and P. Favaro. Self-supervised feature learning by learning to spot artifacts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2733–2742, 2018.
- [30] A. Kanazawa, S. Tulsiani, A. A. Efros, and J. Malik. Learning category-specific mesh reconstruction from image collections. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 371–386, 2018.
- [31] M. Kiapour, S. Zheng, R. Piramuthu, and O. Poursaeed. Generating a digital image using a generative adversarial network, Feb. 4 2020. US Patent 10,552,714.
- [32] D. P. Kingma, S. Mohamed, D. J. Rezende, and M. Welling. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*, pages 3581–3589, 2014.
- [33] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [34] R. Kulkov and V. Lempitsky. Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 863–872, 2017.
- [35] A. Kolesnikov, X. Zhai, and L. Beyer. Revisiting self-supervised visual representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1920–1929, 2019.
- [36] G. Larsson, M. Maire, and G. Shakhnarovich. Learning representations for automatic colorization. In *European Conference on Computer Vision*, pages 577–593. Springer, 2016.
- [37] C.-L. Li, M. Zaheer, Y. Zhang, B. Póczos, and R. Salakhutdinov. Point cloud gan. *arXiv preprint arXiv:1810.05795*, 2018.
- [38] J. Li, K. Xu, S. Chaudhuri, E. Yumer, H. Zhang, and L. Guibas. Grass: Generative recursive autoencoders for shape structures. *ACM Transactions on Graphics (TOG)*, 36(4):1–14, 2017.
- [39] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, and B. Chen. Pointcnn: Convolution on x-transformed points. In *Advances in neural information processing systems*, pages 820–830, 2018.
- [40] R. Liao, A. Schwing, R. Zemel, and R. Urtasun. Learning deep parsimonious representations. In *Advances in Neural Information Processing Systems*, pages 5076–5084, 2016.
- [41] Y. Liu, B. Fan, S. Xiang, and C. Pan. Relation-shape convolutional neural network for point cloud analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8895–8904, 2019.
- [42] A. Mahendran, J. Thewlis, and A. Vedaldi. Cross pixel optical-flow similarity for self-supervised learning. In *Asian Conference on Computer Vision*, pages 99–116. Springer, 2018.
- [43] D. Maturana and S. Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 922–928. IEEE, 2015.
- [44] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019.
- [45] I. Misra, C. L. Zitnick, and M. Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision*, pages 527–544. Springer, 2016.
- [46] T. Nathan Mundhenk, D. Ho, and B. Y. Chen. Improvements to context based self-supervised learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9339–9348, 2018.
- [47] M. Noroozi and P. Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pages 69–84. Springer, 2016.
- [48] M. Noroozi, H. Pirsiavash, and P. Favaro. Representation learning by learning to count. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5898–5906, 2017.
- [49] M. Noroozi, A. Vinjimoor, P. Favaro, and H. Pirsiavash. Boosting self-supervised learning via knowledge transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9359–9367, 2018.
- [50] A. v. d. Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [51] A. Owens, P. Isola, J. McDermott, A. Torralba, E. H. Adelson, and W. T. Freeman. Visually indicated sounds. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2405–2413, 2016.

- [52] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019.
- [53] D. Pathak, R. Girshick, P. Dollár, T. Darrell, and B. Hariharan. Learning features by watching objects move. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2701–2710, 2017.
- [54] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016.
- [55] O. Poursaeed, M. Fisher, N. Aigerman, and V. G. Kim. Coupling explicit and implicit surface representations for generative 3d modeling. *arXiv preprint arXiv:2007.10294*, 2020.
- [56] O. Poursaeed, V. Kim, E. Shechtman, J. Saito, and S. Belongie. Neural puppet: Generative layered cartoon characters. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 3346–3356, 2020.
- [57] O. Poursaeed, G. Yang, A. Prakash, Q. Fang, H. Jiang, B. Hariharan, and S. Belongie. Deep fundamental matrix estimation without correspondences. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018.
- [58] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [59] C. R. Qi, H. Su, M. Nießner, A. Dai, M. Yan, and L. J. Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5648–5656, 2016.
- [60] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108, 2017.
- [61] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [62] A. Ranjan, T. Bolkart, S. Sanyal, and M. J. Black. Generating 3d faces using convolutional mesh autoencoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 704–720, 2018.
- [63] Z. Ren and Y. Jae Lee. Cross-domain self-supervised multi-task feature learning using synthetic imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 762–771, 2018.
- [64] D. J. Rezende, S. A. Eslami, S. Mohamed, P. Battaglia, M. Jaderberg, and N. Heess. Unsupervised learning of 3d structure from images. In *Advances in neural information processing systems*, pages 4996–5004, 2016.
- [65] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pages 1278–1286, 2014.
- [66] G. Riegler, A. Osman Ulusoy, and A. Geiger. Octnet: Learning deep 3d representations at high resolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3577–3586, 2017.
- [67] J. Sauder and B. Sievers. Context prediction for unsupervised deep learning on point clouds. *arXiv preprint arXiv:1901.08396*, 2019.
- [68] P. Sermanet, C. Lynch, Y. Chebotar, J. Hsu, E. Jang, S. Schaal, S. Levine, and G. Brain. Time-contrastive networks: Self-supervised learning from video. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1134–1141. IEEE, 2018.
- [69] A. Sharma, O. Grau, and M. Fritz. Vconv-dae: Deep volumetric shape learning without object labels. In *European Conference on Computer Vision*, pages 236–250. Springer, 2016.
- [70] Y. Shen, C. Feng, Y. Yang, and D. Tian. Mining point cloud local structures by kernel correlation and graph pooling. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4548–4557, 2018.
- [71] N. Srivastava, E. Mansimov, and R. Salakhudinov. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, pages 843–852, 2015.
- [72] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 945–953, 2015.
- [73] Y. Sun, Y. Wang, Z. Liu, J. Siegel, and S. Sarma. Pointgrow: Autoregressively learned point cloud generation with self-attention. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 61–70, 2020.
- [74] M. Tatarchenko, A. Dosovitskiy, and T. Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2088–2096, 2017.
- [75] A. Thabet, H. Alwassel, and B. Ghanem. Self-supervised learning of local features in 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 938–939, 2020.
- [76] A. Van den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves, et al. Conditional image generation with pixelcnn decoders. In *Advances in neural information processing systems*, pages 4790–4798, 2016.
- [77] A. Van Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning*, pages 1747–1756, 2016.
- [78] N. Wang, Y. Zhang, Z. Li, Y. Fu, W. Liu, and Y.-G. Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 52–67, 2018.
- [79] X. Wang and A. Gupta. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2794–2802, 2015.
- [80] X. Wang, K. He, and A. Gupta. Transitive invariance for self-supervised visual representation learning. In *Proceedings of*

- the *IEEE international conference on computer vision*, pages 1329–1338, 2017.
- [81] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 38(5):1–12, 2019.
- [82] D. Wei, J. J. Lim, A. Zisserman, and W. T. Freeman. Learning and using the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8052–8060, 2018.
- [83] L. Wei, Q. Huang, D. Ceylan, E. Vouga, and H. Li. Dense human body correspondences using convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1544–1553, 2016.
- [84] J. Wu, C. Zhang, T. Xue, B. Freeman, and J. Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Advances in neural information processing systems*, pages 82–90, 2016.
- [85] W. Wu, Z. Qi, and L. Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9621–9630, 2019.
- [86] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015.
- [87] Q. Xu, W. Wang, D. Ceylan, R. Mech, and U. Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. In *Advances in Neural Information Processing Systems*, pages 492–502, 2019.
- [88] J. Yang, D. Parikh, and D. Batra. Joint unsupervised learning of deep representations and image clusters. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5147–5156, 2016.
- [89] Y. Yang, C. Feng, Y. Shen, and D. Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 206–215, 2018.
- [90] L. Yi, H. Su, X. Guo, and L. J. Guibas. Syncspecnn: Synchronized spectral cnn for 3d shape segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2282–2290, 2017.
- [91] L. Zhang and Z. Zhu. Unsupervised feature learning for point cloud understanding by contrasting and clustering using graph convolutional neural networks. In *2019 International Conference on 3D Vision (3DV)*, pages 395–404. IEEE, 2019.
- [92] R. Zhang, P. Isola, and A. A. Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016.
- [93] R. Zhang, P. Isola, and A. A. Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1058–1067, 2017.
- [94] Y. Zhao, T. Birdal, H. Deng, and F. Tombari. 3d point capsule networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1009–1018, 2019.
- [95] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2019.
- [96] C. Zhu, K. Xu, S. Chaudhuri, R. Yi, and H. Zhang. Scores: Shape composition with recursive substructure priors. *ACM Transactions on Graphics (TOG)*, 37(6):1–14, 2018.
- [97] C. Zou, E. Yumer, J. Yang, D. Ceylan, and D. Hoiem. 3d-prnn: Generating shape primitives with recurrent neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 900–909, 2017.