

DEEP LEARNING WITH SETS AND POINT CLOUDS

Siamak Ravanbakhsh, Jeff Schneider & Barnabás Póczos

School of Computer Science

Carnegie Mellon University

Pittsburgh, PA 15213, USA

{mravanba, jeff.schneider, bapoczos}@cs.cmu.edu

ABSTRACT

We introduce a simple permutation equivariant layer for deep learning with set structure. This type of layer, obtained by parameter-sharing, has a simple implementation and linear-time complexity in the size of each set. We use deep permutation-invariant networks to perform point-cloud classification and MNIST-digit summation, where in both cases the output is invariant to permutations of the input. In a semi-supervised setting, where the goal is make predictions for each instance within a set, we demonstrate the usefulness of this type of layer in set-outlier detection as well as semi-supervised learning with clustering side-information.

1 INTRODUCTION

Recent progress in deep learning (LeCun et al., 2015) has witnessed its application to structured settings, including graphs (Bruna et al., 2013; Duvenaud et al., 2015) groups (Gens & Domingos, 2014; Christopher, 2014; Cohen & Welling, 2016), sequences and hierarchies (Irsoy & Cardie, 2014; Socher et al., 2013). Here, we introduce a simple permutation-equivariant layer for deep learning with set structure, where the primary dataset is a collection of sets, possibly of different sizes. Note that each instance may have a structure of its own, such as graph, image, or another set. In typical machine-learning applications, iid assumption implies that the entire data-set itself has a set structure. Therefore, our special treatment of the set structure is only necessary due to multiplicity of distinct yet homogeneous (data)sets. Here, we show that a simple parameter-sharing scheme enables a general treatment of sets within supervised and semi-supervised settings.

In the following, after introducing the set layer in Section 2, we explore several novel applications. Section 3 studies supervised learning with sets that requires “invariance” to permutation of inputs. Section 3.1 considers the task of summing multiple MNIST digits, and Section 3.2 studies an important application of sets in representing low-dimensional point-clouds. Here, we show that deep networks can successfully classify objects using their point-cloud representation.

Section 4 presents numerical study in semi-supervised setting where the output of the multi-layer network is “equivariant” to the permutation of inputs. We use permutation-equivariant layer to perform outlier detection on CelebA face dataset in Section 4.1 and improve galaxy red-shift estimates using its clustering information in Section 5.

2 PERMUTATION-EQUIVARIANT LAYER

Let $x_n \in \mathcal{X}$ denote use $\mathbf{x} = [x_1, \dots, x_N]$ to denote a vector of x_n instances. Here, x_n , could be a feature-vector, an image or any other structured object. Our goal is to design neural network layers that are “indifferent” to permutations of instances in \mathbf{x} . Achieving this goal amounts to treating \mathbf{x} as a “set” rather than a vector.

The function $\mathbf{f} : \mathcal{X}^N \rightarrow \mathcal{Y}^N$ is **equivariant** to the permutation of its inputs iff

$$\mathbf{f}(\pi\mathbf{x}) = \pi\mathbf{f}(\mathbf{x}) \quad \forall \pi \in \mathcal{S}_N$$

where the symmetric group \mathcal{S}_N is the set of all permutation of indices $1, \dots, N$. Similarly, the function $f : \mathcal{R}^N \rightarrow \mathcal{R}$ is **invariant** to permutation of its inputs *-i.e.*, a.k.a. a symmetric function (David




Figure 1: Classification accuracy of different schemes in predicting the sum of a (left) $N=3$ and (right) $N=6$ MNIST digits without access to individual image labels. The training set was fixed to 10,000 sets.

et al., 1966)– iff $f(\pi\mathbf{x}) = f(\mathbf{x}) \quad \forall \pi \in \mathcal{S}_N$. Here, the **action** of π on a vector $\mathbf{x} \in \mathbb{R}^N$ can be represented by a permutation matrix. With some abuse of notation, we use $\pi \in \{0, 1\}^{N \times N}$ to also denote this matrix.

Given two permutation equivariance function $\mathbf{f} : \mathcal{X}^N \rightarrow \mathcal{Y}^N$ and $\mathbf{g} : \mathcal{Y}^N \rightarrow \mathcal{Z}^N$, their **composition** is also permutation-equivariance; this is because $\mathbf{g}(\mathbf{f}(\pi\mathbf{x})) = \mathbf{g}(\pi\mathbf{f}(\mathbf{x})) = \pi\mathbf{g}(\mathbf{f}(\mathbf{x}))$.

Consider the standard neural network layer

$$\mathbf{f}_\Theta(\mathbf{x}) \doteq \sigma(\Theta\mathbf{x}) \quad \Theta \in \mathbb{R}^{N \times N} \quad (1)$$

where Θ is the weight vector and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a nonlinearity such as sigmoid function. $\sigma : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is the point-wise application of σ to its input vector. The following theorem states the necessary and sufficient conditions for permutation-equivariance in this type of function.

Theorem 2.1. *The function $\mathbf{f}_\Theta : \mathbb{R}^N \rightarrow \mathbb{R}^N$ as defined in Eq. (1) is permutation equivariant iff all the off-diagonal elements of Θ are tied together and all the diagonal elements are equal as well,*

$$\Theta = \lambda\mathbf{I} + \gamma(\mathbf{1}\mathbf{1}^\top) \quad \lambda, \gamma \in \mathbb{R} \quad \mathbf{1} = [1, \dots, 1]^\top \in \mathbb{R}^N \quad (2)$$

where \mathbf{I} is the identity matrix.¹

This function is simply a non-linearity applied to a weighted combination of I) its input $\mathbf{I}\mathbf{x}$ and; II) the sum of input values $(\mathbf{1}\mathbf{1}^\top)\mathbf{x}$. Since summation does not depend on the permutation, the layer is permutation-equivariant. Therefore we can manipulate the operations and parameters in this layer, for example to get another **variation**

$$\mathbf{f}(\mathbf{x}) \doteq \sigma(\lambda\mathbf{I}\mathbf{x} - \gamma(\max_n \mathbf{x})\mathbf{1}) \quad (3)$$

where the max operation over elements of the set is (similar to summation) commutative and using $-\gamma$ instead of $+\gamma$ amounts to a reparametrization. In practice using this variation performs better in some applications. This may be due to the fact that for $\lambda = \gamma$, the input to the non-linearity is max-normalized.

For **multiple input-output channels**, we may speed up the operation of the layer using matrix multiplication. Suppose we have K input channels –corresponding to K features for each instance in the set– with a set of size N , and K' output channels. Here, $\mathbf{x} \in \mathbb{R}^{N \times K}$ and $\mathbf{f}_\Theta : \mathbb{R}^{N \times K} \rightarrow \mathbb{R}^{N \times K'}$. The permutation-equivariant layer parameters are $\Lambda, \Gamma \in \mathbb{R}^{K, K'}$ (replacing λ and γ in Eq. (2)). The output of this layer becomes

$$\mathbf{y} = \sigma(\mathbf{x}\Lambda + \mathbf{1}\mathbf{1}^\top\mathbf{x}\Gamma)$$

Similarly, the multiple input-output channel variation of the Eq. (3) is

$$\mathbf{y} = \sigma(\mathbf{x}\Lambda - \mathbf{1}\mathbf{x}_{\max}\Gamma) \quad (4)$$

where $\mathbf{x}_{\max} = (\max_n \mathbf{x}) \in \mathbb{R}^{1 \times K}$ is a row-vector of maximum value of $\mathbf{x} \in \mathbb{R}^{N \times K}$ over the “set” dimension. In practice, we may further reduce the number of parameters in favor of better generalization by factoring Γ and Λ and keeping a single $\lambda \in \mathbb{R}^{K, K'}$

$$\mathbf{y} = \sigma(\beta + (\mathbf{x} - \mathbf{1}(\max_n \mathbf{x}))\Gamma) \quad (5)$$

¹See the Appendix A for the proof.




Figure 2: Examples for 8 out of 40 object classes (column) in the ModelNet40. Each point-cloud is produced by sampling 1000 particles from the mesh representation of the original ModelNet40 instances. Two point-clouds in the same column are from the same class. The projection of particles into xy , zy and xz planes are added for better visualization.

where $\beta \in \mathbb{R}^{K'}$ is a bias parameter. This final variation of permutation-equivariant layer is simply a fully connected layer where input features are max-normalized within each set.

With multiple input-output channels, the complexity of this layer for each set is $\mathcal{O}(N K K')$. Subtracting the mean or max over the set also reduces the internal covariate shift (Ioffe & Szegedy, 2015) and we observe that for deep networks (even using tanh activation), batch-normalization is not required.

When applying dropout (Srivastava et al., 2014) to regularize permutation-equivariant layers with multiple output channels, it is often beneficial to *simultaneously dropout* the channels for all instances within a set. In particular, when set-members share similar features, independent dropout effectively does not regularize the model as the network learns to replace the missing features from other set-members.

In the remainder of this paper, we demonstrate how this simple treatment of sets can solve novel and non-trivial problems that occasionally have no alternative working solutions within deep learning.

3 SUPERVISED LEARNING

The permutation-equivariant layers that we introduced so far are useful for semi-supervised learning (or *transductive*) setting, where we intend to predict a value per each instance in every set. In supervised (or *inductive*) setting the task is to make a prediction for each set (rather than instances within them), and we require permutation “invariance” of $f : \mathcal{X}^N \rightarrow \mathcal{Y}$. A pooling operation over the set-dimension can turn any permutation equivariant function $\mathbf{f} : \mathcal{X}^N \rightarrow \mathcal{Y}^N$, permutation invariant $f(\mathbf{x}) \doteq \bigoplus_n \mathbf{f}(\mathbf{x})$. Here \bigoplus is any commutative operation such as summation or maximization.

In a related work Chen et al. (2014) construct deep permutation invariant features by pairwise coupling of features at the previous layer, where $\mathbf{f}_{i,j}([x_i, x_j]) \doteq [|x_i - x_j|, x_i + x_j]$ is invariant to transposition of i and j .²

As we see shortly in Section 3.1, in the supervised setting, even simple application of set-pooling, without max-normalization of Eq. (5) performs very well in practice. However, in semi-supervised setting since there is no pooling operation, the permutation invariant layer requires max-normalization in order to obtain the required information about the *context* of each instance.

3.1 PREDICTING THE SUM OF MNIST DIGITS

MNIST dataset (LeCun et al., 1998) contains 70,000 instances of 28×28 grey-scale stamps of digits in $\{0, \dots, 9\}$. We randomly sample a subset of N images from this dataset to build 10,000 “sets” of training and 10,000 sets of validation images, where the set-label is the sum of digits in that set (*i.e.*, individual labels per image is unavailable).

²Due to change in the number of features/channels in each layer this approach cannot produce permutation-“equivariant” layers. Also, this method requires a graph to guide the multi-resolution partitioning of the nodes, which is then used to define pairing of features in each layer.

Table 1: Classification accuracy and the (size of) representation used by different methods on the ModelNet40 dataset.

model	instance size	representation	accuracy
set-layer + transformation (ours)	5000 × 3	point-cloud	90 ± .3%
set-layer (ours)	1000 × 3	point-cloud	87 ± 1%
set-pooling only (ours)	1000 × 3	point-cloud	83 ± 1%
set-layer (ours)	100 × 3	point-cloud	82 ± 2%
KNN graph-convolution (ours)	1000 × (3 + 8)	directed 8-regular graph	58 ± 2%
3DShapeNets (Wu et al., 2015)	30 ³	voxels (using convolutional deep belief net)	77%
DeepPano (Shi et al., 2015)	64 × 160	panoramic image (2D CNN + angle-pooling)	77.64%
VoxNet (Maturana & Scherer, 2015)	32 ³	voxels (voxels from point-cloud + 3D CNN)	83.10%
MVCNN (Su et al., 2015)	164 × 164 × 12	multi-view images (2D CNN + view-pooling)	90.1%
VRN Ensemble (Brock et al., 2016)	32 ³	voxels (3D CNN, variational autoencoder)	95.54%
3D GAN (Wu et al., 2016)	64 ³	voxels (3D CNN, generative adversarial training)	83.3%

In our first experiment, each set contains $N = 3$ images, and the set label is a number between $0 \leq y \leq 3 * 9 = 27$. We then considered four different models for predicting the sum:

- I) Concatenating instances along the horizontal axis.
- II) Stacking images in each set as different input channels.
- III) Using set-pooling without max-normalization.
- IV) Using the permutation equivariant layer of Eq. (5) with set-pooling.

All models are defined to have similar number of layers and parameters; see Appendix B.1 for details. The output of all models is a $(9N + 1)$ -way softmax, predicting the sum of N digits.

Figure 1(left) show the prediction accuracy over the validation-set for different models for $N = 3$. We see that using the set-layer performs the best. However, interestingly, using set-pooling alone produces similarly good results. We also observe that concatenating the digits eventually performs well, despite its lack of invariance. This is because due to the sufficiently large size of the dataset most permutations of length three appear in the training set.

However, as we increase the size of each set to $N = 6$, permutation invariance becomes crucial; see Figure 1(right). We see that using the default dropout rate of 20%, the model simply memorizes the input instances (indicated by discrepancy of training/validation error) and by increasing this dropout rate, the model simply predicts values close to the mean value. However, the permutation-invariant layer learns to predict the sum of six digits with $> 80\%$ accuracy, without having access to individual image labels. Performance using set-pooling alone is similar.

3.2 POINT CLOUD CLASSIFICATION

A low-dimensional point-cloud is a set of low-dimensional vectors. This type of data is frequently encountered in various applications from robotics and vision to cosmology. In these applications, point-cloud data is often converted to voxel or mesh representation at a preprocessing step (e.g., Maturana & Scherer, 2015; Ravanbakhsh et al., 2016; Lin et al., 2004). Since the output of many range sensors such as LiDAR – which are extensively used in applications such as autonomous vehicles – is in the form of point-cloud, direct application of deep learning methods to point-cloud is highly desirable. Moreover, when working with point-clouds rather than voxelized 3D objects, it is easy to apply transformations such as rotation and translation as differentiable layers at low cost.

Here, we show that treating the point-cloud data as a set, we can use the set-equivariant layer of Eq. (5) to classify point-cloud representation of a subset of ShapeNet objects (Chang et al., 2015), called ModelNet40 (Wu et al., 2015). This subset consists of 3D representation of 9,843 training and 2,468 test instances belonging to 40 classes of objects; see Fig. 2. We produce point-clouds with 100, 1000 and 5000 particles each (x, y, z -coordinates) from the mesh representation of objects using the point-cloud-library’s sampling routine (Rusu & Cousins, 2011). Each set is normalized by the initial layer of the deep network to have zero mean (along individual axes) and unit (global)

variance. Additionally we experiment with the K-nearest neighbor graph of each point-cloud and report the results using graph-convolution; see Appendix B.3 for model details.

Table 1 compares our method against the competition.³ Note that we achieve our best accuracy using 5000×3 dimensional representation of each object, which is much smaller than most other methods. All other techniques use either voxelization or multiple view of the 3D object for classification. Interestingly, variations of view/angle-pooling (e.g., Su et al., 2015; Shi et al., 2015) can be interpreted as set-pooling where the class-label is invariant to permutation of different views. The results also shows that using fully-connected layers with set-pooling alone (without max-normalization over the set) works relatively well.

We see that reducing the number of particles to only 100, still produces comparatively good results. Using graph-convolution is computationally more challenging and produces inferior results in this setting. The results using 5000 particles is also invariant to small changes in scale and rotation around the z -axis; see Appendix B.3 for details.

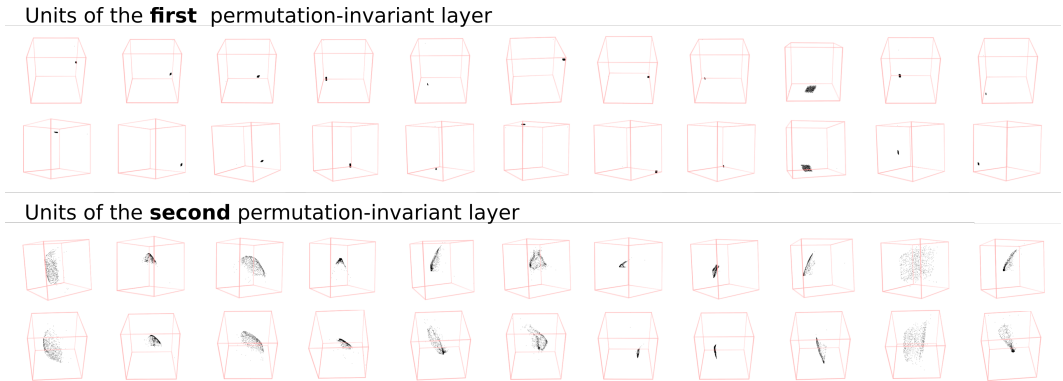


Figure 3: Each box is the particle-cloud maximizing the activation of a unit at the first (top) and second (bottom) permutation-equivariant layers of our model. Two images of the same column are two different views of the same point-cloud.

Features. To visualize the features learned by the set layers, we used Adamax (Kingma & Ba, 2014) to locate 1000 particle coordinates maximizing the activation of each unit.⁴ Activating the tanh units beyond the second layer proved to be difficult. Figure 3 shows the particle-cloud-features learned at the first and second layers of our deep network. We observed that the first layer learns simple localized (often cubic) point-clouds at different (x, y, z) locations, while the second layer learns more complex surfaces with different scales and orientations.

4 SEMI-SUPERVISED LEARNING

In semi-supervised or transductive learning, some/all instances within each training set are labelled. Our goal is to make predictions for individual instances within a test set. Therefore, the permutation equivariant layer leverages the interaction between the set-members to label individual member. Note that in this case, we do not perform any pooling operation over the set dimension of the data.

4.1 SET ANOMALY DETECTION

The objective here is for the deep model to find the anomalous face in each set, simply by observing examples and without any access to the attribute values. CelebA dataset (Liu et al., 2015) contains

³The error-bar on our results is due to variations depending on the choice of particles during test time and it is estimated over three trials.

⁴We started from uniformly distributed set of particles and used a learning rate of .01 for Adamax, with first and second order moment of .1 and .9 respectively. We optimized the input in 10^5 iterations. The results of Fig. 3 are limited to instances where tanh units were successfully activated. Since the input at the first layer of our deep network is normalized to have a zero mean and unit standard deviation, we do not need to constrain the input while maximizing unit’s activation.



Figure 4: Each row shows a set, constructed from CelebA dataset, such that all set members except for an outlier, share at least two attributes (on the right). The outlier is identified with a red frame. The model is trained by observing examples of sets and their anomalous members, *without access to the attributes*. The probability assigned to each member by the outlier detection network is visualized using a red bar at the bottom of each image. The probabilities in each row sum to one. See Appendix B.2 for more examples.

202,599 face images, each annotated with 40 boolean attributes. We use 64×64 stamps and using these attributes we build 18,000 sets, each containing $N = 16$ images (on the training set) as follows: after randomly selecting two attributes, we draw 15 images where those attributes are present and a single image where both attributes are absent. Using a similar procedure we build sets on the test images. No individual person’s face appears in both train and test sets.

Our deep neural network consists of 9 2D-convolution and max-pooling layers followed by 3 permutation-equivariant layers and finally a softmax layer that assigns a probability value to each set member (Note that one could identify arbitrary number of outliers using a sigmoid activation at the output.) Our trained model successfully finds the anomalous face in **75% of test sets**. Visually inspecting these instances suggests that the task is non-trivial even for humans; see Fig. 4. For details of the model, training and more identification examples see Appendix B.2.

As a *baseline*, we repeat the same experiment by using a set-pooling layer after convolution layers, and replacing the permutation-equivariant layers with fully connected layers, with the same number of hidden units/output-channels, where the final layer is a 16-way softmax. The resulting network shares the convolution filters for all instances within all sets, however the input to the softmax is not equivariant to the permutation of input images. Permutation equivariance seems to be crucial here as the baseline model achieves a training and **test accuracy of $\sim 6.3\%$** ; the same as random selection.

5 IMPROVED RED-SHIFT ESTIMATION USING CLUSTERING INFORMATION

An important regression problem in cosmology is to estimate the red-shift of galaxies, corresponding to their age as well as their distance from us (Binney & Merrifield, 1998). Two common types of observation for distant galaxies include a) photometric and b) spectroscopic observations, where the latter can produce more accurate red-shift estimates.

One way to estimate the red-shift from photometric observations is using a regression model (Connolly et al., 1995). We use a multi-layer Perceptron for this purpose and use the more accurate spectroscopic red-shift estimates as the ground-truth. As another baseline, we have a photometric redshift estimate that is provided by the catalogue and uses various observations (including clustering information) to estimate individual galaxy-red-shift. Our objective is to use clustering information of the galaxies to improve our red-shift prediction using the multi-layer Perceptron.

Note that the prediction for each galaxy does not change by permuting the members of the galaxy cluster. Therefore, we can treat each galaxy cluster as a “set” and use permutation-equivariant layer to estimate the individual galaxy red-shifts.

For each galaxy, we have 17 photometric features⁵ from the redMaPPer galaxy cluster catalog (Roza & Rykoff, 2014), which contains photometric readings for 26,111 red galaxy clusters. In this task in contrast to the previous ones, sets have different cardinalities; each galaxy-cluster in this catalog has between $\sim 20 - 300$ galaxies – *i.e.*, $\mathbf{x} \in \mathbb{R}^{N(c) \times 17}$, where $N(c)$ is the cluster-size. See Fig. 5(a) for distribution of cluster sizes. The catalog also provides accurate spectroscopic red-shift estimates for a *subset* of these galaxies as well as photometric estimates that uses clustering information. Fig. 5(b) reports the distribution of available spectroscopic red-shift estimates per cluster.

We randomly split the data into 90% training and 10% test clusters, and use the following simple architecture for semi-supervised learning. We use four permutation-equivariant layers with 128, 128, 128 and 1 output channels respectively, where the output of the last layer is used as red-shift estimate. The squared loss of the prediction for available spectroscopic red-shifts is minimized.⁶ Fig. 5(c) shows the agreement of our estimates with spectroscopic readings on the galaxies in the test-set with spectroscopic readings. The figure also compares the photometric estimates provided by the catalogue (see Roza & Rykoff, 2014), to the ground-truth. As it is customary in cosmology literature, we report the average **scatter** $\frac{|z_{\text{spec}} - z|}{1 + z_{\text{spec}}}$, where z_{spec} is the accurate spectroscopic measurement and z is a photometric estimate. The average scatter using **our model** is .023 compared to the scatter of .025 in the **original photometric estimates** for the redMaPPer catalog. Both of these values are averaged over all the galaxies with spectroscopic measurements in the test-set.

We repeat this experiment, replacing the permutation-equivariant layers with fully connected layers (with the same number of parameters) and only use the individual galaxies with available spectroscopic estimate for training. The resulting average scatter for **multi-layer Perceptron** is .026, demonstrating that using clustering information indeed improves photometric red-shift estimates.




Figure 5: application of permutation-equivariant layer to semi-supervised red-shift prediction using clustering information: **a)** distribution of cluster (set) size; **b)** distribution of reliable red-shift estimates per cluster; **c)** prediction of red-shift on test-set (versus ground-truth) using clustering information as well as RedMaPPer photometric estimates (also using clustering information).

CONCLUSION

We introduced a simple parameter-sharing scheme to effectively achieve permutation-equivariance in deep networks and demonstrated its effectiveness in several novel supervised and semi-supervised tasks. Our treatment of set structure also generalizes various settings in multi-instance learning (Ray et al., 2011; Zhou et al., 2009). In addition to our experimental settings, the permutation-invariant layer can be used for distribution regression and classification which have become popular recently (Szabo et al., 2016). In our experiments with point-cloud data we observed the model to be robust to the variations in the number of particles in each cloud, suggesting the usefulness of our method in the general setting of distribution regression – where the number of samples should not qualitatively affect our representation of a distribution. We leave further investigation of this direction to future work.

⁵We have a single measurement for each u,g,r, i and z band as well as measurement error bars, location of the galaxy in the sky, as well as the probability of each galaxy being the cluster center. We do not include the information regarding the richness estimates of the clusters from the catalog, for any of the methods, so that baseline multi-layer Perceptron is blind to the clusters.

⁶We use mini-batches of size 128, Adam (Kingma & Ba, 2014), with learning rate of .001, $\beta_1 = .9$ and $\beta_2 = .999$. All layers except for the last layer use Tanh units and simultaneous dropout with 50% dropout rate.

ACKNOWLEDGEMENT

We would like to thank Francois Lanusse for the pointing us to the redMaPPer dataset and the anonymous reviewers as well as Andrew Wagner for valuable feedback.

REFERENCES

- Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- James Binney and Michael Merrifield. *Galactic astronomy*. Princeton University Press, 1998.
- Andrew Brock, Theodore Lim, JM Ritchie, and Nick Weston. Generative and discriminative voxel modeling with convolutional neural networks. *arXiv preprint arXiv:1608.04236*, 2016.
- Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- Xu Chen, Xiuyuan Cheng, and Stéphane Mallat. Unsupervised deep haar scattering on graphs. In *Advances in Neural Information Processing Systems*, pp. 1709–1717, 2014.
- Olah Christopher. Groups and group convolutions. <http://colah.github.io/posts/2014-12-Groups-Convolution/>, 2014.
- Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
- Taco S Cohen and Max Welling. Group equivariant convolutional networks. *arXiv preprint arXiv:1602.07576*, 2016.
- AJ Connolly, I Csabai, AS Szalay, DC Koo, RG Kron, and JA Munn. Slicing through multicolor space: Galaxy redshifts from broadband photometry. *arXiv preprint astro-ph/9508100*, 1995.
- F.N. David, M.S. Kendall, and D.E. Barton. *Symmetric Functions and Allied Tables*. University Press, 1966.
- David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems*, pp. 2224–2232, 2015.
- Robert Gens and Pedro M Domingos. Deep symmetry networks. In *Advances in neural information processing systems*, pp. 2537–2545, 2014.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Ozan Irsoy and Claire Cardie. Deep recursive neural networks for compositionality in language. In *Advances in Neural Information Processing Systems*, pp. 2096–2104, 2014.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- Hong-Wei Lin, Chiew-Lan Tai, and Guo-Jin Wang. A mesh reconstruction algorithm driven by an intrinsic property of a point cloud. *Computer-Aided Design*, 36(1):1–9, 2004.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.

- Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pp. 922–928. IEEE, 2015.
- Siamak Ravanbakhsh, Junier Oliva, Sebastien Fromenteau, Layne C Price, Shirley Ho, Jeff Schneider, and Barnabás Póczos. Estimating cosmological parameters from the dark matter distribution. In *Proceedings of The 33rd International Conference on Machine Learning*, 2016.
- Soumya Ray, Stephen Scott, and Hendrik Blockeel. Multi-instance learning. In *Encyclopedia of Machine Learning*, pp. 701–710. Springer, 2011.
- Eduardo Rozo and Eli S Rykoff. redmapper ii: X-ray and sz performance benchmarks for the sdss catalog. *The Astrophysical Journal*, 783(2):80, 2014.
- Radu Bogdan Rusu and Steve Cousins. 3D is here: Point Cloud Library (PCL). In *IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, May 9-13 2011.
- Baoguang Shi, Song Bai, Zhichao Zhou, and Xiang Bai. Deeppano: Deep panoramic representation for 3-d shape recognition. *IEEE Signal Processing Letters*, 22(12):2339–2343, 2015.
- Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the conference on empirical methods in natural language processing (EMNLP)*, volume 1631, pp. 1642. Citeseer, 2013.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1): 1929–1958, 2014.
- Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 945–953, 2015.
- Z. Szabo, B. Sriperumbudur, B. Póczos, and A. Gretton. Learning theory for distribution regression. *Journal of Machine Learning Research*, 2016.
- Jiajun Wu, Chengkai Zhang, Tianfan Xue, William T Freeman, and Joshua B Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *arXiv preprint arXiv:1610.07584*, 2016.
- Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1912–1920, 2015.
- Zhi-Hua Zhou, Yu-Yin Sun, and Yu-Feng Li. Multi-instance learning by treating instances as non-iid samples. In *Proceedings of the 26th annual international conference on machine learning*, pp. 1249–1256. ACM, 2009.

Appendix

A PROOFS

Proof. of the Theorem 2.1

From definition of permutation equivariance $\mathbf{f}_\Theta(\pi\mathbf{x}) = \pi\mathbf{f}_\Theta(\mathbf{x})$ and definition of \mathbf{f} in Eq. (1), the condition becomes $\sigma(\Theta\pi\mathbf{x}) = \pi\sigma(\Theta\mathbf{x})$, which (assuming sigmoid is a bijection) is equivalent to $\Theta\pi = \pi\Theta$. Therefore we need to show that the necessary and sufficient conditions for the matrix $\Theta \in \mathbb{R}^{N \times N}$ to commute with all permutation matrices $\pi \in \mathcal{S}_n$ is given by Eq. (2). We prove this in both directions:

- To see why $\Theta = \lambda\mathbf{I} + \gamma(\mathbf{1}\mathbf{1}^\top)$ commutes with any permutation matrix, first note that commutativity is linear – that is

$$\Theta_1\pi = \pi\Theta_1 \wedge \Theta_2\pi = \pi\Theta_2 \quad \Rightarrow \quad (a\Theta_1 + b\Theta_2)\pi = \pi(a\Theta_1 + b\Theta_2).$$

Since both Identity matrix \mathbf{I} , and constant matrix $\mathbf{1}\mathbf{1}^\top$, commute with any permutation matrix, so does their linear combination $\Theta = \lambda\mathbf{I} + \gamma(\mathbf{1}\mathbf{1}^\top)$.

- We need to show that in a matrix Θ that commutes with “all” permutation matrices
 - *All diagonal elements are identical:* Let $\pi_{k,l}$ for $1 \leq k, l \leq N, k \neq l$, be a transposition (i.e., a permutation that only swaps two elements). The inverse permutation matrix of $\pi_{k,l}$ is the permutation matrix of $\pi_{l,k} = \pi_{k,l}^\top$. We see that commutativity of Θ with the transposition $\pi_{k,l}$ implies that $\Theta_{k,k} = \Theta_{l,l}$:

$$\pi_{k,l}\Theta = \Theta\pi_{k,l} \Rightarrow \pi_{k,l}\Theta\pi_{l,k} = \Theta \Rightarrow (\pi_{k,l}\Theta\pi_{l,k})_{l,l} = \Theta_{l,l} \Rightarrow \Theta_{k,k} = \Theta_{l,l}$$

Therefore, π and Θ commute for any permutation π , they also commute for any transposition $\pi_{k,l}$ and therefore $\Theta_{i,i} = \lambda \forall i$.

- *All off-diagonal elements are identical:* We show that since Θ commutes with any product of transpositions, any choice two off-diagonal elements should be identical. Let (i, j) and (i', j') be the index of two off-diagonal elements (i.e., $i \neq j$ and $i' \neq j'$). Moreover for now assume $i \neq i'$ and $j \neq j'$. Application of the transposition $\pi_{i,i'}$ swaps the rows i, i' in Θ . Similarly, $\Theta\pi_{j,j'}$ switches the j^{th} column with j'^{th} column. From commutativity property of Θ and $\pi \in \mathcal{S}_n$ we have

$$\begin{aligned} \pi_{j',j}\pi_{i,i'}\Theta &= \Theta\pi_{j',j}\pi_{i,i'} \Rightarrow \pi_{j',j}\pi_{i,i'}\Theta(\pi_{j',j}\pi_{i,i'})^{-1} = \Theta & \Rightarrow \\ \pi_{j',j}\pi_{i,i'}\Theta\pi_{i',i}\pi_{j,j'} &= \Theta \Rightarrow (\pi_{j',j}\pi_{i,i'}\Theta\pi_{i',i}\pi_{j,j'})_{i,j} = \Theta_{i,j} & \Rightarrow \Theta_{i',j'} = \Theta_{i,j} \end{aligned}$$

where in the last step we used our assumptions that $i \neq i', j \neq j', i \neq j$ and $i' \neq j'$. In the cases where either $i = i'$ or $j = j'$, we can use the above to show that $\Theta_{i,j} = \Theta_{i'',j''}$ and $\Theta_{i',j'} = \Theta_{i'',j''}$, for some $i'' \neq i, i'$ and $j'' \neq j, j'$, and conclude $\Theta_{i,j} = \Theta_{i',j'}$.

□

B DETAILS OF MODELS

In the following, all our implementations use Tensorflow (Abadi et al., 2016).

B.1 MNIST SUMMATION

All nonlinearities are exponential linear units (ELU Clevert et al., 2015). All models have 4 convolution layers followed by max-pooling. The convolution layers have respectively 16-32-64-128 output channels and 5×5 receptive fields.

Each pooling, fully-connected and set-layer is followed by a 20% dropout. For models III and IV we use simultaneous dropout. In models I and II, the convolution layers are followed by two fully-connected layers with 128 hidden units. In model III, after the first fully connected layer we perform



Figure 6: Each row of the images shows a set, constructed from CelebA dataset images, such that all set members except for an outlier, share at least two attributes. The outlier is identified with a red frame. The model is trained by observing examples of sets and their anomalous members and *without access to the attributes*. The probability assigned to each member by the outlier detection network is visualized using a red bar at the bottom of each image. The probabilities in each row sum to one.

set-pooling followed by another dense layer with 128 hidden units. In the model IV, the convolution layers are followed by a permutation-equivariant layer with 128 output channels, followed by set-pooling and a fully connected layer with 128 hidden units. For optimization, we used a learning rate of .0003 with Adam using the default $\beta_1 = .9$ and $\beta_2 = .999$.

B.2 FACE OUTLIER DETECTION MODEL

Our model has 9 convolution layers with 3×3 receptive fields. The model has convolution layers with 32, 32, 64 feature-maps followed by max-pooling followed by 2D convolution layers with 64, 64, 128 feature-maps followed by another max-pooling layer. The final set of convolution layers have 128, 128, 256 feature-maps, followed by a max-pooling layer with pool-size of 5 that reduces the output dimension to $\text{batch-size} \times N \times 256$, where the set-size $N = 16$. This is then forwarded to three permutation-equivariant layers with 256, 128 and 1 output channels. The output of final layer is fed to the Softmax, to identify the outlier. We use exponential linear units (Clevert et al., 2015), drop out with 20% dropout rate at convolutional layers and 50% dropout rate at the first two set layers. When applied to set layers, the selected feature (channel) is simultaneously dropped in all the set members of that particular set. We use Adam (Kingma & Ba, 2014) for optimization and use batch-normalization only in the convolutional layers. We use mini-batches of 8 sets, for a total of 128 images per batch.

B.3 MODELS FOR POINT-CLOUDS CLASSIFICATION

Set convolution. We use a network comprising of 3 permutation-equivariant layers with 256 channels followed by max-pooling over the set structure. The resulting vector representation of the set is then fed to a fully connected layer with 256 units followed by a 40-way softmax unit. We use Tanh activation at all layers and dropout on the layers after set-max-pooling (*i.e.*, two dropout operations) with 50% dropout rate. Applying dropout to permutation-equivariant layers for point-cloud data deteriorated the performance. We observed that using different types of permutation-equivariant layers (see Section 2) and as few as 64 channels for set layers changes the result by less than 5% in classification accuracy.

For the setting with 5000 particles, we increase the number of units to 512 in all layers and randomly rotate the input around the z -axis. We also randomly scale the point-cloud by $s \sim \mathcal{U}(.8, 1./8)$. For this setting only, we use Adamax (Kingma & Ba, 2014) instead of Adam and reduce learning rate from .001 to .0005.

Graph convolution. For each point-cloud instance with 1000 particles, we build a sparse K-nearest neighbor graph and use the three point coordinates as input features. We normalized all graphs at the preprocessing step. For direct comparison with set layer, we use the exact architecture of 3 graph-convolution layer followed by set-pooling (global graph pooling) and dense layer with 256 units. We use exponential linear activation function instead of Tanh as it performs better for graphs. Due to over-fitting, we use a heavy dropout of 50% after graph-convolution and dense layers. Similar to dropout for sets, all the randomly selected features are simultaneously dropped across the graph nodes. We use a mini-batch size of 64 and Adam for optimization where the learning rate is .001 (the same as that of permutation-equivariant counter-part).

Despite our efficient sparse implementation using Tensorflow, graph-convolution is significantly slower than the set layer. This prevented a thorough search for hyper-parameters and it is quite possible that better hyper-parameter tuning would improve the results that we report here.