Building Efficient Deep Neural Networks with Unitary Group Convolutions

Ritchie Zhao   Yuwei Hu   Jordan Dotzel   Christopher De Sa   Zhiru Zhang
Cornell University

{rz252, yh457, jad443}@cornell.edu, cdesa@cs.cornell.edu, zhiruz@cornell.edu

Abstract

We propose unitary group convolutions (UGConvs), a building block for CNNs which compose a group convolution with unitary transforms in feature space to learn a richer set of representations than group convolution alone. UGConvs generalize two disparate ideas in CNN architecture, channel shuffling (i.e. ShuffleNet [29]) and block-circulant networks (i.e. CirCNN [6]), and provide unifying insights that lead to a deeper understanding of each technique. We experimentally demonstrate that dense unitary transforms can outperform channel shuffling in DNN accuracy. On the other hand, different dense transforms exhibit comparable accuracy performance. Based on these observations we propose HadaNet, a UGConv network using Hadamard transforms. HadaNets achieve similar accuracy to circulant networks with lower computation complexity, and better accuracy than ShuffleNets with the same number of parameters and floating-point multiplies.

1. Introduction

Deep convolutional neural networks (CNNs) have proven extremely successful at large-scale computer vision problems. Research over the past few years has made steady progress on improving CNN accuracy [26]. Concurrently, efforts have been made to reduce the number of parameters and floating-point multiplies (fpmuls) in CNNs. One major trend in this research space is the sparsification of layer connections. Early networks such as AlexNet [13] and VGG [19] exclusively utilize dense mappings, i.e. convolutional (conv) or fully-connected (FC) layers that form a weight connection between every input and every output feature. More advanced architectures such as Xception [2] and MobileNets [8] make use of *depthwise separable* convolutions, which consist of a sparse spatial mapping (depthwise convolution) and a dense cross-channel mapping (pointwise convolution). Even more recently, ShuffleNet [29] replaces the pointwise convolutions with sparse *group convolutions*, and additionally proposes a channel

shuffle to allow information to flow between groups. These changes to layer structure look to remove weight connections while retaining accuracy performance.

A different line of efficient CNNs research looks to train networks with circulant or block-circulant [1] weights [1, 20, 6, 22]. An n-by-n circulant matrix is dense and full-rank in general, but contains only n unique elements. Moreover, every circulant matrix $\mathbf{C}$ can be diagonalized by the normalized discrete Fourier matrix $\mathbf{F}$ as follows:

$$\mathbf{C} = \mathbf{F}^*\mathbf{DF} \qquad (1)$$

giving rise to an asymptotically faster algorithm for matrix multiplications involving $\mathbf{C}$ via the fast Fourier transform (FFT). By exploiting these properties of circulant weights, works in this area can also reduce CNN complexity and model size.

In this paper, we propose the concept of *unitary group convolution* (UGConv), defined as a building block for neural networks that combines a weight layer (most commonly a group convolution) with unitary transforms in feature space. We show that group convs with channel shuffle (ShuffleNet) and block-circulant networks (CirCNN) are specific instances of UGConvs. By unifying two different lines of work in CNN literature, we gain a deeper understanding into the basic underlying idea — that group convolutions exhibit improved learning ability when performed in a transformed feature basis. Through a series of experiments, we then investigate how different transforms and UGConv structures affect the learning performance. Specifically, our contributions are as follows:

1. We propose the concept of unitary group convolutions. We show that ShuffleNets and circulant networks, techniques from two disparate lines of research, are in fact both instances of UGConv networks. This lets us unify the conceptual insights of both works.

---

[1]In this paper, block-circulant, block-diagonal, etc. refers to matrices consisting of square submatrices which are circulant, diagonal, etc. This is different from the canonical definition of a block-diagonal matrix.

2. We evaluate how different unitary transforms affect learning performance. Our experiments show that when the weight layer is highly sparse (i.e. the number of groups is large), dense transforms outperform simple permutations.

3. We propose HadaNets, UGConv networks using the easy-to-compute Hadamard transform. HadaNets obtain similar accuracy as circulant networks at a lower computation complexity, and outperform ShuffleNets with identical parameter and fpmul counts.

2. Related Work

2.1. Depthwise Separable and Group Convolutions

In a traditional convolutional layer, each 3D filter must learn both spatial and cross-channel correlations. A depthwise separable convolution decouples this into two steps: a depthwise convolution which only performs spatial filtering, and a pointwise convolution which only learns cross-channel mappings. The idea originated in Sifre 2014 [18] and was subsequently popularized by networks like Xception [2] and MobileNets [8]. These and other examples showed that depthwise separable convolutions can outperform traditional convolutions using fewer parameters and fpmuls.

A group convolution divides the input and output features into mutually independent groups and performs a convolution in each one. Depthwise convs are specific cases of group convs with group size 1. Group convolutions were part of the original AlexNet, but only to facilitate training on multiple GPUs [13]; they gained popularity as a building block of efficient CNNs as part of ResNeXt [25] and ShuffleNet [29]. The latter proposed channel shuffling to promote cross-channel information flow, surpassing MobileNets in accuracy and parameter efficiency.

Interleaved group convolutions [28, 24, 21] is a line of work that studied interleaving group convs and channels shuffles, and showed how a specific combination of width and sparsity (i.e. number of groups) can maximize accuracy. Deep Roots [10] uses group convolutions with increasing group size deeper into the network to improve numerous existing models. Different from these works, we study the composition of group convs with dense unitary transforms.

2.2. Circulant and Block-Circulant Networks

An n-by-n circulant matrix requires only $O(n)$ storage space and $O(n \log n)$ operations for the matrix-vector product (see Equation (1)). Circulant weights can reduce the model size and computational complexity of CNNs in a deterministic manner. Cheng et al. in 2015 applies this to achieve 18x parameter reduction on AlexNet with only 0.7% Top-1 accuracy loss [1]. Other authors proposed variations of circulant structure. Moczulski et al.'s ACDC used cosine transforms to avoid complex values that arise with DFTs and added a second channel-wise filter [16]. Sindhwani et al. studied the superset of generalized Toeplitz-like matrices [20]. These works exclusively worked on structured FC layers.

More recently, Wang et al. [6, 23] proposed to use block-circulant matrices and applied them to both FC and convolutional layers. Block-circulant structure elegantly addresses the long-standing issue of non-square weight matrices. The same authors also leveraged the butterfly structure of the DFT to construct efficient accelerators for circulant nets in dedicated hardware [6, 22]. A more recent followup in this line of work proposed to use permuted block-diagonal matrices [4] in specialized hardware; this method bears resemblance to ShuffleNet in weight structure.

2.3. Random Projections and Hadamard Networks

Our study on random orthogonal and Hadamard transforms is partly inspired by the Fastfood transform [14] and its application to CNNs [27]. This work is a well-known example of using random embeddings and Hadamard transforms in machine learning.

A recent work from Devici et al. [5] used Hadamard-transformed images as CNN inputs. Their work differs significantly from ours; they applied a single 2D Hadamard on the input image to extract frequency features while we use Hadamard throughout the network for channel mixing.

3. Unitary Group Convolutions

The basic idea of a UGConv is a group convolution sandwiched between two unitary transforms in feature space. Figure 1(a) illustrates a UGConv block. A formal definition is as follows.

Let $\mathbf{X} = \{\mathbf{x}^{(i)}\}_{i=1}^{M}$ be a multi-dimensional input tensor with M input features. Each input feature $\mathbf{x}^{(i)}$ is itself a vector of length R. Define $\mathbf{X}_k$ as a depthwise slice, i.e. $\mathbf{X}_k$ is an M-length vector formed by the k-th elements of each feature vector. We use bracketed superscripts to indicate different features and subscripts to indicate elements in a feature. Similarly, let $\mathbf{Y} = \{\mathbf{y}^{(j)}\}_{j=1}^{N}$ be the tensor of N output features, each feature vector having length S. In addition, $\mathbf{W} = \{\mathbf{W}^{(ij)}\}_{i,j=1}^{M,N}$ is the weight tensor of $M \times N$ filters. We can now define an ordinary conv layer (Figure 2(a)) as:

$$\mathbf{y}^{(j)} = \sum_{i=1}^{M} \mathbf{x}^{(i)} * \mathbf{W}^{(ij)}$$

Finally, a unitary group convolution consists of three oper-
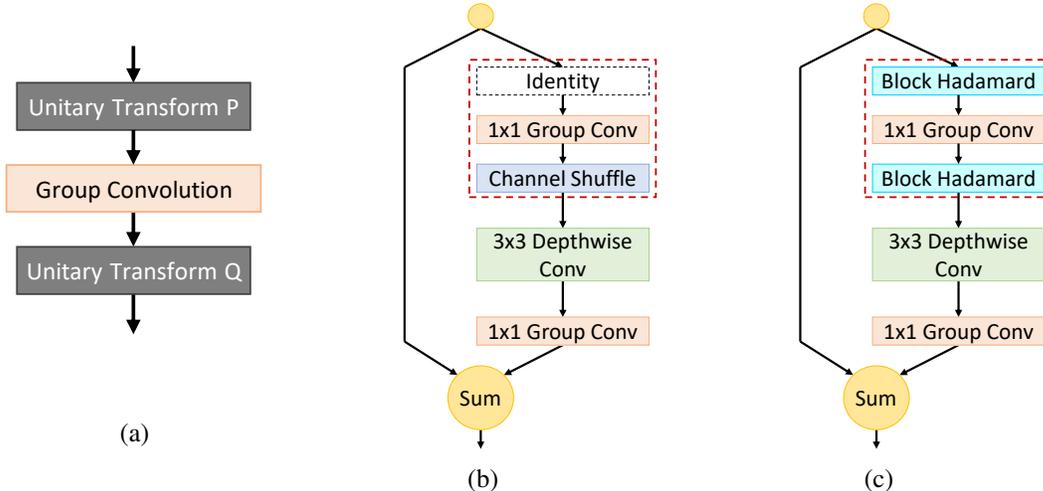
Figure 1. **CNN block architectures –** (a) a general block for unitary group convolutions; (b) a ShuffleNet block reproduced from the original paper [29]; (c) our proposed HadaNet variation. Note that both ShuffleNet and HadaNet blocks contain the UGConv pattern.

ations:

$$\tilde{\mathbf{X}}_k = \mathbf{P}\mathbf{X}_k$$

$$\tilde{\mathbf{y}}^{(g,j)} = \sum_{i=1}^{M/G} \tilde{\mathbf{x}}^{(g,i)} * \tilde{\mathbf{W}}^{(g,ij)} \qquad (2)$$

$$\mathbf{Y}_l = \mathbf{Q}\tilde{\mathbf{Y}}_l$$

The second equation expresses a group convolution with G groups, which simply performs an independent conv in each group. Figure 2(b) shows the weight layer associated with a group conv. $\mathbf{P}$ and $\mathbf{Q}$ are unitary matrix transforms applied element-wise over the features dimension, and tilde indicates the transformed features. Figure 1(a) illustrates the general idea of a UGConv layer. Note that the $\mathbf{P}$ and $\mathbf{Q}$ can be identity transforms, and thus UGConv is a generalization of group convolutions. We also note that unitary transforms preserve inner products; therefore their presence before and after a linear layer should not contribute to the vanishing gradients effect. Our notation and equations can be applied to both 2D convolutions (where R and S are the input and output width times height) and FC layers (where $R = S = 1$ and we have a 1×1 group convolution).

In Figure 2(c), we take the same group layer from 2(b) and relabel the indices to show that a group convolution is equivalent to a block-diagonal weight layer. Note there is no physical shuffling — the two layers are completely identical. We can show this mathematically as well. Equation 3 below describes a block-diagonal matrix multiplication:

$$\mathbf{y}^{(j,b)} = \sum_{i=1}^{B} \mathbf{x}^{(i,b)} \times \mathbf{W}^{(ij,b)} \qquad (3)$$

where i, j iterates over the blocks, b iterates down each diagonal, and B is the block size. It is easy to see that Equation 3 matches the weight multiply in Equation 2.

3.1. UGConv and ShuffleNet

ShuffleNet is a variant of the MobileNets architecture in which the pointwise convolutions (which take up 93.4% of the multiply-accumulate operations [29]) are converted into group convolutions. Figure 2(a) shows a pointwise convolution and 2(b) a group convolution where connections are only made between channels in the same group. Each square in the figure can represent a 2D or 1D feature/weight and the idea of groups applies to both conv and FC layers.

Compared to a regular weight layer, a group layer reduces parameters and fpmuls by a factor of B, where B is the number of groups (or equivalently, the block size in 2(b)). However, when multiple group layers are stacked together, the lack of connections between groups over many layers prevents the learning of any cross-group correlations. To address this, ShuffleNet shuffles channels between groups in a fixed, round-robin manner. For every group, the first channel is shuffled into group 1, the second channel into group 2, etc. The shuffle can be expressed as a fixed permutation in feature space, and ShuffleNets are thus an example of of UGConvNets where $\mathbf{P}$ is the identity matrix and $\mathbf{Q}$ is the permutation in Equation (2).

ShuffleNet's results demonstrate that it is beneficial to mix information across groups when stacking many group convolutional layers. However, the shuffling becomes less effective when the number of groups is large — e.g. when each channel is in its own group, shuffling the channels around does nothing. There may be other unitary transforms that can better accomplish cross-channel mixing.

3.2. UGConv and Circulant Networks

Circulant and block-circulant networks [6, 23] utilize layers that impose a block-circulant structure on
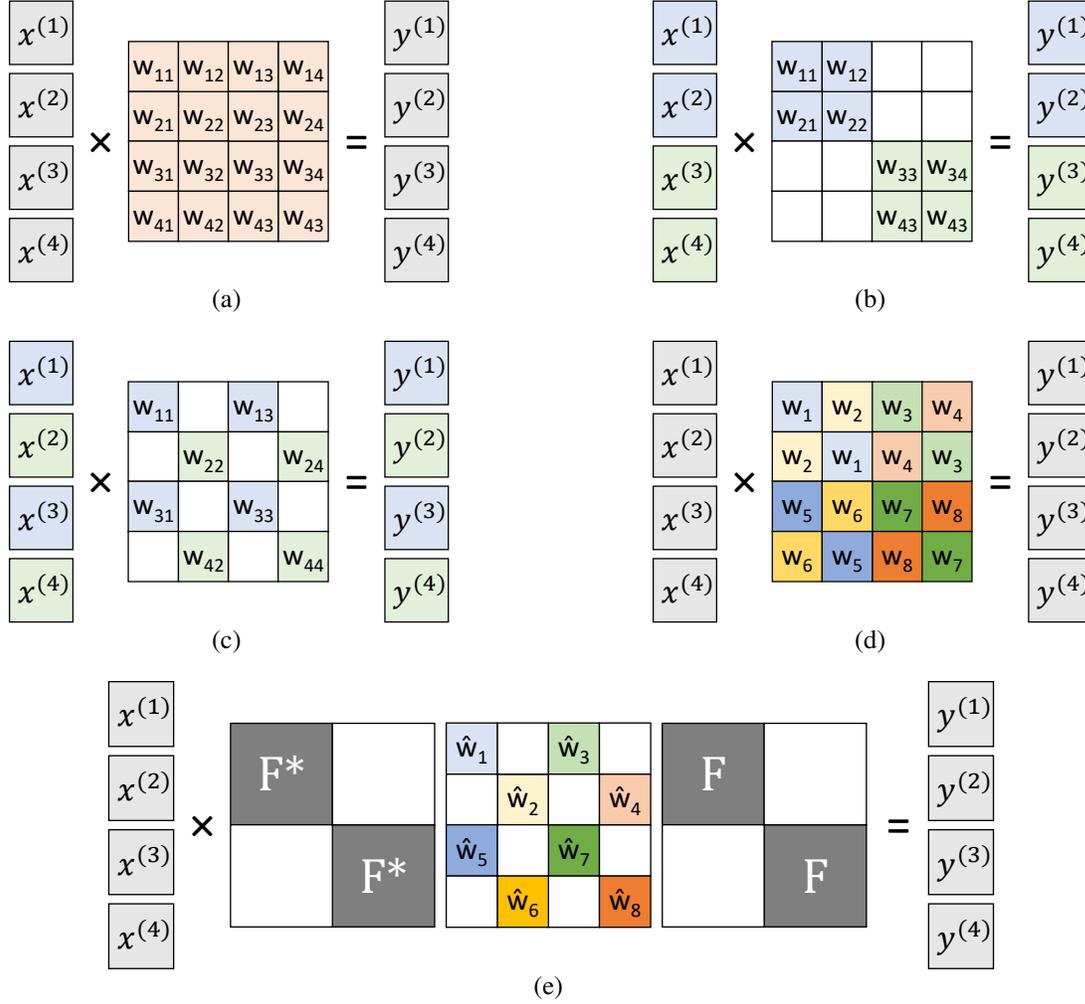
Figure 2. **Relation between group convs and circulant weights** – each square can be either a 2D feature/filter (conv layer) or a single activation/weight (FC layer). (a) regular layer with a dense weight matrix; (b) group layer with 2 groups; (c) same group layer reordered, showing a block-diagonal weight structure; (d) layer with a block-circulant weight matrix; (e) same circulant layer decomposed into block-DFTs and a block-diagonal weight matrix.

the weight tensor. For an FC layer, the 2D weight matrix becomes circulant. For a conv layer, the circulant structure is applied over the input and output channels axes. That is to say, given a 4D convolutional weight tensor with shape (`height`, `width`, `in_channels`, `out_channels`), each 2D slice of this tensor [`i,j,:,:`] becomes circulant.

Figure 2(d) shows a block-circulant layer where each 2-by-2 sub-block of the weight matrix is circulant. By Equation (1), each B-by-B circulant matrix can be decomposed into a diagonal matrix sandwiched between a B-length DFT and a corresponding IDFT. In Figure 2(e), each B-by-B sub-block is diagonalized in this fashion. We use tilde to indicate weight values in the DFT-transformed space. The resulting weight structure is block-diagonal, and the weight

layer sits between block-DFT transforms. We know from the previous section that block-diagonal weights correspond to group convolutions. Therefore, *a block-circulant layer is just a group convolution in a transformed feature space.* This of course fits the definition of a UGConv, with $\mathbf{P}$ and $\mathbf{Q}$ being block-DFT/IDFT transforms. Note that these DFTs are applied along the channels, and so circulant networks are *not* examining the spatial frequency components of the image.

We make a few additional notes about CirCNN's block-circulant layers. First, the size of the diagonal blocks B is equal to the number of groups (not the group size). Thus each B-length DFT is computed over a single channel from each of the B groups and fully mixes information between every group. Second, though our example uses a square

matrix, non-square block-circulant matrices can be diagonalized without issue. For a layer with M inputs and N outputs, the block-IDFT on the input $\mathbf{Q}$ is M-by-M while the block-DFT on the output is N-by-N. The DFT-length remains B in both cases. Here although $\mathbf{Q} \neq \mathbf{P}^{-1}$, each sub-block along the diagonal in $\mathbf{Q}$ is the inverse of the corresponding sub-block in $\mathbf{P}$. In this paper we say that $\mathbf{Q}$ is the *block-inverse* of $\mathbf{P}$.

Because $\mathbf{P}$ and $\mathbf{Q}$ are block-inverses, if we directly stack multiple such blocks many of the transforms will cancel out. However, practical DNNs include batch norm and/or non-linearities between linear layers. The block-DFTs (and orthogonal transforms in general) do not commute with channel-wise or pointwise operations, which prevents trivial cancellation. However, note that channel shuffles *do* commute and cancel out in this manner.

3.3. Discussion of UGConvs

We have provided two specific examples from literature (ShuffleNet [29] and CirCNN [6]) which combine a structured sparse weight layer (group convolution) with unitary transforms. The transforms help to improve cross-channel representation learning without adding additional parameters. However, the two techniques have important differences. ShuffleNet focuses on a very lightweight transform, hence fixed channel permutation. Permutations do not change the level of sparsity in the UGConv layer. On the other hand, CirCNN uses dense block-DFTs implicitly defined by the circulant structure, and the weight connections are dense.

We hypothesize that the representation learning capability of a UGConv layer is a function of both the sparsity of the weight matrix as well as the complexity of the transform. An unstructured dense weight matrix offers the best learning capability; grouping introduces sparsity and degrades cross-channel learning performance, some of which can be recovered via transforms. Because dense transforms create dense weight structures (i.e. circulant weights), we believe they enable the weight layer to learn a richer set of representations compared to simple permutations. When the weight sparsity is low (i.e. few groups), the difference may not be large enough to affect network accuracy. However, we expect dense transforms to outperform shuffling in DNNs using a large number of groups.

Another key difference is that ShuffleNet applies channel shuffle on only one side of the weight layer, while CirCNN effectively applies transformations on both sides. We use the terms 1-sided and 2-sided UGConvs to refer to these two cases, and test both in our experiments.

3.4. The Hadamard Transform

One drawback of dense transforms compared to shuffling is that they require computational overhead. Consider

Table 1. **Hadamard vs. Discrete Fourier transforms** – The entries of the DFT matrix are the complex roots of unity. The entries of the Hadamard matrix are $+1$ or -1. The last column shows the structure of $\mathbf{P}^*\mathbf{DP}$ where $\mathbf{D}$ is a diagonal matrix and $\mathbf{P}$ is the transform; differences are bolded.

Transform	Matrix				Wt Structure			
DFT	1	1	1	1	a	b	c	d
	1	ω	ω^2	ω^3	d	a	b	c
	1	ω^2	ω^4	ω^6	c	d	a	b
	1	ω^3	ω^6	ω^9	b	c	d	a
Hadamard	1	1	1	1	a	b	c	d
	1	-1	1	-1	$\mathbf{b}$	a	$\mathbf{d}$	c
	1	1	-1	-1	c	d	a	b
	1	-1	-1	1	$\mathbf{d}$	c	$\mathbf{b}$	a

a group layer with N input and output channels and group size B. The weight layer requires N^2/B fpmuls. Adding a general block-orthogonal transform like in Figure 2(e) incurs an overhead of NB fpmuls, and using a 'fast' transform like block-DFT requires $O(N \log B)$ fpmuls.

This is reasonable when $N \gg B$, but we can do better. The Hadamard transform [17] is defined as a matrix containing only $+1/-1$ elements and whose rows and columns are mutually orthogonal. Table 1 shows a 4-by-4 Hadamard matrix. Because all coefficients have magnitude 1, the transform can be computed using adds/subtracts only (i.e. without multiplies). In addition, the Hadamard transform can be generated recursively like the Fourier transform. This means that (1) the transform can be computed without storing the coefficients; (2) a fast Hadamard transform (FHT) exists similar to the FFT to compute a B-length Hadamard transform in $O(B \log B)$ adds/subtracts [17]. These factors make the Hadamard transform far more efficient than general orthogonal transforms (or even the DFT) if implemented optimally. Further discussion on using the Hadamard transform in DNNs can be found in Section 4.4.

Our intention to use the Hadamard transform also motivates the question of whether different dense transforms achieve different learning performance. One observation is that the order of features in the channels dimension is random (disregarding cases such as flattening the spatial dimensions into channels). The lack of semantic information to exploit among channel orders means no transform should work extraordinarily well. Another piece of intuition comes from the structure imposed on each sub-block of weights. In Figure 2(d) we see that the DFT imposes a circulant structure. Table 1 shows this structure for the Hadamard matrix — it is very close to circulant. We thereby hypothesize that all dense orthogonal transforms achieve comparable statistical performance, averaged across many training runs.

Table 2. **Test error on a toy MNIST network –** the model is *not* meant to be practical. **P** and **Q** refer to the pre-conv and post-conv transformations, respectively. E.g. the **P** column applies **P** only while **Q** is set to identity. All values are averaged over 5 runs and the 90% confidence bounds are up to ±5%.

Layer 2	Layer 3	Layer 4	None	Rand Ortho			Rand Perm		
				P	Q	PQ	P	Q	PQ
20conv3	**20Gfc1**	10fc1	6%	**4%**	**4%**	**4%**	5%	6%	5%
20conv3	**20Gfc1**	10Gfc1	27%	10%	8%	**4%**	27%	26%	25%
20Gconv3	**20Gfc1**	10fc1	25%	**10%**	**10%**	**10%**	27%	20%	21%
20Gconv3	**20Gfc1**	10Gfc1	60%	23%	**17%**	20%	57%	55%	57%

4. Experimental Validation

We begin by presenting experiments on a toy MNIST network, followed by deeper CIFAR-10 models. These tests are meant to build up insights, not to construct practical networks. We then demonstrate the utility of the Hadamard transform using a ShuffleNet model on ImageNet taken from literature.

4.1. Dense Transforms vs. Shuffle

Our first experiment uses a tiny toy network trained on MNIST. This allows us to isolate the UGConv block and to compare the relative learning capabilities of dense orthogonal transforms and permutations in a simple setting. We stress that the goal is *not* to build a realistic classifier. The layer architecture is denoted below, where the number before the layer is the channels/units and the number after is the filter width:

$$10\text{conv3} - 20\text{conv3} - \mathbf{20\text{fc1}} - 10\text{fc1}$$

We perform 2×2 max pooling before each 3×3 conv layer, and a global average pool before the first FC layer. We use batch normalization and ReLU nonlinearity.

We convert the first FC layer of the network (20fc1, shown in bold) into a UGConv block (i.e. it becomes a grouped FC with transforms). The group number is equal to the number of channels to maximize sparsity. From this base architecture we derive three additional variations: (1) convert the preceding conv layer into group conv; (2) convert the following dense layer into group FC; (3) convert both surrounding layers into group layers. This tests the performance of transforms in the context of stacked group layers. Two types of transform are evaluated: randomly generated dense orthogonal and random permutation transforms. We test with both 1-sided (using one of **P** or **Q** and setting the other to identity) and 2-sided UGConvs ($\mathbf{P} = \mathbf{Q}^{-1}$). All results are the average over five runs, and regenerate the random transformation matrices between runs.

Table 2 shows our results. Due to the small size of the network, the 90% confidence bound for these values can be as large as ±5%. Nevertheless, differences between transforms are clearly demonstrated. When the 20fc is the only grouped layer in the network (row 1), transforms have little to no effect. However, when two or more group layers are stacked together, the dense orthogonal transforms achieve improved accuracy. Meanwhile, the permutations have no accuracy impact. This is a clear (albeit artificial) demonstration that when the number of groups is very large, dense transforms outperform permutations in learning ability.

Another interesting observation we make is that there is little difference between 1-sided and 2-sided transforms, regardless of whether the UGConv block is stacked before or after another group layer. For example, in Table 2 row 3, a dense orthogonal transform improves accuracy even when it is placed *after* both group layers. At first glance it may be surprising for the transform to affect layers before it, but keep in mind that the transform also affects the gradients on the backwards pass. Alternatively, we can understand the UGConv layer as a learnable structured weight layer formed by the composition of transforms and block-diagonal weights (see Section 3.2).

4.2. Evaluation of Different Transforms

We have shown that dense orthogonal transforms can improve over shuffles in small DNNs with large group sizes. To validate our results on more realistic architectures, we perform experiments on CIFAR-10 [12] using ResNet [7]. The basic building block of the ResNet architecture contains two 3×3 convs (see Figure 1(a)). We use UGConvs to replace the two convolutions in each block, as well as replace the 1×1 projection layers. The blocks in these ResNets are divided into three stages (S1, S2, S3), with later stages having more channels. The number of groups used in each stage is chosen to be a fixed ratio of the channels. Two models are tested: **ResNet-20** (3 block per stage) and **ResNet-56** (9 blocks per stage). We also experiment with the same high-level architecture but using the building block from ShuffleNet [29]. This block which contains two 1×1 convs and a 3×3 depthwise conv (see Figure 1(b)). Following ShuffleNet we apply transformations around the first 1×1 group conv only and make no changes to the second group conv. Again two models are tested: **ShuffleNet-29** (3 block per stage) and **ResNet-56** (6 blocks per stage).

Table 3. **Test error for UGConvs on CIFAR-10 –** The first three columns show the number of groups used in the three stages (S1-S3). The **Base** column shows the test error with no transforms, and the other columns show *improvement* in test error over this baseline. Some entries are blank due to insufficient time to complete the experiments.

	# of Groups			Base	1-sided Transforms			2-sided Transforms				Params
	S1	S2	S3		Shuffle	Hada	Ortho	Shuffle*	Fourier	Hada	Ortho	
ResNet-20	4	8	16	19.5%	3.3%	4.0%	4.0%	3.1%	4.1%	**4.2%**		25K
	8	16	32	23.8%	2.9%	4.3%	3.9%	4.1%	**5.4%**	**5.4%**		14K
ResNet-56	4	8	16	16.0%	4.0%	4.4%	4.2%	4.0%	**4.7%**	4.5%	4.6%	76K
	8	16	32	20.6%	5.4%	6.1%	6.4%	5.8%	7.1%	**7.2%**	6.8%	41K
ShuffleNet-29	4	8	16	18.3%	2.7%	2.4%	3.1%	3.8%	**4.9%**	4.5%		23K
	8	16	32	22.1%	0.6%	3.4%	3.6%	3.8%	**5.1%**	5.0%		17K
ShuffleNet-56	4	8	16	16.2%	3.6%	3.5%	3.4%	3.9%	4.6%	4.5%	**4.7%**	41K
	8	16	32	19.7%	4.3%	4.4%	4.9%	5.2%	**6.0%**	**6.0%**	**6.0%**	29K
Mean	4	8	16	17.5%	3.4%	3.6%	3.7%	3.7%	**4.6%**	4.4%		
	8	16	32	21.5%	3.3%	4.6%	5.0%	4.7%	**5.9%**	**5.9%**		

We use layer widths and training hyperparameters from [7] and make use of standard data augmentations: padding 8 pixels on each side and randomly cropping back to original size, combined with a random horizontal flip [7, 9, 15]. Each network is trained for 200 epochs, and we report the mean test error over the last 5 epochs.

Another goal of this experiment is to compare different transforms, including the Hadamard transform. We test the following: identity (None), ShuffleNet permutation (Shuffle), block-Hadamard (Hada), block-DFT (Fourier), or block-random-orthogonal (Ortho). The block transform follow the same scheme described in Section 3.2. For each transform, both 1-sided (letting $\mathbf{Q}$ be the transform and $\mathbf{P}$ identity) and 2-sided ($\mathbf{P}$ and $\mathbf{Q}$ are block-inverses) versions are tested where reasonable. The 1-sided DFT is left out because it introduced complex numbers into the network. For the 2-sided channel shuffle (Shuffle*), we set $\mathbf{P} = \mathbf{Q}$ to essentially perform additional shuffling; this is done since using block-inverse shuffle transforms will lead to trivial cancellation. All results are displayed in Table 3 — the error rate with no transforms is given first followed by the accuracy improvement achieved with each UGConv setup. Our base error rates are high for CIFAR-10 because group convolutions significantly compress the network

A key result from this experiment is that dense orthogonal transforms perform similarly in accuracy. Fourier, Hada, and Ortho obtain results which are within a spread of 0.4% in both 1-sided and 2-sided settings. On the other hand, the shuffle transforms (1 and 2-sided) clearly perform worse for larger group sizes. This confirms that different dense UGConvs have comparable learning performance, and that the Hadamard transform is comparable to the Fourier transform while being much easier to compute.

Another observation is that 2-sided transforms outper-

form their 1-sided variants, which is different from the MNIST data. Indeed, the performance of 1-sided shuffle and 1-sided dense transforms were comparable in many cases. We currently do not have an explanation for this effect. One speculation was that 2-sided transforms perform better when the number of input and output channels did not match. However, further testing with the small MNIST network showed that this did not appear to be the case.

Finally, note that the accuracy trends remained the same whether the transforms were applied to 3×3 group convs in ResNet or 1×1 group convs in ShuffleNet. This is evidence that spatial and cross-channel dependencies are effectively decoupled in convolutional layers, and the size of the filter does not affect the action of channel-space transforms.

4.3. Hadamard Networks on ImageNet

The data from previous sections point to the existence of two regimes: at low weight sparsity (i.e. small group numbers) a simple shuffle is sufficient to maximize accuracy. At large group numbers, however, dense transforms outperform shuffles. This section evaluates the 2-sided block-Hadamard transform against shuffle on ImageNet. Hadamard was chosen as it is more efficient than other dense unitary transforms, and ShuffleNet was used for comparison as it is highly related work and a strong baseline. We refer to networks using Hadamard UGConvs as HadaNets.

Due to hardware constraints, we chose small models with fairly large group size — this is the setting where dense transforms should perform the best compared to shuffle. We evaluate with a ResNet-18 following the ImageNet architecture from [7]. The ResNet uses group size 8 throughout the network. We also test with the ShuffleNet-x0.25 g8, which is the smallest ShuffleNet variant from [29]. This

Table 4. **Top-1 classification error on ImageNet –** we include data on both the original ShuffleNet (with our own code) and our pre-activation variation. Despite best efforts we could not reproduce literature results for ShuffleNet (52.7%). For each model we show the number of parameters and fpmuls, as well as the overhead in additions from the Hadamard transform.

	Shuffle	Hada	Delta	Params	FPmuls	Hada Adds
ResNet-18 g8	46.4%	**44.6%**	(-1.8%)	1.9M	330M	3.8M
ShuffleNet-x0.25 g8 (Our impl.)	58.1%	**57.2%**	(-0.9%)	0.46M	17M	0.95M
ShuffleNet-x0.25 g8 (pre-act)	54.5%	**53.9%**	(-0.6%)	0.46M	17M	0.95M

network has 50 layers and also uses 8 groups. Each network was trained following the hyperparameters and learning rate schedule described in the respective paper. We compare 1-sided shuffle to 2-sided block-Hadamard (note that ShuffleNet from literature already contains the 1-sided shuffle). The residual blocks for ShuffleNet and HadaNet are in Figure 1(b) and 1(c). All results are displayed in Table 4.

Despite best efforts, we were unable to replicate the performance of ShuffleNet-x0.25-g8 reported in the original paper. The reported Top-1 error is 52.7% (Table 2 in [29]) while we obtained a much worse 58.1%. We changed the residual block to a pre-activation shortcuts and was able to improve this to 54.5%. Both versions are reported.

The results demonstrate that the Hadamard transform can indeed outperform shuffling in terms of accuracy on large scale datasets. ResNet-18 with group convolutions is a non-standard model — it mostly serves to show that the trends observed using ResNets on CIFAR-10 appear to carry over to ImageNet. On the other hand, ShuffleNet is a well-optimized baseline which obtains good accuracy performance on a very tight parameter and fpmul budget. In addition, we did not do much hyperparamter tuning to any of the networks. Still, HadaNet was able to obtain a small improvement over ShuffleNet both with and without our architecture changes.

4.4. Practicality of HadaNet

HadaNet slightly outperforms ShuffleNet on accuracy, but it also requires extra additions. Section 3.4 gives formulae for the overhead — an N-channel group conv with B groups requires N^2/B fpmuls for the weight layer and $2N \log B$ adds for the two block-Hadamard transforms. Compared to multiplies, additions are already much cheaper in practical hardware. The last column of Table 4 shows the number of additions needed for each network if the fast Hadamard transform is used. This data shows that the relative overhead of HadaNet is fairly small: 1.6% of the existing multiply-accumulates in ResNet-18, and 5.8% in ShuffleNet-x0.25.

That said, the overhead of the Hadamard transform being small depends on a well-optimized implementation. On GPU, this means a fast Hadamard kernel operating along the channels dimension. This is not currently available and our own HadaNet implementation is slow as a result.

On the other hand, we believe the Hadamard transform might be useful for specialized DNN accelerators implemented with FPGAs [3] or ASICs [11]. Top computer hardware conferences already contain works demonstrating the use of structured matrices for DNN compression in dedicated hardware [6, 22, 4]. This is because transform kernels can be very efficiently implemented on a dedicated chip. Our study reveals that there is a high weight sparsity regime at which dense transforms outperform simple shuffling. In this regime HadaNet is more efficient than existing state-of-the-art (i.e. DFT transforms) while achieving similar accuracy performance.

5. Conclusions and Future Work

We introduce the concept of unitary group convolutions, a composition of group convolutions with unitary transforms in feature space. We use the UGConv framework to unify two disparate ideas in CNN literature, ShuffleNets and block-circulant networks, and provide valuable insights into both techniques. UGConvs with dense unitary transforms demonstrate superior ability to learn cross-channel mappings versus ordinary and shuffled group convolutions. Based on these these observations we propose HadaNet, a variant of ShuffleNet that improves accuracy on the ImageNet dataset without incurring additional parameters or floating-point multiplies.

One future work is to replace the Hadamard transform with a trained $0, +1, -1$ transform; training may allow the transform to adapt to the weights, and introducing zeros enables sparse compute reduction.

References

[1] Y. Cheng, F. X. Yu, R. S. Feris, S. Kumar, A. Choudhary, and S.-F. Chang. An Exploration of Parameter Redundancy in Deep Networks with Circulant Projections. *Int'l Conf. on Computer Vision (ICCV)*, pages 2857–2865, 2015. 1, 2

[2] F. Chollet. Xception: Deep learning with Depthwise Separable Convolutions. *Conf. on Computer Vision and Pattern Recognition (CVPR)*, Jun 2016. 1, 2

[3] E. Chung, J. Fowers, K. Ovtcharov, M. Papamichael, A. Caulfield, T. Massengill, M. Liu, D. Lo, S. Alkalay, M. Haselman, M. Abeydeera, L. Adams, H. Angepat, C. Boehn, D. Chiou, O. Firestein, A. Forin, K. S. Gatlin, M. Ghandi, S. Heil, K. Holohan, A. E. Husseini, T. Juhasz, K. Kagi, R. K. Kovvuri, S. Lanka, F. van Megen, D. Mukhortov, P. Patel, B. Perez, A. G. Rapsang, S. K. Reinhardt, B. D. Rouhani, A. Sapek, R. Seera, S. Shekar, B. Sridharan, G. Weisz, L. Woods, P. Y. Xiao, D. Zhang, R. Zhao, , and D. Burger. Serving DNNs in Real Time at Datacenter Scale with Project Brainwave . *IEEE Micro*, 38(2):8–20, 2018. 8

[4] C. Deng, S. Liao, Y. Xie, K. K. Parhi, X. Qian, and B. Yuan. PermDNN: Efficient Compressed Deep Neural Network Architecture with Permuted Diagonal Matrices . *Int'l Symp. on Microarchitecture (MICRO)*, Oct 2019. 2, 8

[5] T. C. Deveci, S. Cakir, and A. E. Cetin. Energy Efficient Hadamard Neural Networks. *arXiv preprint*, arXiv:1805.05421, May 2018. 2

[6] C. Ding, S. Liao, Y. Wang, Z. Li, N. Liu, Y. Zhuo, C. Wang, X. Qian, Y. Bai, G. Yuan, X. Ma, Y. Zhang, J. Tang, Q. Qiu, X. Lin, and B. Yuan. CirCNN: Accelerating and Compressing Deep Neural Networks using Block-Circulant Weight Matrices. *Int'l Symp. on Microarchitecture (MICRO)*, pages 395–408, 2017. 1, 2, 3, 5, 8

[7] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. *arXiv e-print*, arXiv:1512.0338, Dec 2015. 6, 7

[8] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient Convolutional Neural Networks for Nobile Vision Applications. *arXiv e-print*, arXiv:1704.04861, 2017. 1, 2

[9] G. Huang, Y. Sun, Z. Liu, D. Sedra, and K. Q. Weinberger. Deep Networks with Stochastic Depth. *European Conference on Computer Vision (ECCV)*, pages 646–661, 2016. 7

[10] Y. Ioannou, D. Robertson, R. Cipolla, and A. Criminisi. Deep roots: Improving cnn efficiency with hierarchical filter groups. *Conf. on Computer Vision and Pattern Recognition (CVPR)*, Jun 2017. 2

[11] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, et al. In-Datacenter Performance Analysis of a Tensor Processing Unit. *Int'l Symp. on Computer Architecture (ISCA)*, pages 1–12, 2017. 8

[12] A. Krizhevsky and G. Hinton. Learning Multiple Layers of Features from Tiny Images. *Tech report*, 2009. 6

[13] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems (NIPS)*, pages 1097–1105, 2012. 1, 2

[14] Q. Le, T. Sarlós, and A. Smola. Fastfood —Approximating Kernel Expansions in Loglinear Time. *Int'l Conf. on Machine Learning (ICML)*, 85, 2013. 2

[15] M. Lin, Q. Chen, and S. Yan. Network in Network. *arXiv e-print*, arXiv:1312.4400, 2013. 7

[16] M. Moczulski, M. Denil, J. Appleyard, and N. de Freitas. ACDC: A Structured Efficient Linear Layer. *Int'l Conf. on Learning Representations (ICLR)*, 2016. 2

[17] W. K. Pratt, J. Kane, and H. C. Andrews. Hadamard Transform Image Coding. *Proceedings of the IEEE*, 57(1):58–68, 1969. 5

[18] L. Sifre. Rigid-Motion Scattering for Image Classification. *Ph.D. thesis*, 2014. 2

[19] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv e-print*, arXiv:1409.15568, Apr 2015. 1

[20] V. Sindhwani, T. Sainath, and S. Kumar. Structured Transforms for Small-Footprint Deep Learning. *Advances in Neural Information Processing Systems (NIPS)*, pages 3088–3096, 2015. 1, 2

[21] K. Sun, M. Li, D. Liu, and J. Wang. Igcv3: Interleaved low-rank group convolutions for efficient deep neural networks. *arXiv e-print*, arXiv:1806.00178, Jun 2018. 2

[22] S. Wang, Z. Li, C. Ding, B. Yuan, Q. Qiu, Y. Wang, and Y. Liang. C-LSTM: Enabling Efficient LSTM using Structured Compression Techniques on FPGAs. *to appear in International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2018. 1, 2, 8

[23] Y. Wang, C. Ding, Z. Li, G. Yuan, S. Liao, X. Ma, B. Yuan, X. Qian, J. Tang, Q. Qiu, and X. Lin. Towards Ultra-High Performance and Energy Efficiency of Deep Learning Systems: An Algorithm-Hardware Co-Optimization Framework. *AAAI Conf' on Artificial Intelligence (AAAI)*, Feb 2018. 2, 3

[24] G. Xie, J. Wang, T. Zhang, J. Lai, R. Hong, and G.-J. Qi. Igcv2: Interleaved structured sparse convolutional neural networks. *arXiv e-print*, arXiv:1804.06202, Apr 2018. 2

[25] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated Residual Transformations for Deep Neural Networks. *Conf. on Computer Vision and Pattern Recognition (CVPR)*, Jun 2017. 2

[26] X. Xu, Y. Ding, S. X. Hu, M. Niemier, J. Cong, Y. Hu, and Y. Shi. Scaling for Edge Inference of Deep Neural Networks. *Nature Electronics*, 1(4):216, 2018. 1

[27] Z. Yang, M. Moczulski, M. Denil, N. de Freitas, A. Smola, L. Song, and Z. Wang. Deep Fried Convnets. *Int'l Conf. on Computer Vision (ICCV)*, pages 1476–1483, 2015. 2

[28] T. Zhang, G.-J. Qi, B. Xiao, and J. Wang. Interleaved group convolutions. *Conf. on Computer Vision and Pattern Recognition (CVPR)*, Jun 2017. 2

[29] X. Zhang, X. Zhou, M. Lin, and J. Sun. ShuffleNet: An Extremely Efficient Convolutional Neural Network for Mobile Devices. *arXiv e-print*, arXiv:1707.01083, Aug 2017. 1, 2, 3, 5, 6, 7, 8