

# Learning Non-Parametric Invariances from Data with Permanent Random Connectomes

Dipan K. Pal  
dipanp@andrew.cmu.edu  
Akshay Chawla  
akshaych@andrew.cmu.edu  
Marios Savvides  
marios@andrew.cmu.edu

Dept. Electrical and Computer Engg.  
Carnegie Mellon University  
Pittsburgh, PA, USA

---

## Abstract

Learning non-parametric invariances directly from data remains an important open problem. In this paper, we introduce a new architectural layer for convolutional networks which is capable of learning general invariances from data itself. This layer can learn invariance to non-parametric transformations and interestingly, motivates and incorporates permanent random connectomes, thereby being called Permanent Random Connectome Non-Parametric Transformation Networks (PRC-NPTN). PRC-NPTN networks are initialized with random connections (not just weights) which are a small subset of the connections in a fully connected convolution layer. Importantly, these connections in PRC-NPTNs once initialized remain permanent throughout training and testing. Permanent random connectomes make these architectures loosely more biologically plausible than many other mainstream network architectures which require highly ordered structures. We motivate randomly initialized connections as a simple method to learn invariance from data itself while invoking invariance towards multiple nuisance transformations simultaneously. We find that these randomly initialized permanent connections have positive effects on generalization, outperform much larger ConvNet baselines and the recently proposed Non-Parametric Transformation Network (NPTN) on benchmarks such as augmented MNIST, ETH-80 and CIFAR10, that enforce learning invariances from the data itself.

## 1 Introduction

**The Problem of Invoking Invariances.** Learning invariances to nuisance transformations in data has emerged to be a core problem in machine learning [1, 2, 3, 4, 5]. Moving towards real-world data of different modalities, it is a daunting task to theoretically model all nuisance transformations. Towards this goal, methods which *learn* non-parametric invariances from the data itself without any change in architecture will be critical. However, before delving into methods which learn such invariances, it is important to study methods which incorporate *known* invariances in data. An early method to incorporate the translation prior was the Convolutional Neural Network (ConvNet) [6]. Over the years, there have been efforts in investigating what other transformations result in useful hand-crafted priors

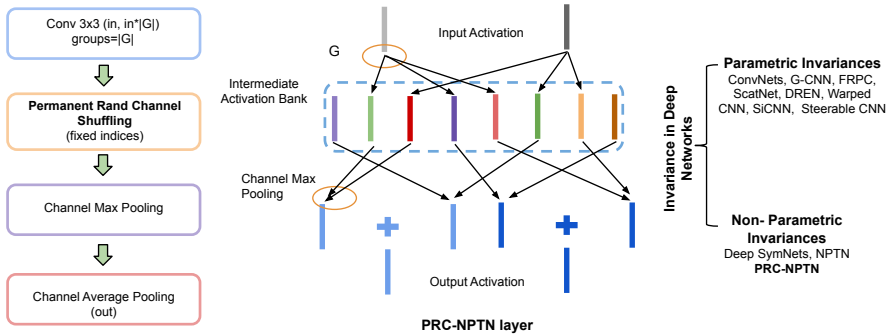


Figure 1: **Left:** Operations comprising the PRC-NPTN layer. The number of input and output channels in the Conv layer is  $(inch)$  and  $(inch * |G|)$  respectively.  $G$  is the number of filters (linear transformations learnt) for each input channel. The key operation proposed is the Permanent Random Channel shuffling operation with a fixed index mapping for every forward pass. This indexing or connectome is initialized randomly during network initialization. **Center:** Architecture of the PRC-NPTN layer. Each input channel is convolved with a number of filters (parameterized by  $G$ ). Each of the resultant activation maps is connected to a one of the channel max pooling units randomly (initialized once, fixed during training and testing). Each channel pooling unit pools over a fixed random support of a size parameterized by  $CMP$ . **Right:** Explicit invariances enforced within deep networks in prior art are mostly parametric in nature. The important problem of learning *non-parametric* invariances from data has not received a lot of attention.

in data such as rotation and scale [6, 9, 6, 10, 16, 19, 28]. It is important to note however that these methods ultimately were limited to hand-crafted invariances assumed to be useful for the task at hand.

**Motivation of this Study:** In this study, we motivate and investigate one possible architecture that can learn invariances towards multiple transformations from data itself. At the heart of the architecture is the structure called the *permanent random* connectome. This simply refers to a channel shuffling layer that uses a fixed shuffling schedule throughout the life of the network (including training and testing) resulting in a *permanent* connectome. Importantly however, this shuffling indexing is chosen at random at the initialization of the network. Thereby leading to the layer being referred to as the *permanent random* connectome. We find that layers utilizing the permanent random operation allow architectures to learn multiple invariances efficiently from data itself. Our motivation also loosely stems from observations regarding connectomes in the cortex.

**Prior Art Learning Invariances from Data or using Random Connectomes.** A different class of architectures that have been recently proposed explicitly attempt to *learn* the transformation invariances directly from the data, with the only inductive bias being the *structure or architecture* that allows them to do so. One of the earliest attempts using back-propagation was the SymNet [4], which utilized kernel based interpolation to learn general invariances. Although given the interesting nature of the study, the method was limited in scalability. Spatial Transformer Networks [12] were also designed to learn activation normalization from data itself, however the transformation invariance learned was parametric in nature. A more recent effort was through the introduction of the Transformation Network paradigm [22]. Non-Parametric Transformation Networks (NPTN) were introduced as a generalization of the convolution layer to model general symmetries from data [22]. It was also introduced as an alternate direction of network development other than skip connections.

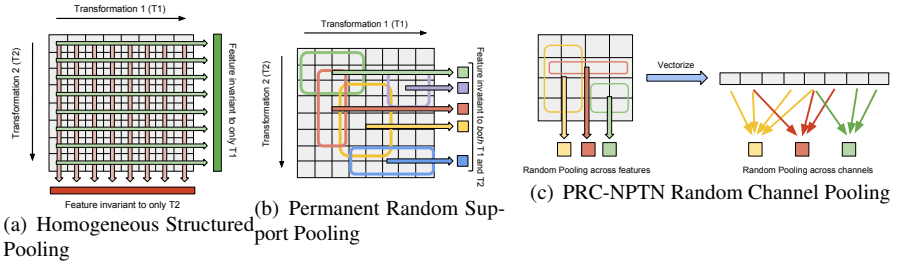


Figure 2: **(a) Homogeneous Structured Pooling** pools across the entire range of transformations of the *same kind* leading to feature vectors invariant only to that particular transformations. Here, two *distinct* feature vectors are invariant to transformation  $T_1$  and  $T_2$  independently. **(b) Heterogeneous Random Support Pooling** pools across randomly selected ranges of multiple transformations *simultaneously*. The pooling supports are defined during initialization and remain fixed during training and testing. This results in a *single* feature vector that is invariant to multiple transformations simultaneously. Here, each colored box defines the support of the pooling and pools across features only inside the boxed region leading to one single feature. **(c) Vectorized Random Support Pooling** extends this idea to convolutional networks, where one realizes that the random support pooling on the feature grid (on the left) is equivalent to random support pooling of the vectorized grid. Each element of the vector (on the right) now represents a single *channel* in a convolutional network and hence random support pooling in PRC-NPTNs occur across channels.

The convolution operation followed by pooling was re-framed as pooling across outputs from the translated versions of a filter. Translation forming a unitary group generates invariance through group symmetry [10]. The NPTN framework has the important advantage of learning general invariances without any change in architecture while being scalable. Given this is an important open problem, we introduce an extension of the Transformation Network (TN) paradigm with an enhanced ability to learn non-parametric invariances through permanent random connectivity. There have been many seminal works that have indeed explored the role of *temporary* random connections in deep networks such as DropOut [26], DropConnect [29] and Stochastic Pooling [54]. However, unlike the proposed approach, the connections in these networks randomly change at *every* forward pass, hence are temporary. More recently, random permanent connections were explored for large scale architectures [53]. It is important however to note that the basic unit of computation, the convolutional layer, remained unchanged. Our study explores permanent random connectomes *within* the convolutional layer itself, and explores how it can learn non-parametric invariances to multiple transformations simultaneously in a simple manner. We briefly visit other deep architectures that have been proposed over the years in the supplementary.

**Relaxed Biological Motivation for Randomly Initialized Connectomes.** Although not central to our motivation, the observation that the cortex lacks precise *local* pathways for back-propagation provided the initial inspiration for this study. It further garnered pull from the observation that random unstructured local connections are indeed common in many parts of the cortex [9, 24]. Moreover, it has been shown that orientation selectivity can arise in the visual cortex even through local random connections [10]. Though we do not explore these biological connections in more detail, it is still an interesting observation. There has also been some interesting work which explored the use of random weight matrices for back propagation [18]. Here, the forward weight matrices were updated so as to fruitfully use the random weight matrices during back propagation. The motivation of the [18] study was to

address the biological implausibility of the transport of precise gradients through the cortex due to the lack of exact connections and pathways [8, 20, 27, 32]. The common presence of random connections in the cortex at a *local* level leads us to ask: Is it possible that such locally random connectomes improve generalization in deep networks? We provide evidence for answering this question in the positive.

**Contributions.** 1) We motivate permanent random connectomes from the perspective of learning invariance to multiple transformations directly from data. The fundamental problem of learning non-parametric invariances in perception has not received enough attention. We present an architectural prior capable of such a task with loose biological motivation. 2) We present a theoretical result on learning invariances to transformations which do not obey a group structure in contrast to prior work. 3) We provide results on learning invariances to individual and multiple transformations in data without any change in architecture whatsoever. Further, we demonstrate improvements in generalization while using PRC-NPTN as a drop in replacement to conv layers in DenseNets. 4) Finally, as an engineering effort, we develop fast and efficient CUDA kernels for random channel pooling which result in efficient implementations of PRC-NPTNs in terms of computational speed and memory requirements compared to traditional Pytorch code.

## 2 Permanent Random Connectome NPTNs

We begin by motivating permanent random connectomes from the perspective of selecting each specific support for pooling. We find that permanent random channel pooling invokes invariance to multiple transformations simultaneously. Investigating idea of pooling across transformations to invoke invariance, permanent random pooling emerges naturally. As part of our contribution, we present a theoretical result which confirms a long standing intuition that max pooling invokes invariance.

**Invoking Invariance through Max Pooling.** In previous years a number of theories have emerged on the mechanics of generating invariance through pooling. [11, 12] develop a framework in which the transformations are modelled as a group comprised of unitary operators denoted by  $\{g \in \mathcal{G}\}$ . These operators transform a given filter  $w$  through the operation  $gw^1$ , following which the dot-product between these transformed filters and an novel input  $x$  is measured through  $\langle x, gw \rangle$ . It was shown by [11] that any moment such as the mean or max (infinite moment) of the distribution of these dot-products in the set  $\{\langle x, gw \rangle | g \in \mathcal{G}\}$  is an invariant. These invariants will exhibit robustness to the transformation in  $\mathcal{G}$  encoded by the transformed filters in practice, as confirmed by [13, 14]. Though this framework did not make any assumptions on the distribution of the dot-products, it imposed the restricting assumption of group symmetry on the transformations. We now show that invariance can be invoked even when *avoiding the assumption that the transformations in  $\mathcal{G}$  need to form a group*. Nonetheless, we assume that the distribution of the dot-product  $\langle x, gw \rangle$  is uniform and thus we have the following result<sup>2</sup>.

**Lemma 2.1.** (*Invariance Property*) *Assume a novel test input  $x$  and a filter  $w$  both fixed vectors  $\in \mathbb{R}^d$ . Further, let  $g$  denote a random variable representing unitary operators with*

<sup>1</sup>The action of the group element  $g$  on  $w$  is denoted by  $gw$  to promote clarity.

<sup>2</sup>We provide a proof in the supplementary and thank Purvasha Chakravarti at CMU for the proof.. The assumption of the distribution being uniform is meant to provide insight into the general behavior of the max pooling operation, rather than a statement that deep learning features are uniformly distributed.

some distribution. Finally, let  $\Upsilon(x) = \langle x, gw \rangle$ , with  $\Upsilon(x) \sim U(a, b)$  i.e. a Uniform distribution between  $a$  and  $b$ . Then, we have

$$\text{Var}(\max \Upsilon(x)) \leq \text{Var}(\Upsilon(x)) = \text{Var}(\langle g^{-1}x, w \rangle)$$

This result is interesting because it shows that the max operation of the dot-products has less variance due to  $g$  than the pre-pooled features. Though this is largely known empirical result, a concrete proof for invoking invariance was so far missing. More importantly, it bypasses the need for a group structure on the nuisance transformations  $\mathcal{G}$ . Practical studies such as [17, 23] had ignored the effects of non-group structure in theory while demonstrating effective empirical results. Also note that the variance of the max is less than the variance of the quantity  $\langle g^{-1}x, w \rangle$ , which implies that  $\max \Upsilon(x)$  is more robust to  $g$  even in test, though it has never observed  $gx$ . This useful property is due to the unitarity of  $g$ .

**Connection to Deep Networks.** PRC-NPTN as we will show, perform max pooling across *channels* not space, to invoke invariance. In the framework  $\langle x, gw \rangle$ ,  $w$  would be one convolution filter with  $g$  being the transformed version of it. Note that this modelling is done only to satisfy a theoretical construction, we do not actually transform filters in practice. All transformed filters are learnt through backpropagation. This framework is already utilized in ConvNets. For instance, ConvNets [24] pool only across translations (convolution operation itself followed by *spatial* max pooling implies  $g$  to be translation).

**Invoking Invariance through Channel Pooling in Deep Networks.** Consider a grid of features that have been obtained through a dot product  $\langle x, gw \rangle$  (for instance from a convolution activation map, where the grid is simply populated with each  $k \times k \times 1$  filter, not  $k \times k \times c$ ) (see Fig. 2(a)). Assume that along the two axes of the grid, two different kinds of transformation are acted.  $T_1$  along the horizontal axis and  $T_2$  along the vertical.  $T_1 = g_1(\cdot; \theta_1)$  where  $g_1$  is a transformation parameterized by  $\theta_1$  that acts on  $w$  and similarly  $T_2 = g_2(\cdot; \theta_2)$ . Now, pooling homogeneously across one axis invokes invariance *only* to the corresponding  $g$  (for a more in depth analysis see [10]). Similarly, pooling along  $T_2$  only will result in a feature vector (Feature 2) invariant *only* to  $T_2$ . These representations (Feature 1 and 2) have complimentary invariances and can be used for complimentary tasks *e.g.* face recognition (invariant to pose) versus pose estimation (invariant to subject). This approach has one major limitation that this scales linearly with the number of transformations which is impractical. One therefore would need features that are invariant to multiple transformations simultaneously. A simple yet effective approach is to pool along all axes thereby being invariant to all transformations simultaneously. However, doing so will result in a degenerative feature (that is invariant to everything and discriminative to nothing). Therefore, the key is to limit the *range* of pooling performed for each transformation.

**Choosing the Support for Pooling at Random: Permanent Random Connectomes.** A solution to trivial feature problem described above, is to limit the *range* or support of pooling as illustrated in Fig.2(b). One simple way of selecting such a support for pooling is at random. This selection would happen only once during initialization of the network (or any other model), and will remain fixed through training and testing. In order to increase the selectivity of such features, multiple such pooling units are needed with such a randomly initialized support [11, 25]. These multiple pooling units together form the feature that is invariant to multiple transformations simultaneously, which improves generalization as we find in our experiments. This is called heterogeneous pooling and Fig. 2(b) illustrates this more concretely. We therefore find that permanent random pooling is motivated naturally through the need to attain invariance to multiple transformations simultaneously.

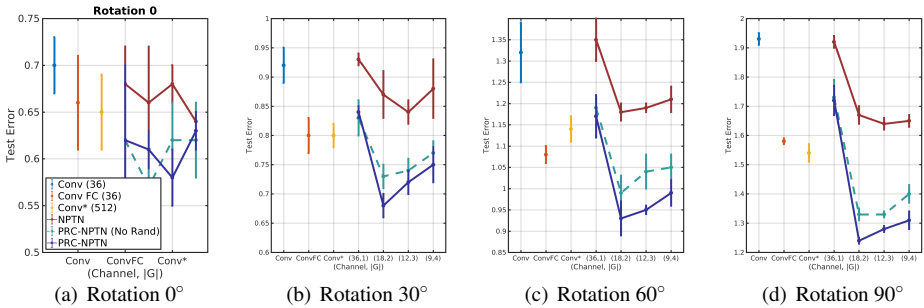


Figure 3: **Individual Transformation Results:** Test error statistics with mean and standard deviation on MNIST with progressively extreme transformations with a) **random rotations** and b) **random pixel shifts** (in supplementary). Only for PRC-NPTN and NPTN the brackets indicate the number of channels in the layer 1 and  $G$ . ConvNet FC denotes the addition of a 2-layered pooling  $1 \times 1$  pooling network after every layer. Note that for this experiment,  $\text{CMP} = |G|$ . Permanent Random Connectomes help with achieving better generalization despite increased nuisance transformations. We provide experiments on CIFAR 10 using DenseNets in the supplementary.

**The PRC-NPTN layer.** Fig. 1 shows the the architecture of a single PRC-NPTN layer<sup>3</sup>. The PRC-NPTN layer consists of a set of  $N_{in} \times G$  filters of size  $k \times k$  where  $N_{in}$  is the number of input channels and  $G$  is the number of filters connected to each input channel. More specifically, each of the  $N_{in}$  input channels connects to  $|G|$  filters. Then, a number of channel max pooling units randomly select a fixed number of activation maps to pool over. This is parameterized by Channel Max Pool (CMP). Note that this random support selection for pooling is the reason a PRC-NPTN layer contains a permanent random connectome. These pooling supports once initialized do not change through training or testing. Once max pooling over CMP activation maps completes, the resultant tensor is average pooled across channels with a average pool size such that the desired number of outputs is obtained. After the CMP units, the output is finally fed through a two layered network with the same number of channels with  $1 \times 1$  kernels, which we call a pooling network. This small pooling network helps in selecting non-linear combinations of the invariant nodes generated through the CMP operation, thereby enriching feature combinations downstream. For experimental rigor, we also benchmark against the baseline ConvNet supplemented with this  $1 \times 1$  pooling network.

### 3 Empirical Evaluation and Discussion

**Goal.** The goal of our evaluation study is to demonstrate PRC-NPTNs as capable of learning transformations from data and to showcase improvements in generalization in supervised classification over relevant baselines. The goal is not in fact, to compete with the state of the art approaches for any dataset.

**General Experimental Settings.** For all experiments, we run all models for 300 epochs trained using SGD. The initial learning rate was kept at 0.1 and decreased by 10 at 50% and 75% epoch completion. Momentum was kept at 0.9 with a weight decay of  $10^{-5}$ . Batch size

<sup>3</sup>We provide pseudo-code in the supplementary

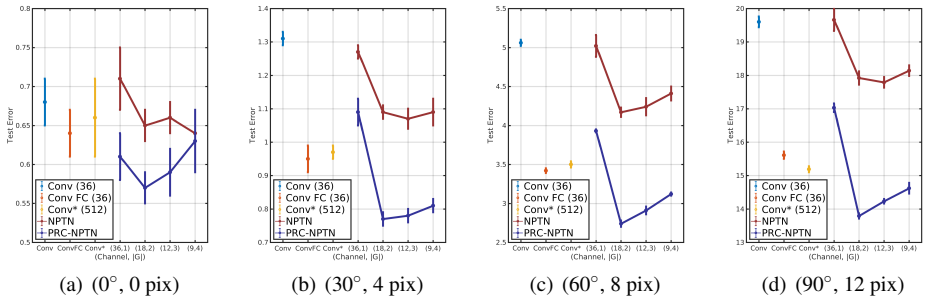


Figure 4: **Simultaneous Transformation Results:** Test error statistics with mean and standard deviation on MNIST with progressively extreme transformations with **random rotations** and **random pixel shifts simultaneously**. For PRC-NPTN and NPTN the brackets indicate the number of channels in the layer 1 and  $G$ . Note that for this experiment,  $\text{CMP}=|G|$ . We provide experiments on CIFAR 10 with DenseNets in the supplementary.

was kept at 64 for both MNIST and ETH-80<sup>4</sup>. For the MNIST experiments, gradients were clipped to norm 1. Each block for all baselines for ConvNet and PRC-NPTN had either a convolution layer or PRC-NPTN layer followed by batch normalization, PReLU and spatial max pooling. The convolutional kernel size for all models was kept at  $5 \times 5$  for all MNIST experiments and  $3 \times 3$  for all other models. Spatial max pooling of size  $3 \times 3$  was performed after every layer, BN and PReLU for MNIST models.

**Limitations in Typical Implementations and Developing Faster Kernels.** Our implementation with traditional PyTorch still suffered from heavy GPU memory use and slower run times despite optimizing code at the PyTorch abstraction level. The key bottleneck in computational and memory efficiency was found to be the randomized channel pooling operation. The issue was addressed by developing CUDA kernels that performed pooling on non-contiguous blocks of memory without creating copies of the same. This allowed for faster non-contiguous pooling over feature and activation maps with a significant reduction in memory usage. The operation was built as a CUDA-kernel that is interfaced with PyTorch through CuPy. This engineering effort is part of our contribution and we demonstrate significant improvements in memory and computational efficiency in our experiments. We show improvements in memory and speed as different aspects of the PRC-NPTN networks are changed in the supplementary<sup>5</sup>. We observe a consistent speed up of at least 1.5x and a significant reduction in memory usage.

**Efficacy in Learning Arbitrary and Unknown Transformations Invariances from Data.** We evaluate on one of the most important tasks of any perception system, *i.e.* being invariant to nuisance transformations *learned* from the data itself. Most other architectures based on vanilla ConvNets learn these invariances through the implicit neural network functional map rather than explicitly through the architecture as PRC-NPTNs. Moreover, most previous approaches needed hand crafted architectures to handle different transformations. We benchmark our networks based on tasks where nuisance transformations such as large amounts of in-plane rotation and translation are steadily increased, with no change in architecture whatsoever. For this purpose, we utilize MNIST where it is straightforward to add such transformations without any artifacts.

<sup>4</sup>We provide experiments on CIFAR-10 in the supplementary.

<sup>5</sup>We provide details of the architecture, benchmarking and additional experiments in the supplementary.



Architecture	
ConvNet B	C(12) - C(24) - C(48) - C(48) - GAP - FC(300) - FC(200) - FC(8)
NPTN-large B	C(12) - NPTN(24) - NPTN(48) - NPTN(48) - GAP - FC(300) - FC(200) - FC(8)
NPTN-small B	C(12) - NPTN(8) - NPTN(24) - NPTN(48) - GAP - FC(300) - FC(200) - FC(8)
PRC-NPTN B	C(12) - PRC(24) - PRC(48) - PRC(48) - GAP - FC(300) - FC(200) - FC(8)
ConvNet C	C(12) - C(24) - C(48) - C(64) - C(128) - GAP - FC(8)
PRC-NPTN C	C(12) - PRC(24) - PRC(48) - PRC(64) - PRC(128) - GAP - FC(8)

Table 1: **Architectures tested on ETH-80.** C - convolution layer, FC - fully connected layer, PRC - PRC-NPTN layer, NPTN - NPTN layer, GAP - global average pooling layer. Every Conv, NPTN and PRC-NPTN layer was followed by a spatial pooling layer of kernel size 2 except the last laer before the GAP. The  $1 \times 1$  versions of these architectures have a  $1 \times$  conv layer after every  $3 \times 3$  layer except the first C(12) layer. ConvNet B was designed to be similar to the architecture explored in [13] however with more layers. ConvNet C was designed to be more aligned with modern architecture choices such as global average pooling followed by just one FC layer.

We benchmark on such a task as described in [24] and for fair comparisons, we follow the exact same protocol. We train *and* test on MNIST augmented with progressively increasing transformations *i.e.* **1)** extreme random translations (up to 12 pixels in a 28 by 28 image), **2)** extreme random rotations (up to  $90^\circ$  rotations) and finally **3)** both transformations simultaneously. *Both* train and test data were augmented randomly for every sample leading to an increase in overall complexity of the problem. No architecture was altered in anyway between the two transformations *i.e.* they were not designed to specifically handle either. The same architecture for all networks is expected to learn invariances directly from data unlike prior art where such invariances are hand crafted in [8, 11, 16, 25, 28, 33].

For this experiment, we utilize a two layered network with the intermediate layer 1 having up to 36 channels and layer 2 having exactly 16 channels for all networks (similar to the architectures in [24]) except a wider ConvNet baseline with 512 channels. All ConvNet, NPTN and PRC-NPTN models have the similar number of parameters (except the ConvNet with 512 channels). For PRC-NPTN, the number of channels in layer 1 was decreased from 36, through to 9 while  $|G|$  was increased in order to maintain similar number of parameters. All PRC-NPTN networks have a two layered  $1 \times 1$  pooling network with same number of channels as that layer. For a fair benchmark, Convnet FC has 2 two-layered pooling networks with 36 channels each. Average test errors are reported over 5 runs for all networks.

**Discussion.** We present all test errors for this experiment in Fig. 3 and Fig. 4<sup>6</sup>. From both figures, it is clear that as more nuisance transformations act on the data, PRC-NPTN networks outperform other baselines with the same number of parameters. In fact, even with significantly more parameters, ConvNet-512 performs worse than PRCN-NPTN on this task for all settings. Since the testing data has nuisance transformations similar to the training data, the only way for a model to perform well is to learn invariance to these transformations. It is also interesting to observe that permanent *random* connectomes do indeed help with generalization. Indeed, without randomization the performance of PRCN-NPTNs drop substantially. The performance improvement of PRC-NPTN also increases with nuisance transformations, showcasing the benefits arising from modelling such invariances. This is

<sup>6</sup>We display only the (12, 3) configuration for NPTN as it performed the best. The translation results and more benchmarks with NPTNs are provided in the supplementary. We obtain similar performnce improvements with extreme translation as well.



Method (Protocol 1)	Accuracy (%)	#params	Factor	#filters	Reduction
ConvNet [13]	93.69	1.4M		230	-
NPTN* [22]	96.2	1.4M		230	-
ConvNet B	95.61	110K	1×	3780	1×
ConvNet B (1 × 1)	94.54	115K	0.95×	3780	1×
NPTN-large B (3)	94.63	189K	0.58×	11268	0.33×
NPTN-small B (3)	95.09	120K	0.91×	4356	0.86×
PRC-NPTN B (8, 2)	<b>96.72</b>	<b>97K</b>	<b>1.13×</b>	<b>708</b>	<b>5.33×</b>
ConvNet C	93.90	116K	1×	12740	1×
ConvNet C (1 × 1)	95.98	138K	0.84×	12740	1×
PRC-NPTN C (8, 2)	95.93	64K	1.81×	<b>1220</b>	<b>10.44×</b>
PRC-NPTN C (8, 4)	<b>96.40</b>	<b>39K</b>	<b>2.97×</b>	<b>1220</b>	<b>10.44×</b>

Table 2: **Test accuracy on ETH-80 Protocol 1.** \* indicates the result was obtained from the corresponding paper, and is not on the split used for our experiments. For NPTN is number in the bracket denotes  $|G|$ , for PRC-NPTN the numbers denote  $|G|$  and CMP respectively.

particularly apparent from Fig. 4, where the two simultaneous nuisance transformations pose a significant challenge. Yet, as the transformations increase, the performance improvements increase as well.

**Evaluation on the ETH-80 dataset** The ETH-80 dataset was introduced in [15] as a benchmark to test models against 3D pose variation of various objects. The dataset contains 80 objects belonging to 8 different classes. Each object has images from different viewpoints on a hemisphere for a total of 41 images per object. The images were resized to  $50 \times 50$  following [13]. This dataset is perfectly poised to test how efficiently a model can learn invariance to 3D viewpoint variation.

**Protocol:** For this experiment, we follow the evaluation protocol as described in [13] 7. We randomly select 2,300 images to train and test on the rest. For a fair comparison we retrain the ConvNet described in [13]. We design two ConvNet architectures which reflect more modern architecture choices such as a smaller FC layer or having only a global average pooling after a number of conv layers. Table. 1 presents the architectures that we train for this experiment. Every conv layer (except the first conv layer *within* a PRC-NPTN layer) is followed by BatchNorm and ReLU. We train corresponding PRC-NPTN models that have fewer parameters. For these experiments with PRC-NPTN, we replace the average pooling across channels (not max pooling or CMP) with a  $1 \times 1$  convolution layer. We do this to explore the effect of weighted pooling instead of vanilla channel average pooling. To maintain a fair comparison, we compare against equivalent ConvNet baselines with an extra  $1 \times 1$  added. We also perform an ablation study with the randomization removed. All models were trained with Adam with a learning rate of 0.01 for 100 epochs and a batch size of 64. Each architecture was trained 10 separate times with the mean of the runs being reported. We showcase the results in Table. 2.

**Discussion:** We find that of the two different types of architectures that we explore, PRC-NPTNs outperform both corresponding ConvNet architectures. Further, they do so not only with fewer number of parameters, but also fewer number of  $3 \times 3$  filters. In fact, PRC-NPTN C for  $|G| = 8$  and CMP=4 outperforms the corresponding ConvNet C architectures with a  $2.97 \times$  reduction in the number of parameters and a  $10.44 \times$  reduction in the number

<sup>7</sup>We also present results on a harder protocol we devised in the supplementary

of  $3 \times 3$  convolution filters. Similarly, PRC-NPTNs outperform two architectures of NPTNs [22] both with significantly fewer parameters and  $3 \times 3$  filters. These results illustrate that PRC-NPTN can utilize filters and parameters more efficiently on a classification problem which requires 3D pose invariance. This efficiency we conjecture, is due to the fact that permanent random pooling results in an inductive bias that explicitly helps the learning of multiple invariances within the same layer thereby vastly increasing model capacity. Note that the almost 3X reduction in the number of parameters and 10X reduction in the number of filters is achieved without the use of any network pruning or post processing methods. Note that competing methods presented in [13] all perform comparably however with 1.4 million parameters each with the highest result being TIGradNet [13] at 95.1, HarmNet at 94.0 [50]. PRC-NPTN outperforms these methods which were designed to invoke invariances through inductive biases while using a fraction of the number of parameters.

## References

- [1] Fabio Anselmi, Joel Z Leibo, Lorenzo Rosasco, Jim Mutch, Andrea Tacchetti, and Tomaso Poggio. Unsupervised learning of invariant representations in hierarchical architectures. *arXiv preprint arXiv:1311.4158*, 2013.
- [2] Fabio Anselmi, Georgios Evangelopoulos, Lorenzo Rosasco, and Tomaso Poggio. Symmetry regularization. Technical report, Center for Brains, Minds and Machines (CBMM), 2017.
- [3] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International Conference on Machine Learning*, pages 2990–2999, 2016.
- [4] Taco S Cohen and Max Welling. Steerable cnns. *arXiv preprint arXiv:1612.08498*, 2016.
- [5] Joe Corey and Benjamin Scholl. Cortical selectivity through random connectivity. *Journal of Neuroscience*, 32(30):10103–10104, 2012.
- [6] Sander Dieleman, Kyle W Willett, and Joni Dambre. Rotation-invariant convolutional neural networks for galaxy morphology prediction. *Monthly notices of the royal astronomical society*, 450(2):1441–1459, 2015.
- [7] Robert Gens and Pedro M Domingos. Deep symmetry networks. In *Advances in neural information processing systems*, pages 2537–2545, 2014.
- [8] Stephen Grossberg. Competitive learning: From interactive activation to adaptive resonance. *Cognitive science*, 11(1):23–63, 1987.
- [9] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *Computer vision and pattern recognition, 2006 IEEE computer society conference on*, volume 2, pages 1735–1742. IEEE, 2006.
- [10] David Hansel and Carl van Vreeswijk. The mechanism of orientation selectivity in primary visual cortex without a functional map. *Journal of Neuroscience*, 32(12):4049–4064, 2012.

- [11] João F Henriques and Andrea Vedaldi. Warped convolutions: Efficient invariance to spatial transformations. In *International Conference on Machine Learning*, 2017.
- [12] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *Advances in Neural Information Processing Systems*, pages 2017–2025, 2015.
- [13] Renata Khasanova and Pascal Frossard. Graph-based isometry invariant representation learning. In *International Conference on Machine Learning*, pages 1847–1856, 2017.
- [14] 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.
- [15] Bastian Leibe and Bernt Schiele. Analyzing appearance and contour based methods for object categorization. In *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on*, volume 2, pages II–409. IEEE, 2003.
- [16] Junying Li, Zichen Yang, Haifeng Liu, and Deng Cai. Deep rotation equivariant network. *arXiv preprint arXiv:1705.08623*, 2017.
- [17] Qianli Liao, Joel Z Leibo, and Tomaso Poggio. Learning invariant representations and applications to face verification. In *Advances in Neural Information Processing Systems*, pages 3057–3065, 2013.
- [18] Timothy P Lillicrap, Daniel Cownden, Douglas B Tweed, and Colin J Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7:13276, 2016.
- [19] Stéphane Mallat and Irene Waldspurger. Deep learning by scattering. *arXiv preprint arXiv:1306.5532*, 2013.
- [20] Pietro Mazzoni, Richard A Andersen, and Michael I Jordan. A more biologically plausible learning rule for neural networks. *Proceedings of the National Academy of Sciences*, 88(10):4433–4437, 1991.
- [21] Dipan Pal, Ashwin Kannan, Gautam Arakalgud, and Marios Savvides. Max-margin invariant features from transformed unlabelled data. In *Advances in Neural Information Processing Systems 30*, pages 1438–1446, 2017.
- [22] Dipan K Pal and Marios Savvides. Non-parametric transformation networks for learning general invariances from data. *AAAI*, 2019.
- [23] Dipan K Pal, Felix Juefei-Xu, and Marios Savvides. Discriminative invariant kernel features: a bells-and-whistles-free approach to unsupervised face recognition and pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5590–5599, 2016.
- [24] Manuel Schottdorf, Wolfgang Keil, David Coppola, Leonard E White, and Fred Wolf. Random wiring, ganglion cell mosaics, and the functional architecture of the visual cortex. *PLoS computational biology*, 11(11):e1004602, 2015.

- [25] Laurent Sifre and Stéphane Mallat. Rotation, scaling and deformation invariant scattering for texture discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1233–1240, 2013.
- [26] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [27] David G Stork. Is backpropagation biologically plausible. In *International Joint Conference on Neural Networks*, volume 2, pages 241–246, 1989.
- [28] Damien Teney and Martial Hebert. Learning to extract motion from videos in convolutional neural networks. In *Asian Conference on Computer Vision*, pages 412–428. Springer, 2016.
- [29] Li Wan, Matthew Zeiler, Sixin Zhang, Yann Le Cun, and Rob Fergus. Regularization of neural networks using dropconnect. In *International Conference on Machine Learning*, pages 1058–1066, 2013.
- [30] Daniel E Worrall, Stephan J Garbin, Daniyar Turmukhambetov, and Gabriel J Brostow. Harmonic networks: Deep translation and rotation equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5028–5037, 2017.
- [31] Saining Xie, Alexander Kirillov, Ross Girshick, and Kaiming He. Exploring randomly wired neural networks for image recognition. *arXiv preprint arXiv:1904.01569*, 2019.
- [32] Xiaohui Xie and H Sebastian Seung. Equivalence of backpropagation and contrastive hebbian learning in a layered network. *Neural computation*, 15(2):441–454, 2003.
- [33] Yichong Xu, Tianjun Xiao, Jiaying Zhang, Kuiyuan Yang, and Zheng Zhang. Scale-invariant convolutional neural networks. *arXiv preprint arXiv:1411.6369*, 2014.
- [34] Matthew D Zeiler and Rob Fergus. Stochastic pooling for regularization of deep convolutional neural networks. *arXiv preprint arXiv:1301.3557*, 2013.