

LeanConvNets: Low-cost Yet Effective Convolutional Neural Networks

Jonathan Ephrath, Moshe Eliasof, Lars Ruthotto, Eldad Haber and Eran Treister

Abstract—Convolutional Neural Networks (CNNs) have become indispensable for solving machine learning tasks in speech recognition, computer vision, and other areas that involve high-dimensional data. A CNN filters the input feature using a network containing spatial convolution operators with compactly supported stencils. In practice, the input data and the hidden features consist of a large number of channels, which in most CNNs are fully coupled by the convolution operators. This coupling leads to immense computational cost in the training and prediction phase. In this paper, we introduce LeanConvNets that are derived by sparsifying fully-coupled operators in existing CNNs. Our goal is to improve the efficiency of CNNs by reducing the number of weights, floating point operations and latency times, with minimal loss of accuracy. Our lean convolution operators involve tuning parameters that controls the trade-off between the network’s accuracy and computational costs. These convolutions can be used in a wide range of existing networks, and we exemplify their use in residual networks (ResNets). Using a range of benchmark problems from image classification and semantic segmentation, we demonstrate that the resulting LeanConvNet’s accuracy is close to state-of-the-art networks while being computationally less expensive. In our tests, the lean versions of ResNet in most cases outperform comparable reduced architectures such as MobileNets and ShuffleNets.

I. INTRODUCTION

CONVOLUTIONAL neural networks (CNNs) [1] are among the most effective machine learning approaches for processing structured, high-dimensional data such as voice recordings, images, and videos and have become indispensable in, e.g., speech recognition [2], [3], audio processing [4], and image classification [5].

In the forward propagation, a CNN filters the input features through a sequence of layers, which are composed of convolution operators, biases, normalization layers, nonlinear activation functions, and pooling operators. In imaging tasks, the input features and the hidden features at each layer can be grouped into several channels, each of which can be interpreted as an image. The stencils that parameterize the convolution operators are typically chosen to have a small support around the origin. Hence, each feature in an image interacts with features from a small neighborhood in its channel and, in the standard, *fully-coupled* approach, the features

J. Ephrath, M. Eliasof, and E. Treister are with the Department of Computer Sciences at the Ben-Gurion University of the Negev, Be’er Sheva, Israel. emails: [ephrathj,eliasof]@post.bgu.ac.il, erant@cs.bgu.ac.il. JE and ME contributed equally to this work.

L. Ruthotto is with the Departments of Mathematics and Computer Science, Emory University, Atlanta, GA, USA. email: lruthotto@emory.edu.

E. Haber is with the Department of Earth, Ocean and Atmospheric Sciences, University of British Columbia, and with Xtract AI, Vancouver, Canada. email: ehaber@eos.ubc.ca.

from the same neighborhood in the remaining channels [6], [7]. A drawback of the fully-coupled approach is that the number of convolution operators in a layer is proportional to the product of the number of input and output channels. This scaling can render wide architectures (i.e., architectures whose layers contain a large number of channels) prohibitively expensive in training and inference. It also complicates the deployment of such CNNs, especially on devices with limited memory and computing resources like autonomous vehicles, drones, and smartphones.

In recent years there has been an effort to improve the efficiency of CNNs. Common approaches to reduce the number of weights in CNNs are pruning [8]–[13], sparsity [14]–[16], and quantization [17]–[19]. Pruning reduces the number of weights in the network after training. The fact that in many cases large portions of the networks’ weights can be removed with minimal reduction of its accuracy indicates a considerable redundancy and over-parameterization of standard CNNs [20]. While pruning is effective in reducing the number of weights and floating point operations (FLOPs), it generally leads to an unstructured non-zero pattern of the weights, which increases the memory access costs. The lack of structure also complicates the efficient deployment of the CNN on hardware.

Another approach to improve the efficiency of CNNs is to replace the fully-coupled convolution operators by sparser convolution operators (i.e., operators with fewer non-zero elements) before training. One typical building block is known as a *grouped* convolution operator, which partitions the channels into groups and only allow grouped coupling; see, e.g., [5]. When the number of groups equals the number of channels, one obtains a *depth-wise* convolution operator, which is a block diagonal matrix whose blocks are spatial convolution operators. The depth-wise convolution operator filters each channel of the image data separately and thus restricts the interaction of each feature to its nearby features in the same channel. It is common to use the depth-wise operator in conjunction with fully-connected point-wise 1×1 convolutions to introduce coupling across the channels.

A few CNN architectures have been derived using depth-wise and 1×1 convolution operators, often augmented with bottleneck or shuffling techniques; see, e.g., [21]–[25]. These works use the depth-wise and 1×1 convolution separately, with activation and batch normalization layers in between them. Although in this paper we use 3×3 convolutional stencils to parameterize the depth-wise convolution, other choices are possible; in fact, mixing stencils of different sizes has shown promising results [26]. This typically requires a redesign of existing CNN architectures. To reduce the ratio between FLOPs

and memory access in depth-wise convolution operators, replacing convolutions with shifts has been proposed in [27]. Exploiting the multiscale structure of image data provides an alternative way to derive more efficient architectures; see the contemporary work [28]. It is known, however, that the memory access is often the real bottleneck in modern parallel hardware, and *not necessarily FLOPs*. In fact, state-of-the-art implementations of depth-wise convolution operators on GPUs involve more FLOPs than necessary to achieve lower runtimes; see, e.g., [29]. Nevertheless, whether the dominant cost is the storage of the parameters, the FLOPs, or the memory access, is highly dependent on the hardware at hand. Therefore, it is desirable to build convolution operators and architectures that are flexible in their definition so that they can be configured as necessary on any specific hardware.

In this paper, we introduce LeanConvNets, a new family of CNNs built as lean versions of known networks, using lean convolution operators. These operators reduce the number of weights, computation time, and FLOPs while achieving competitive results. The lean operators preserve the overall network structure and can thus be applied to a variety of networks, e.g., residual networks (ResNets, ResNeXt) [30]–[32], which have been two of the most reliable architectures in the literature. The following aspects set our work apart from other approaches:

- We obtain a new operator as the sum of the grouped and 1×1 convolution operators (a schematic description of a ResNet block with lean convolutions appears in Fig. 1; more details later). Using a prototype implementation, we show that handling both operations simultaneously reduces the computation time required to apply the operator. Also, this design introduces several opportunities for optimization in hardware through its parallelism, minimal number of memory accesses, and slightly reduced number of weights. Using grouped instead of depth-wise convolution operators allows one to gradually enlarge the portion of spatial convolutions in order to improve the performance of the lean networks. Our networks are mostly suitable to parallel devices that are bandwidth bounded and not computation bounded.

- We present two ways to reduce the spatial kernel size that further decrease the number of weights and FLOPs and are easy to implement efficiently. In the first method, we replace the standard 3×3 by a 5-point stencil. In the second method, we filter two-dimensional images using a one-dimensional convolution operator (3×1 or 1×3 , depending on memory layout) and its transpose applied at the memory write. This operator can be implemented with the same number of memory accesses as the 1×1 convolution since the memory of the feature maps is sequential in the one dimension.

The remainder of the paper is organized as follows: In Sec. II, we discuss existing convolution operators and their computational costs in the context of residual neural networks. In Sec. III, we introduce a family of lean convolution operators, analyze their costs, and outline their implementation. In Sec. IV, we provide extensive numerical evidence for the efficacy of the resulting LeanConvNets for image classification and semantic segmentation. In Sec. V, we summarize the paper and discuss directions for future research.

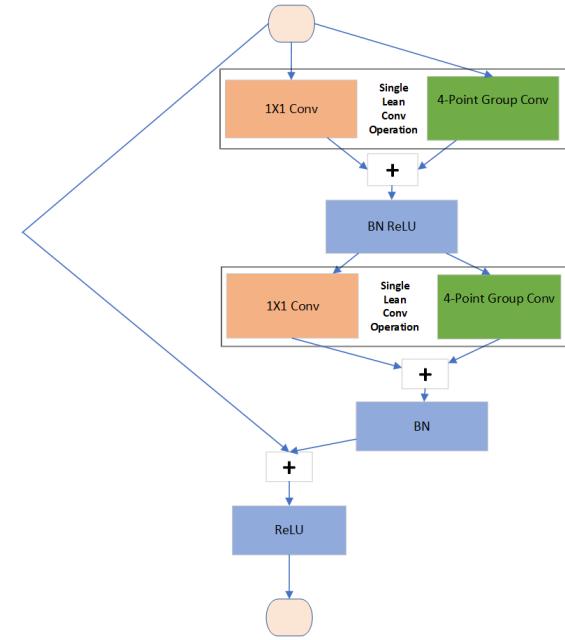


Fig. 1: Building blocks of basic LeanResNet step. The 1×1 and spatial grouped convolutions are applied simultaneously.

II. PRELIMINARIES AND NOTATION

We now introduce our main notation and define the supervised classification and semantic segmentation problems that we use to validate our methods; for more details see [7]. For brevity, we restrict the discussion to images although the techniques derived here can also be used for other structured data types such as audio or video data. In supervised learning, we are given a set of training data consisting of pairs, $\{(\mathbf{y}_0^{(k)}, \mathbf{c}^{(k)})\}_{k=1}^s \subset \mathbb{R}^{n_f} \times \mathbb{R}^{n_c}$. In our case, $\mathbf{y}_0^{(k)}$ is the k -th input image and $\mathbf{c}^{(k)}$ either represents the probabilities for the entire image (in classification) or each pixel (in segmentation) to belong to one of the pre-defined classes. Our goal is to define a neural network architecture and train its weights $\boldsymbol{\theta} \in \mathbb{R}^p$ and the weights of a linear classifier, denoted by $\mathbf{W} \in \mathbb{R}^{n_c \times n_{out}}$ and $\boldsymbol{\mu} \in \mathbb{R}^{n_c}$, such that

$$\mathbf{c}^{(k)} \approx S(\mathbf{W}\mathbf{y}^{(k)}(\boldsymbol{\theta}) + \boldsymbol{\mu}), \quad \text{for all } k = 1, 2, \dots, s.$$

Here, S is a softmax hypothesis function and $\mathbf{y}^{(k)}(\boldsymbol{\theta}) \in \mathbb{R}^{n_{out}}$ denotes the output of the network applied to the k th sample.

The learning problem can be phrased as a minimization problem of a regularized empirical loss function

$$\min_{\boldsymbol{\theta}, \mathbf{W}, \boldsymbol{\mu}} \frac{1}{s} \sum_{k=1}^s L(S(\mathbf{W}\mathbf{y}^{(k)}(\boldsymbol{\theta}) + \boldsymbol{\mu}), \mathbf{c}^{(k)}) + R(\boldsymbol{\theta}, \mathbf{W}, \boldsymbol{\mu}),$$

where L is the cross entropy loss and R is a regularization function. The optimization problem is usually solved using variants of stochastic gradient descent (SGD); see the original work [33] and the survey [34].

As a baseline architecture, we consider residual networks (ResNet) [30], [31], which have been successful in many imaging tasks. Given a data sample, \mathbf{y}_0 , the forward propagation through an N -layer ResNet is defined as

$$\mathbf{y}_{l+1} = \mathbf{y}_l + \mathbf{F}(\boldsymbol{\theta}_l, \mathbf{y}_l), \quad \text{for } l = 0, \dots, N-1, \quad (1)$$

where θ_l is the set of weights associated with the l -th layer and we define $\mathbf{y}(\theta) = \mathbf{y}_0$. There are different choices for the nonlinear term in (1), e.g.,

$$\mathbf{F}(\theta_l, \mathbf{y}_l) = \mathbf{K}_2(\theta_{l,2})\sigma(\mathcal{N}(\mathbf{K}_1(\theta_{l,1})\sigma(\mathcal{N}(\mathbf{y}_l)))). \quad (2)$$

Here, $\sigma(x) = \max\{x, 0\}$ denotes an element-wise rectified linear unit (ReLU) activation function and the weights are partitioned into $\theta_{l,1}$ and $\theta_{l,2}$ that parameterize the two linear operators \mathbf{K}_1 and \mathbf{K}_2 , respectively. For brevity, we omit the weights of the normalization layer \mathcal{N} .

In convolutional ResNets, the operators \mathbf{K}_i in (2) are composed of spatial convolution operators. If the input \mathbf{y} has c_{in} channels, and the output $\mathbf{K}\mathbf{y}$ has c_{out} channels, then the common choice for an operator \mathbf{K} is a $c_{\text{out}} \times c_{\text{in}}$ block matrix of convolutions, introducing full coupling across the channels. For example, if $c_{\text{in}} = c_{\text{out}} = 4$, then the convolution operators in (2) can be written in matrix form as

$$\mathbf{K}_{\text{full}}(\theta) = \begin{pmatrix} \mathbf{C}^{(1,1)} & \mathbf{C}^{(1,2)} & \mathbf{C}^{(1,3)} & \mathbf{C}^{(1,4)} \\ \mathbf{C}^{(2,1)} & \mathbf{C}^{(2,2)} & \mathbf{C}^{(2,3)} & \mathbf{C}^{(2,4)} \\ \mathbf{C}^{(3,1)} & \mathbf{C}^{(3,2)} & \mathbf{C}^{(3,3)} & \mathbf{C}^{(3,4)} \\ \mathbf{C}^{(4,1)} & \mathbf{C}^{(4,2)} & \mathbf{C}^{(4,3)} & \mathbf{C}^{(4,4)} \end{pmatrix}, \quad (3)$$

where $\mathbf{C}^{(i,j)} = \mathbf{C}(\theta^{(i,j)})$ denotes the sparse matrix associated with the spatial convolution kernel parameterized by the 3×3 filter $\theta^{(i,j)} \in \mathbb{R}^9$. For ease of notation, we do not explicitly denote the dependency on θ in the following. The sparsity pattern of this operator is visualized in the leftmost subplot of Fig. 2 for an image size of 6×6 . Applying \mathbf{K}_{full} requires $\mathcal{O}(c_{\text{in}} \cdot c_{\text{out}})$ FLOPs, and \mathbf{K}_{full} has $9 \cdot c_{\text{in}} \cdot c_{\text{out}}$ weights. In practice, each \mathbf{K}_{full} can have millions of weights.

Grouped convolutions are popular alternatives to \mathbf{K}_{full} as they reduce the number of weights and computations. In our example, we can restrict the interaction of the channels to $g = 2$ groups, which leads to the block diagonal matrix

$$\mathbf{K}_{g=2} = \begin{pmatrix} \mathbf{C}^{(1,1)} & \mathbf{C}^{(1,2)} & 0 & 0 \\ \mathbf{C}^{(2,1)} & \mathbf{C}^{(2,2)} & 0 & 0 \\ 0 & 0 & \mathbf{C}^{(3,3)} & \mathbf{C}^{(3,4)} \\ 0 & 0 & \mathbf{C}^{(4,3)} & \mathbf{C}^{(4,4)} \end{pmatrix}. \quad (4)$$

This reduces the number of weights and FLOPs by a factor of g compared to the full convolution. Clearly, $\mathbf{K}_{\text{full}} = \mathbf{K}_{g=1}$ and for $g = c_{\text{in}}$, we get the depth-wise convolution; the sparsity pattern of $\mathbf{K}_{g=2}$ and $\mathbf{K}_{g=4}$ are shown in Fig. 2. Grouped convolutions can be extended to rectangular operators when g divides both c_{in} and c_{out} .

III. LEAN CONVOLUTIONAL OPERATORS

We now introduce a family of lean convolutional operators that achieve competitive performance and reduce the number of weights, memory access, and FLOPs. It has been shown that 1×1 convolutions can be effective if complemented with a relatively small number of spatial convolutions [21]–[25]. Then, the computational cost of the 1×1 convolution (in terms of FLOPs and weights) dominates the cost of the spatial convolutions as the number of channels grows. It has also been observed that the accuracy of the network suffers from the relative shortage of spatial convolutions, which is

often explained by a relatively small number of weights. To increase the accuracy, our lean convolution operators aim to allocate the weights more efficiently between grouped and 1×1 convolutions. To this end, we reduce the kernel size on the one hand and add spatial convolutions on the other. The group size of the spatial convolution is a hyper parameter that trades off between accuracy and computational efficiency. This also allows us to accommodate different computational devices without changing the high-level structure of the network.

We obtain lean convolution operators in three steps. First, lean operators are a sum of 1×1 and grouped spatial convolutions—see Fig. 1. If implemented efficiently the lean convolution allows one to reuse of memory access, increase parallelism, and further reduce the number of weights. Second, the group size parameter g allows the user to balance between spatial filtering (by a grouped operator) and coupling (by 1×1 operators) and control the performance of the network. Third, we use convolutional filters with only five or three elements instead of 9 for common 3×3 filters, which reduces the number of weights and FLOPs.

Continuing our example from above, we define the lean analogue to (3)–(4) as

$$\mathbf{K}_{\text{lean},g=2} = \begin{pmatrix} \mathbf{C}^{(1,1)} & \mathbf{C}^{(1,2)} & \alpha_{1,3}\mathbf{I} & \alpha_{1,4}\mathbf{I} \\ \mathbf{C}^{(2,1)} & \mathbf{C}^{(2,2)} & \alpha_{2,3}\mathbf{I} & \alpha_{2,4}\mathbf{I} \\ \alpha_{3,1}\mathbf{I} & \alpha_{3,2}\mathbf{I} & \mathbf{C}^{(3,3)} & \mathbf{C}^{(3,4)} \\ \alpha_{4,1}\mathbf{I} & \alpha_{4,2}\mathbf{I} & \mathbf{C}^{(4,3)} & \mathbf{C}^{(4,4)} \end{pmatrix}, \quad (5)$$

where \mathbf{I} is a scaled identity matrix and $\alpha_{i,j} \in \mathbb{R}$ are weights. The identity operators represent the 1×1 convolution and the convolution operators $\mathbf{C}^{(i,j)}$ enable spatial filtering. Similarly, the lean operator with $g = 4$ groups is

$$\mathbf{K}_{\text{lean},g=4} = \begin{pmatrix} \mathbf{C}^{(1,1)} & \alpha_{1,2}\mathbf{I} & \alpha_{1,3}\mathbf{I} & \alpha_{1,4}\mathbf{I} \\ \alpha_{2,1}\mathbf{I} & \mathbf{C}^{(2,2)} & \alpha_{2,3}\mathbf{I} & \alpha_{2,4}\mathbf{I} \\ \alpha_{3,1}\mathbf{I} & \alpha_{3,2}\mathbf{I} & \mathbf{C}^{(3,3)} & \alpha_{3,4}\mathbf{I} \\ \alpha_{4,1}\mathbf{I} & \alpha_{4,2}\mathbf{I} & \alpha_{4,3}\mathbf{I} & \mathbf{C}^{(4,4)} \end{pmatrix}, \quad (6)$$

The sparsity patterns of these operators are shown in fourth and fifth subplots of Fig. 2.

The setting in (6) with $g = 4$ can be seen as the sum of depth-wise and 1×1 convolutions, which are also used in [21], [22], [25]. These works perform the depth-wise and 1×1 convolutions separately—the depth-wise convolution is applied between two 1×1 convolutions with ReLU operations between them. Since we sum the operators we can apply them both simultaneously, which allows us to optimize memory access and improve parallelism (more can be done at once).

A. The argument for groups in compact networks

Our experiments suggest that if we take a given compact network that utilizes depth-wise and 1×1 operations, and define its *full* version by placing a 3×3 convolution instead of each 1×1 convolution, we get networks that are significantly more expensive, but perform better in terms of accuracy. Employing such compact schemes may result in a spatial component that is too small, especially when the number of channels is large and the 1×1 operators dominate the spatial convolutions. This motivates us to add a small number (e.g.,

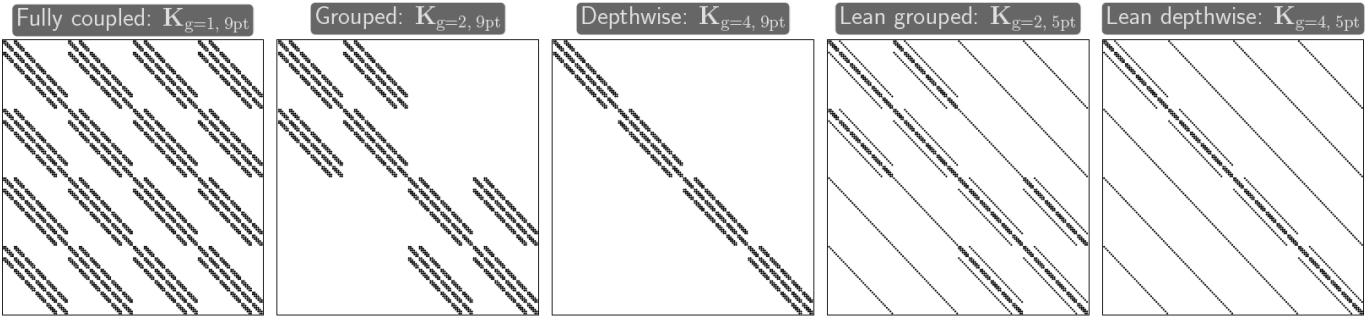


Fig. 2: Sparsity patterns of different convolution operators for 6×6 images with four input and output channels. The leftmost subplot shows the sparsity pattern of a 3×3 fully-coupled convolution operator. The next two subplots depict the grouped convolution operators for $g = 2$ and $g = 4$, respectively. The remaining two subplots show the proposed lean grouped and depth-wise operators that are a sum of a fully coupled 1×1 and a grouped (or depth-wise) spatial convolution operators.

$\frac{c_{in}c_{out}}{g}$) of spatial convolutions to improve the performance of the network compared to depth-wise operators.

The motivation for using this operator is as follows: first, the implementation of the grouped convolution works best in groups of intermediate size, and it is often even more efficient to zero-pad the groups (artificially enlarge them) to get better computational performance on GPUs [29]. Second, in the standard combination of depth-wise and 1×1 convolutions, the former becomes negligible compared to the latter as the number of channels grow, hurting accuracy without providing considerable savings. Our proposal in this context is to use the grouping mechanism to keep a constant ratio of operations between the two types of convolutions, such that the 1×1 convolution that has $c_{in} \cdot c_{out}$ weights is more dominant than the grouped convolution that has $\frac{(r-1)c_{in}c_{out}}{g}$ weights¹, where r is the stencil size (e.g., $r = 9$ for a 3×3 stencil). For example, if we choose a ratio of $\frac{1}{g}$, then we set $g \approx 8(r-1)$ s.t. the number of channels is divisible by g . We subtract 1 from r since the middle weight is included in the 1×1 convolution.

We note that enhancing the lean convolution with the grouping mechanism is similar to enhancing the depth-wise convolution in networks such as MobileNetV2. This would result in a network that is similar to the ResNeXt networks [32], which applies a grouped convolution instead of the full 3×3 convolution in the bottleneck version of the original ResNet. The grouping helps to enlarge the bottleneck expansion while keeping a low additional cost, and without adding many weights. Although the works were proposed independently, maximizing the number of groups in ResNeXt leads to a network which is similar to MobileNetV2.

B. LeanConv 5-pt: lean convolutions with 5-point stencils.

The first version of LeanConvNets is based on 5-point convolution stencils. The idea is to replace the stencils of $C^{(i,j)}$ in (5) and (6) by the 5-point stencil

$$\begin{bmatrix} 0 & c_{i,1} & 0 \\ c_{i,2} & \alpha_{i,i} & c_{i,3} \\ 0 & c_{i,4} & 0 \end{bmatrix}, \quad (7)$$

¹The cost in operations is proportional to the number of weights.

where $\alpha_{i,i}$ is the i, i -th entry of the 1×1 convolution, and $c_{i,1}, \dots, c_{i,4}$ are additional four weights per input channel i . An example for the sparsity pattern of the resulting lean operator with 5-point convolutions is shown in Fig. 2. The operator $K_{\text{lean},g}$ with 5-point stencils and g groups has $(1 + \frac{4}{g})(c_{in} \cdot c_{out})$ weights. We note that if the number of channels and g are large, then the 1×1 convolution is the dominating operator both in terms of weights and FLOPs.

The lean convolution can replace fully-coupled convolution operators in many existing CNNs without any structural changes to the architecture. A straightforward way to implement a grouped lean convolution like in (5) is to use the package cudnn to perform the 1×1 and spatial convolutions separately. As we show later, our custom implementation, which simultaneously applies both operations, outperforms the cudnn approach for $g = c_{in}$.

To motivate the use of 5-point stencils, consider the ResNet architectures that have been recently interpreted as time-dependent nonlinear ordinary differential equations (ODEs); see, e.g., [35]–[41]. This allows the community to analyze and extend ResNets using theoretical and practical ideas from the world of ODEs and PDEs [42]. In this point of view, the five-point stencil is able to express a mass term, and a discretization of first and second spatial derivatives in the x and y dimensions. That is, the first and second derivatives in the x dimension can be approximated by

$$\frac{\partial}{\partial x} \approx \frac{1}{2h_x}[-1, 0, 1] \quad \text{and} \quad \frac{\partial^2}{\partial x^2} \approx \frac{1}{h_x^2}[1, -2, 1], \quad (8)$$

where h_x is the edge length of a pixel. This, together with $\frac{\partial}{\partial y}$, $\frac{\partial^2}{\partial y^2}$ and the mass term (or 5-point low-pass filter) are included in the span of the 5-point stencil (7). The remaining four entries of a full 3×3 stencil correspond to mixed partial derivatives, which rarely occur in PDE-models, and are thus good candidates for reducing computations and weights.

Implementation: The standard 3×3 convolution is implemented using a shift per stencil parameter (known as the `shiftIm2col` operation), and a matrix-matrix multiplication using the function `gemm`. In the same way we can multiply a 5-point convolution operator, trivially saving $4/9$ of the operations. For small group sizes, and in particular for the

depth-wise setting (group size of 1), we found that a direct implementation of the convolutions is faster than the standard implementation with `shiftIm2col`. As the groups get larger, the approach using `shiftIm2col` is more preferable, due to the efficiency of `gemm`. Even more savings can be realized in 3D CNNs where the standard 27-point stencils are replaced with 7-point stencils (the 3D version of (7)), saving 20/27 of the operations and weights.

C. LeanConv 3-pt: lean convolutions with 1D 3-point stencils

In this section, we present a more sophisticated lean convolution operator that can be applied almost at the same cost as a 1×1 convolution as both operators use the same memory accesses. This convolution is based on *1D* convolution operators, either 1×3 or 3×1 , which can be applied efficiently if the memory is continuous in the direction of the 1D kernel.

In addition to the benefits from an implementation perspective, the use of 1D kernels can also be motivated as follows: It is known that in 2D, a large portion of the 3×3 kernel can be parameterized by a multiplication of 1×3 and 3×1 kernels, also called *separable* kernels. Separable kernels can represent many of the important operators, such as low-pass filters, and the spatial derivatives in (8). Our idea is to use two convolutions, such as \mathbf{K}_1 and \mathbf{K}_2 in (2), to effectively apply separable operators: \mathbf{K}_1 applies a 1×3 kernel in the horizontal direction, and \mathbf{K}_2 applies a 3×1 kernel in the vertical direction. We note that 1D stencils were also used in a small section of the InceptionV4 network [43]. There, 1×7 and 7×1 were used in together with 3×3 convolutions to increase the field of view of the network. Here we show that even if we use 1×3 and 3×1 convolutions *only*, we can still get an effective network, while reducing the number of weights, FLOPs and (most importantly) memory access. The latter can be saved if the memory of the feature maps is aligned with the direction of the kernel. Our idea here to maintain the memory alignment is to apply the convolutions together with channel transposition. That is, if the 1D convolution operator \mathbf{K}_1 is aligned with the memory, then the feature maps are transposed during the WRITE operation to prepare the result to \mathbf{K}_2 that is aligned in the other direction, and vice-versa.

Custom GPU Implementation: To explain the 3-pt lean convolution, we first briefly describe one of the approaches for multiplying matrices on a GPU—that is essentially the 1×1 convolution operator. We follow the description of the `cutlass` library [44], and the implementation of the MAGMA open source project [45]. Given two matrices $\mathbf{K} \in \mathbb{R}^{c_{\text{out}} \times c_{\text{in}}}$ and $\mathbf{Y} \in \mathbb{R}^{c_{\text{in}} \times n}$, we first divide their product $\mathbf{KY} \in \mathbb{R}^{c_{\text{out}} \times n}$ into tiles of size $t_n \times t_o$. Each of these tiles is computed by a multiplication of a block of t_o columns of \mathbf{K} and t_n columns of \mathbf{Y} . These sub-matrices are also divided into sub-blocks of size t_i . Each group of physical cores gets a task of computing a tile of $t_n \times t_o$ output numbers. To apply this, we first fetch the relevant tiles into shared memory, and then multiply them in parallel. Algorithm 1 summarizes the procedure, ignoring the underlined parts; see [44], [45] for more details.

Now we explain how to apply the 3-pt lean operator, assuming that the number of convolutions is small compared to $c_{\text{in}} \times c_{\text{out}}$. An important consideration for GPU implementations is that fetching memory from global memory into shared memory is slow, while accessing the shared memory is fast. Our idea is to add a small memory fetch into the procedure above, and apply the 1D convolution to the already-fetched tile of \mathbf{Y} , assuming that the memory of \mathbf{Y} is continuous in the same direction of the 1D kernel.

To finish the operation, we now wish that the direction of the next kernel is aligned with the direction of the data. This will be true only for one direction, and we handle that by transposing the feature maps during the *write* phase at the end of the convolution. Thus, when multiplying \mathbf{K}_1 in (2) we have the maps aligned in one direction, but during the multiplication, we transpose the data in shared memory and write it transposed. After a ReLU operation (for which the direction does not matter) the input to \mathbf{K}_2 is ready to be multiplied and is aligned in the other direction. At the end of the same multiplication of \mathbf{K}_2 , the result is again transposed back to the original alignment. Algorithm 1 summarizes the 3-pt lean convolution procedure. Compared to the 5-pt lean convolution, the 3-pt conv requires that at least two kernels are applied one after the other before a skip connection, such that the maps are transposed back to their original form. We note that this algorithm is mostly beneficial with large number of groups, and in particular with depth-wise configuration. If the groups are large, the `shiftIm2Col` approach is preferable.

Algorithm 1 Tiled 3-point LeanConv Multiplication

```

1: # Computation of  $\mathbf{X} = \mathbf{KY}$ .
2: #  $t_n, t_o, t_i$ : tile sizes.  $(i, j)$ : thread id.
3: # The underline parts add over simple  $1 \times 1$  conv.
4: procedure LEAN3PTGEMM( $\mathbf{X}, \mathbf{K}, \mathbf{Y}, i, j$ )
5:   Fetch boundary values from  $\mathbf{Y}$  and spatial
6:   convolution parameters into shared memory  $C$ .
7:   for  $k = 1, \dots, \lceil \frac{c_{\text{in}}}{t} \rceil$  do
8:     # Each thread in the tile fetches two blocks:
9:     Fetch tile  $i, k$  from  $\mathbf{K}$  to shared  $A \in \mathbb{R}^{t_n \times t_i}$ .
10:    Fetch tile  $k, j$  from  $\mathbf{Y}$  to shared  $B \in \mathbb{R}^{t_i \times t_o}$ .
11:    # Multiply  $AB$  by  $t_i$  outer products:
12:    Multiply  $AB$  into local memory.
13:    If relevant to the output tile:
14:      Apply convolution to  $B$  into local memory.
15:      Write transposed local memory to  $\mathbf{X}$ .
16:    end for
17: end procedure

```

IV. EXPERIMENTS

We demonstrate the proposed LeanConvNet approach and compare the lean versions of ResNet and ResNeXt [32], called "LeanResNet", and "LeanResNeXt" to a fully-coupled ResNet, and other recent state-of-the-art compact architectures: ShuffleNetV2 [25], MobileNetV2 [22], and ShiftResNet [27]. We consider the image classification and semantic segmentation tasks using several data sets. Our primary focus is to

Type	Layer width (channels)	# Steps	Strides
Res18	32-64-128-256	2-2-2-2	1-2-2-2
Res24-narrow	12-24-48-96	2-3-3-3	1-2-2-2
Res24	32-64-128-256	2-3-3-3	1-2-2-2
Res34	64-128-256-512	3-4-6-3	1-2-2-2
Res38-narrow	24-48-96-192-384	4-5-5-3-1	1-2-2-2-2
Res38	64-128-256-512-1024	4-5-5-3-1	1-2-2-2-2
Res40-narrow	24-48-96-192	3-5-7-4	1-2-2-2
Res40	64-128-256-512	3-5-7-4	1-2-2-2

TABLE I: Network configurations.

compare how different architectures perform using a relatively small number of parameters and FLOPs (we count floating point multiplications). Our experiments are performed with the PyTorch software [46]. In a third experiment, we also show that our lean operators can be implemented efficiently and, for $g = c_{in}$, outperform the separate application of depth-wise and 1×1 operators using the highly optimized package cudnn.

As our focus is on the performance of the lean convolution operators, we use the established ResNet architectures as baseline for comparison, and we use the same structure of those ResNets, only with lean convolutions. Our ResNet networks consist of several blocks that are preceded by an opening convolutional 3×3 layer, which initially increases the number of channels. Then, there are several blocks, each consisting of a ResNet-based steps like Eq. (1). Each convolution is followed by ReLU and batch normalization as described in (1). To increase the number of channels and to down sample the image, we concatenate the feature maps with a depth-wise convolution applied to the same channels, thus doubling the number of channels. This is followed by an average pooling layer. The last block consists of a pooling layer that averages each channel's feature map to a single pixel, and we use a linear classifier with softmax and cross entropy loss. For the ImageNet dataset, we used the ResNeXt configuration [32], which starts with an opening layer of a strided 7×7 convolution followed by max pooling. Its basic step is combined from 3 convolutions (with ReLU and normalization), where the middle is a grouped 3×3 and the rest are 1×1 convolutions. To limit the number of parameters in our LeanResNeXt, we replace each 1×1 convolution with the lean convolution, and the 3×3 convolution with a 5-pt convolution. The last block is identical to that of ResNet. In this ResNeXt architecture, we reduced the number of channels compared to the original network, since its cost is dominated by 1×1 convolutions.

Although the architectures of LeanResNets and ResNet are similar, the former employs efficient convolutions such as (6). The convolution sizes of MobileNetV2 and ShiftResNet were chosen such that the size of each expanded (by 6) 1×1 convolution is equivalent to the size of a square 1×1 convolution of LeanResNet. The architecture ShuffleNetV2 is evaluated with the configurations (0.5x, 1.0x, 1.5x, 2.0x) in [25].

A. Image Classification

We consider the CIFAR10, CIFAR100, STL10, ImageNet, and tinyImageNet200 datasets. The CIFAR-10/100 datasets [47] consists of 60k natural images of size 32×32 with labels assigning each image into one of ten categories

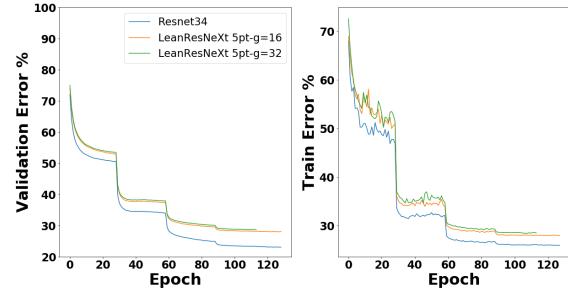


Fig. 3: Validation (left) and Train (right) error per epoch for the ImageNet dataset.

(for CIFAR10) or 100 categories (for CIFAR100). The data are split into 50K training and 10K test sets. The STL-10 dataset [48] contains 13K color-images each of size 96×96 that are divided into 5K training and 8K test images that are split into the ten categories. The ImageNet [49] challenge ILSVRC consists of over 1.28M images of size 224×224 with labels assigning each image into one of 1000 classes where each class has 50 validation images. The tinyImageNet200 [50] is a subset of the ImageNet dataset, and consists of 110K labeled images of size 64×64 belonging to 200 classes, where each class has 500 training images and 50 validation images.

For each of the data sets we used a different configuration, according to the difficulty of that data set. Table I summarizes the network weights that we use, which differ in the number of channels and the number of repetitions for each layer. As optimization strategy for TinyImageNet200 we use momentum SGD with a mini-batch size of 64 for 300 epochs. The learning rate start at 0.05 and is reduced to 0.01, 0.005 and 0.001 after the epochs 75, 150 and 225 respectively. The weight decay is 0.0001 and the momentum is 0.9. The strategy for the other data sets is similar, with slight changes in the number of epochs, batch sizes and the timing for reducing the learning rate. We use standard data augmentation, i.e., random resizing, cropping and horizontal flipping.

Our classification results are given in Tables II- III, where we chose several representative configurations of groups for the lean convolutions. The results show that our architecture is on par with and in some cases better than other networks. There is no preferred architecture between all options, but our architecture has the advantage of simplicity and resemblance to a standard and reliable ResNet network, which, as expected, yields better accuracy than all the other network at the expense of more parameters and cost. In Fig. 3 we show the training and validation convergence plots of the architectures for the ImageNet. The plots show that the convergence of the LeanResNeXt is similar to that of the standard ResNet.

The influence of groups and stencils size: In this set of experiments we demonstrate the classification accuracy of LeanResNet with different configuration of grouping and stencil sizes on CIFAR10 and CIFAR100 data sets. We use small networks so that the differences in performance are more obvious. Table IV presents the classification results. The configuration of $g = c_{in}/q$ indicates that the group sizes are equal throughout the layers, and leads to more FLOPs but less

Architecture	Network	CIFAR10		CIFAR100	
		Params\ FLOPs[M]	Test acc.	Network	Params\ FLOPs[M]
ResNet	Res24	4.7\212	94.5%	Res40	28.9\1490
ResNet (small)	Res24-narrow	0.66\53	92.0%	Res40-narrow	3.8\239
MobileNetV2	Res24*	0.50\33	91.7%	Res40*	3.1\167
ShuffleNetV2	0.5x	0.35\42	91.6%	1.5x	2.6\375
ShiftResNet	Res24*	0.49\31	90.7%	Res40*	3.1\201
LeanResNet 5-pt _{DW} [ours]	Res24	0.53\26	92.8%	Res40	3.3\167
LeanResNet 5-pt _{g=16} [ours]	Res24	0.65\31	93.7%	Res40	4.0\203
LeanResNet 3-pt _{g=8} [ours]	Res24	0.66\31	93.4%	Res40	4.1\203
Architecture	Network	STL10		TinyImageNet200	
		Params\ FLOPs[M]	Test acc.	Network	Params\ FLOPs[M]
ResNet	Res24	4.7\1908	86.6%	Res38	40.9\4816
ResNet (small)	Res24-narrow	0.66\277	82.5%	Res38-narrow	5.8\831
MobileNetV2	Res24*	0.50\302	84.0%	1.4 [22]	4.7\661
ShuffleNetV2	1.0x	1.2\608	81.7%	2.0X [25]	5.7\740
ShiftResNet	Res24*	0.49\361	84.0%	Res38*	4.5\793
LeanResNet 5-pt _{DW} [ours]	Res24	0.53\235	84.0%	Res38	4.7\488
LeanResNet 5-pt _{g=16} [ours]	Res24	0.65\275	86.5%	Res38	5.9\590
LeanResNet 3-pt _{g=8} [ours]	Res24	0.66\275	85.4%	Res38	5.9\590

TABLE II: Comparison of classification results for small datasets. To make a fair comparison, we seek to match the number of parameters and FLOPs for each network. For MobileNetV2 and ShiftResNet we use expansion of $\epsilon = 6$, and choose the width (number of channels) to be approximately $\sqrt{6}$ smaller than the width of LeanResNet, so that their number of parameters and FLOPs are comparable. We denote this by *. Since the images here are smaller than those of ImageNet, we removed the strides from the first steps of the networks MobileNetV2 and ShuffleNetV2.

Architecture	Network	ImageNet	
		Params\ FLOPs[M]	Val. acc.
Resnet	Res34 [30]	21.8\3600	74.0%
ResNeXt	Res50 [32]	25.0\4100	77.8%
LeanResNeXt 5-pt _{g=32}	Res34	3.6\680	71.7%
LeanResNeXt 5-pt _{g=16}	Res34	3.9\630	72.1%
MobileNetV2	1.0 [22]	3.4\300	71.9%
ShuffleNetV2	1.5x [25]	3.5\299	72.6%
ShiftResNet	ShiftNet-A[27]	4.1\-	70.1%

TABLE III: Comparison of classification results for ImageNet using different compact networks.

Architecture	Groups	Network	CIFAR10		CIFAR100	
			Params \ FLOPs[M]	Test acc.	Network	Params \ FLOPs [M]
ResNet 9-pt	—	Res18	2.7\181	94.3%	Res34	21.1\1325
ResNet 5-pt	—	Res18	1.5\101	94.0%	Res34	11.8\739
ResNet 3-pt	—	Res18	0.92\62	93.4%	Res34	7.1\445
LeanResNet 9-pt	DW (g = c_{in})	Res18	0.33\25	91.1%	Res34	2.5\160
LeanResNet 9-pt	g = 32	Res18	0.39\27	91.6%	Res34	3.0\188
LeanResNet 9-pt	g = 16	Res18	0.46\32	92.0%	Res34	3.6\225
LeanResNet 9-pt	g = 8	Res18	0.62\42	92.9%	Res34	4.8\298
LeanResNet 9-pt	g = $c_{in}/32$	Res18	0.80\107	93.6%	Res34	4.3\426
LeanResNet 9-pt	g = $c_{in}/16$	Res18	0.56\64	93.0%	Res34	3.4\289
LeanResNet 9-pt	g = $c_{in}/8$	Res18	0.43\43	92.5%	Res34	2.9\220
LeanResNet 5-pt	DW (g = c_{in})	Res18	0.32\23	91.0%	Res34	2.4\156
LeanResNet 5-pt	g = 32	Res18	0.35\25	91.7%	Res34	2.7\170
LeanResNet 5-pt	g = 16	Res18	0.39\27	92.5%	Res34	3.0\188
LeanResNet 5-pt	g = 8	Res18	0.46\32	92.8%	Res34	3.6\225
LeanResNet 5-pt	g = $c_{in}/32$	Res18	0.56\64	93.9%	Res34	3.4\289
LeanResNet 5-pt	g = $c_{in}/16$	Res18	0.43\43	93.2%	Res34	2.9\220
LeanResNet 5-pt	g = $c_{in}/8$	Res18	0.37\33	92.8%	Res34	2.6\186
LeanResNet 3-pt	DW (g = c_{in})	Res18	0.31\23	90.5%	Res34	2.4\153
LeanResNet 3-pt	g = 32	Res18	0.33\23	91.0%	Res34	2.6\160
LeanResNet 3-pt	g = 16	Res18	0.35\24	91.4%	Res34	2.7\170
LeanResNet 3-pt	g = 8	Res18	0.39\27	92.4%	Res34	3.0\188
LeanResNet 3-pt	$c_{in}/32$	Res18	0.43\43	92.5%	Res34	2.9\220
LeanResNet 3-pt	$c_{in}/16$	Res18	0.37\33	92.2%	Res34	2.6\186
LeanResNet 3-pt	$c_{in}/8$	Res18	0.35\27	92.0%	Res34	2.5\169

TABLE IV: Classification results for the CIFAR10/100 datasets. Keeping the same basic architectures we study the impact on groups and stencil sizes on the test accuracy.

weights than the constant number of groups $g = q$. Since g linearly increases as a function of number of channels, we get relatively dense convolutions at the first layers of the network (large maps, small # of channels) and sparser convolutions at the last layers of the network (small maps, large # of channels). In these examples, having more parameters at the beginning of the network increases the accuracy, at the expense of more FLOPs. The configuration is advantageous when having a low number of parameters is more crucial than FLOPs. On the other hand, keeping the number of groups constant adds a fixed proportion of parameters and FLOPs to the 1×1 convolution, and should be chosen in cases where FLOPs cost as considerably as number of parameters. As a result, the optimal configuration for an application can be wisely chosen based on the limitations of the target device. If there is a constraint on the number of FLOPs, then a constant number of groups can be beneficial, but if the emphasis is on a lower number of parameters, then, then a configuration of $g = c_{in}/q$ will be more suitable for the application. In addition, the table shows that by adding a small addition of parameters to the lean network yields higher accuracy, which gets closer to the considerably larger fully-coupled network.

B. Semantic Segmentation

We demonstrate the effectiveness of our proposed network for the semantic segmentation task, which is used for autonomous vehicles, for example, which require real-time inference, and by design have less computational power. We adopt two popular segmentation architectures, namely U-net [51] and DeepLabV3 [52]. Specifically, we use the general U-net architecture built on top of ResNet as a backbone. That is, we adopt an encoder-decoder scheme, where the encoder is of a standard ResNet architecture and the decoder is based on upscaling operations and transposed convolutions within a ResNet block. Similarly to the classification task, the U-net based on ResNet is used as a baseline. With these settings, we use the baseline with similar networks incorporating various backbones as encoders: MobileNetV2, ShuffleNetV2, ShiftResNet, and ours. As part of the decoders, we perform convolutions to decrease the number of channels and then perform upsampling, such that in the last layer we have an image, the same size of the labeled image. For the second batch of experiments, we adopt DeepLabV3's ASPP module to be our decoder, and compare the various encoders. We test the networks using two popular datasets - Cityscapes (fine annotated) and PASCAL VOC 2012. Cityscapes [53] contains 5000 finely-annotated images with 19 categories ranging from road, vehicles, trees and humans. We use the standard train-validation data split as in [53], i.e.; 2975 and 500 for training and validation, respectively. We resize the images from 1024×2048 to 512×1024 due to memory and computational limitations. As shown in [54] the reduction in performance is only marginal when down-sampling the images this way. In addition, we use standard augmentations like random horizontal flips and random rotation of 10 degrees. The PASCAL VOC 2012 dataset contains 1,464 training, 1,449 validation, and 1,456 test images over 21 object classes

Type	# Channels	# Steps	Strides
MobileNetV2	32-64-128-256	1-2-3-2	1-2-2-2
ShuffleNetV2	116-232-464-512	7-10-10-1	1-2-2-2
ShiftResNet	64-128-256-320	3-4-6-4	1-2-2-2
(Lean)ResNet34	48-96-192-384	3-4-6-4	1-2-2-2
MobileNetV2	16-24-32-64-96-160-320	1-2-3-4-3-3-1	1-2-2-2-1-1-1
ShuffleNetV2 X1.0	24-116-232-464	1-3-7-3	2-2-2-1
LeanResNets	64-128-256-512	3-4-6-3	2-2-2-1
ResNet34	64-128-256-512	3-4-6-3	2-2-2-1

TABLE V: Semantic segmentation network configurations. The upper and lower tables refer to the backbones in the U-net and DeepLabV3 architectures, respectively.

U-net's backbone architecture	Cityscapes			
	Params[M]	FLOPs[B]	Val. acc.	mIoU
ResNet34	25.95	228.4	94.1%	65.1%
MobileNetV2	3.50	31.0	92.1%	56.9%
ShuffleNetV2	3.43	36.5	90.7%	53.5%
ShiftResNet	3.82	48.0	93.0%	60.0%
LeanResNet34 5pt _{DW}	3.53	31.6	92.8%	57.9%
LeanResNet34 5pt _{g=16}	4.12	36.0	92.8%	60.2%
LeanResNet34 3pt _{DW}	3.41	30.1	92.8%	59.2 %
LeanResNet34 3pt _{g=8}	3.96	34.9	93.1%	61.7%
DeepLabV3 backbone architecture	PASCAL VOC 2012			
	Params[M]	FLOPs[B]	Val. acc.	mIoU
ResNet34	25.42	39.5	94.1%	73.1%
MobileNetV2	4.52	8.35	92.18%	66.9%
ShuffleNetV2 [1.0]	4.55	9.04	90.57%	65.54%
LeanResNet34 5pt _{DW}	5.12	8.17	92.48%	67.40%
LeanResNet34 5pt _{g=8}	6.41	9.05	93.30%	70.31%
LeanResNet34 3pt _{DW}	4.91	6.35	92.11%	66.23%

TABLE VI: Comparison of our semantic segmentation results with other compact networks.

(including a background class). We follow standard usage of the augmented PASCAL dataset [55] which brings the training set to 10,582 images. In the training process, we use the ADAM [56] optimizer with a minibatch size of 8 and weight decay of 0.01. The initial learning rate is 1e-4 and we employ an adaptive learning rate reduction, where upon stagnation of the mIoU metric for more than 5 epochs, the learning rate is decreased by a factor of 10. We use the Focal loss [57] as it penalizes wrong segmentations more than correct ones, relative to Cross-Entropy loss. In table V we summarize the configurations used for the segmentation experiments, where again, we tried to configure the sizes of all the compact architectures to have similar number of parameters and FLOPs.

Table VI shows the segmentation results. Similarly to the classification results, the lean networks yield performance that is comparable to the other compact architectures. In particular, the grouped lean versions again yield the best accuracy among compact networks, with a slight increase in the parameters and FLOPs. Table VII shows the segmentation accuracy per class, and Fig. 4 shows two example images from the data set and their segmentation result using different networks. For the U-net architecture, the encoder networks are trained from scratch to give equal starting point to all experiments, employing the same training scheme. The networks which adopt the ASPP module as decoder were trained starting from pre-trained models on ImageNet [49].

Method	road	swalk	build.	wall	fence	pole	light	sign	veg.	terrain	sky	person	rider	car	truck	bus	train	mbike	bike	mIoU
ResNet34	97.2	78.3	89.1	40.5	44.2	49.8	53.2	63.9	90.0	58.0	92.4	70.9	47.1	91.8	62.6	66.6	36.8	40.0	63.9	65.1
MobileNetV2	95.5	70.4	86.4	31.9	34.3	41.4	41.3	56.3	88.4	54.3	89.6	63.3	34.8	87.8	39.4	52.1	35.0	20.9	57.7	56.9
ShuffleNetV2	95.6	69.5	83.4	35.9	34.5	16.9	26.9	39.4	83.7	49.0	85.0	49.8	27.7	84.9	52.3	62.3	44.9	25.4	49.6	53.5
ShiftResNet	96.5	74.3	87.1	38.1	37.6	43.8	43.9	55.0	88.7	54.7	92.0	64.5	39.3	88.9	49.9	62.2	32.8	31.5	59.0	60.0
LeanResNet 5-pt _{DW}	96.6	74.5	87.0	41.5	40.7	38.0	44.3	54.0	88.2	53.5	91.0	62.6	36.7	88.8	50.7	62.4	31.0	28.0	56.1	59.2
LeanResNet 5-pt _{g=16}	96.6	74.3	86.9	44.0	42.6	39.1	41.7	53.7	88.1	54.5	90.8	62.6	37.7	89.0	53.5	61.8	42.0	28.3	55.8	60.2
LeanResNet 3-pt _{DW}	96.6	74.5	87.0	41.5	40.7	38.0	44.3	54.0	88.2	53.5	91.0	62.6	36.7	88.8	50.7	62.4	31.0	28.0	56.1	59.2
LeanResNet 3-pt _{g=8}	96.5	74.2	87.2	43.9	40.4	41.7	43.4	56.2	88.6	54.9	91.4	63.1	37.8	89.3	58.9	66.8	49.7	30.7	58.4	61.7

TABLE VII: Per-class results on Cityscapes validation set using the U-net architecture.

Method	background	airplane	bicycle	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motorbike	person	plant	sheep	sofa	train	monitor	mIoU
ResNet34	93.6	88.2	40.1	84.1	63.8	77.9	93.7	82.9	88.0	36.7	74.1	56.0	79.4	77.7	81.0	84.3	56.5	80.8	42.9	82.4	70.8	73.1
MobileNetV2	91.7	77.6	38.7	72.4	55.6	66.3	85.1	79.3	80.7	32.7	73.2	48.9	73.8	72.0	75.8	78.1	46.6	79.4	37.9	76.7	62.8	66.9
ShuffleNetV2	91.5	79.8	38.1	73.5	57.2	58.2	86.1	81.6	82.1	27.0	65.1	41.7	72.5	73.0	76.8	79.0	42.3	72.7	36.3	77.8	63.4	65.5
LeanResNet 5-pt _{DW}	92.0	81.3	38.3	75.5	60.2	63.7	86.0	81.3	80.9	28.4	71.9	40.7	71.7	76.3	78.6	79.6	48.8	79.4	38.5	77.5	64.6	67.4
LeanResNet 5-pt _{g=8}	92.9	84.7	39.0	79.3	66.6	70.9	89.8	79.6	86.3	32.3	69.7	49.2	75.0	76.4	80.9	81.9	50.4	78.2	45.0	80.6	67.8	70.3
LeanResNet 3-pt _{DW}	91.8	78.4	36.9	76.9	62.7	66.9	88.1	79.6	82.3	25.1	67.2	44.5	75.2	72.9	74.4	77.9	44.7	73.1	30.1	79.6	62.1	66.2

TABLE VIII: Per-class results on PASCAL VOC 2012 validation set using the DeepLabV3 architecture.

C. Computational Performance

We compare the latency of our CUDA implementation of the lean convolution with two other combination of layers, comprised of a 1×1 convolution that is followed by a depthwise convolution. In one combination we use $c_{in} = c_{out}$, and in the other $c_{in} \approx 6c_{out}$, but with the same number of weights. Such layers are applied in [22]. We compare the runtime of a typical network: the first layer consists of 16 channels of 512×512 maps, and the maps are coarsened by a factor of 2 when the channels increase by a factor of 2 (i.e., for 512 channels the images are of size 16). We use a batch size of 64, and compare the runtime of a NVIDIA GeForce 1080Ti GPU for the task. The implementation for the other convolutions is based on PyTorch's 1×1 and grouped convolutions using CUDA 9.2. Figure 5 summarizes the results. The depthwise convolutions dominate the low channels layers, while all combination converge to the cost of the 1×1 convolution as the channels increase (and the depthwise layer becomes negligible). Our implementation of (6) is comprised of a standard 4-point convolution for each channel followed by a matrix multiplication using cublas for the 1×1 part, to use the highly optimized gemm kernel. Our implementation is faster because the shiftIm2Col approach is not efficient for small group sizes (1 in this case). The clear advantage of the lean operator over the expanded combination is the less feature maps that undergo the spatial convolution. Although our implementation applies the 1×1 and depthwise convolutions separately for each sample, our experiments show that this yields a performance gain compared to a completely separate multiplication for the whole mini-batch.

V. CONCLUSION

We present LeanConvNets, a family of efficient CNNs that reduce the number of weights, and floating point operations with minimal loss of accuracy. LeanConvNets can be obtained from existing CNNs by replacing fully-coupled convolution operators by lean operators that are the sum of grouped and 1×1 convolutions. The group size serves as a hyperparameter that allows the user to trade off computational cost and accuracy. Additional savings can be realized by the proposed five-point and three-point stencils. Those savings will be more pronounced for 3D and 4D imaging data.

In our experiments, we apply various configurations of LeanConvNets to image classification and segmentation tasks. In our tests, the LeanConvNets perform slightly better than other reduced architectures, and is almost as effective as their fully-coupled counterpart. We also demonstrate in a direct comparison that the addition of depth-wise and 1×1 convolution reduces the computational time.

Our future research aims to further optimize implementation of the lean convolutions on GPUs, as well as investigate optimization of such implementation on other devices. In addition, it is worthy to investigate and characterize the hyper-parameter choices of the lean convolution (groups, stencil size, multiplication algorithm), as these choices should be guided by the hardware [58]. We also plan to examine the efficiency of the lean operators in challenging 3D applications such as video analysis on limited devices [59], where the small stencil size is more beneficial.

ACKNOWLEDGEMENTS

LRs work is supported by the US National Science Foundation (NSF) award DMS 1751636. This research was partially supported by grant no. 2018209 from the United States - Israel Binational Science Foundation (BSF), Jerusalem, Israel. ME is supported by Kreitman High-tech scholarship.

REFERENCES

- [1] Y. LeCun, B. E. Boser, and J. S. Denker, "Handwritten digit recognition with a back-propagation network," in *Advances in neural information processing systems*, 1990, pp. 396–404.
- [2] R. Raina, A. Madhavan, and A. Y. Ng, "Large-scale deep unsupervised learning using graphics processors," in *Proceedings of the 26th annual international conference on machine learning*. ACM, 2009, pp. 873–880.
- [3] P. Agrawal and S. Ganapathy, "Modulation filter learning using deep variational networks for robust speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 244–253, May 2019.
- [4] H. Purwins, B. Li, T. Virtanen, J. Schilter, S. Chang, and T. Sainath, "Deep learning for audio signal processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 206–219, May 2019.
- [5] A. Krizhevsky, I. Sutskever, and G. Hinton, "Imagenet classification with deep convolutional neural networks," *Adv Neural Inf Process Syst*, vol. 61, p. 10971105, 2012.
- [6] J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy, B. Shuai, T. Liu, X. Wang, G. Wang, J. Cai, and T. Chen, "Recent advances in convolutional neural networks," *Pattern Recognition*, vol. 77, pp. 354–377, May 2018.

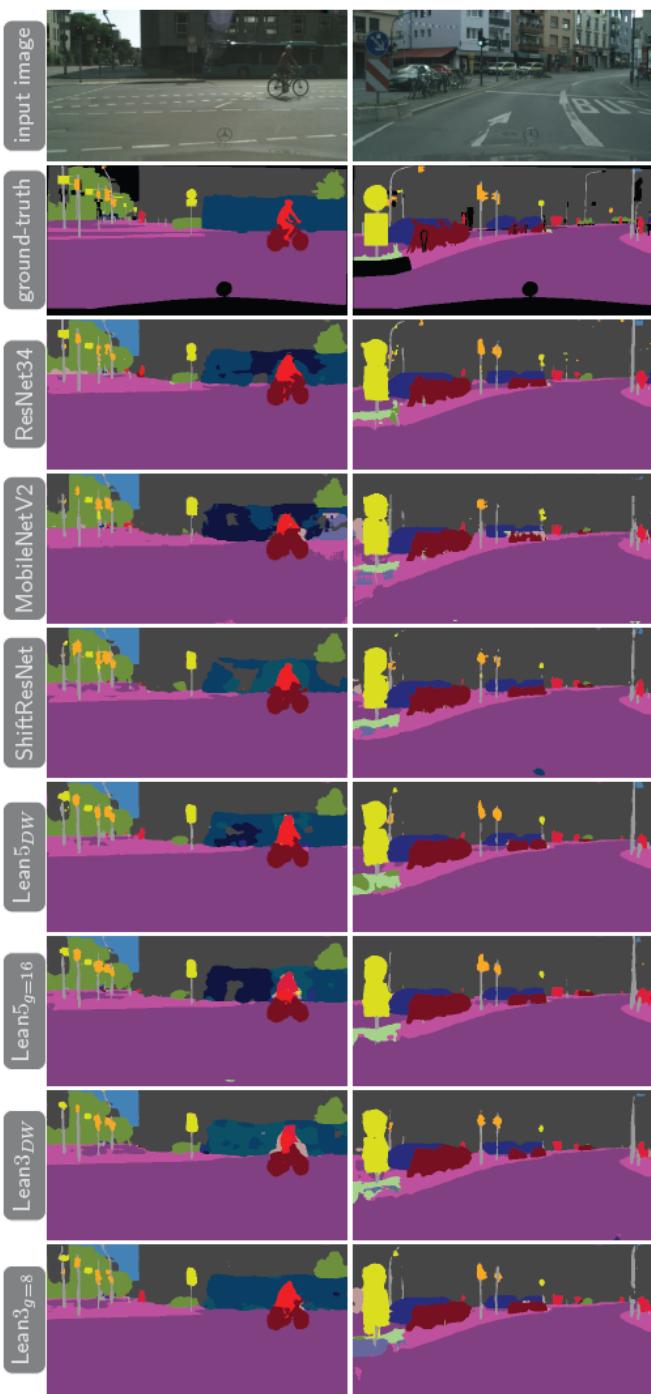


Fig. 4: Visualization of the semantic segmentation results of different networks for two images from the Cityscapes dataset.

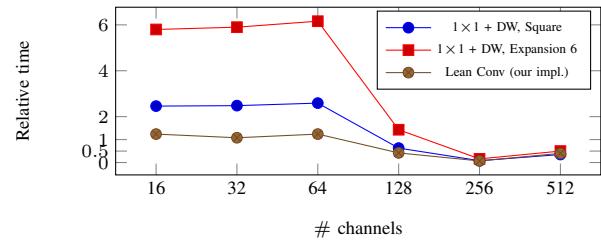


Fig. 5: Relative timings of reduced convolutions compared to a 3×3 convolution (lower is faster). The expanded and square 1×1 convolutions has the same number of weights.

- filters for efficient ConvNets,” <https://arxiv.org/abs/1608.08710>, 2017.
- [12] J.-H. Luo, J. Wu, and W. Lin, “ThiNet: A filter level pruning method for deep neural network compression,” in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
 - [13] T.-W. Chin, C. Zhang, and D. Marculescu, “Layer-compensated pruning for resource-constrained convolutional neural networks,” *arXiv preprint arXiv:1810.00518*, 2018.
 - [14] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, “Learning structured sparsity in deep neural networks,” in *Advances in Neural Information Processing Systems*, 2016, pp. 2074–2082.
 - [15] S. Changpinyo, M. Sandler, and A. Zhmoginov, “The power of sparsity in convolutional neural networks,” *arXiv preprint arXiv:1702.06257*, 2017.
 - [16] S. Han, J. Pool, S. Narang, H. Mao, S. Tang, E. Elsen, B. Catanzaro, J. Tran, and W. J. Dally, “Dsd: regularizing deep neural networks with dense-sparse-dense training flow,” *arXiv preprint arXiv:1607.04381*, vol. 3, no. 6, 2016.
 - [17] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, “Binarized neural networks,” in *Advances in neural information processing systems*, 2016, pp. 4107–4115.
 - [18] H. Li, S. De, Z. Xu, C. Studer, H. Samet, and T. Goldstein, “Training quantized nets: A deeper understanding,” in *Advances in Neural Information Processing Systems*, 2017, pp. 5811–5821.
 - [19] R. Banner, I. Hubara, E. Hoffer, and D. Soudry, “Scalable methods for 8-bit training of neural networks,” in *Advances in Neural Information Processing Systems*, 2018, pp. 5145–5153.
 - [20] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz, “Pruning convolutional neural networks for resource efficient transfer learning,” *arXiv preprint arXiv:1611.06440*, vol. 3, 2016.
 - [21] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, “MobileNets: Efficient convolutional neural networks for mobile vision applications,” *arXiv preprint arXiv:1704.04861*, 2017.
 - [22] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, “MobileNetV2: Inverted residuals and linear bottlenecks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4510–4520.
 - [23] M. Wang, B. Liu, and H. Foroosh, “Design of efficient convolutional layers using single intra-channel convolution, topological subdivisioning and spatial ‘bottleneck’ structure,” *arXiv preprint arXiv:1608.04337*, 2016.
 - [24] X. Zhang, X. Zhou, M. Lin, and J. Sun, “ShuffleNet: An extremely efficient convolutional neural network for mobile devices,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6848–6856.
 - [25] N. Ma, X. Zhang, H.-T. Zheng, and J. Sun, “ShuffleNet V2: Practical guidelines for efficient CNN architecture design,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 116–131.
 - [26] M. Tan and Q. V. Le, “MixConv: Mixed Depthwise Convolutional Kernels,” [arXiv.org](https://arxiv.org/), Jul. 2019.
 - [27] B. Wu, A. Wan, X. Yue, P. Jin, S. Zhao, N. Golmant, A. Gholaminejad, J. Gonzalez, and K. Keutzer, “Shift: A zero flop, zero parameter alternative to spatial convolutions,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9127–9135.
 - [28] Y. Chen, H. Fan, B. Xu, Z. Yan, Y. Kalantidis, M. Rohrbach, S. Yan, and J. Feng, “Drop an Octave: Reducing Spatial Redundancy in Convolutional Neural Networks with Octave Convolution,” [arXiv.org](https://arxiv.org/), Apr. 2019.
 - [29] Z. Qin, Z. Zhang, D. Li, Y. Zhang, and Y. Peng, “Diagonalwise refactoring: An efficient training method for depthwise convolutions,”

- in *2018 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2018, pp. 1–8.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [31] —, “Identity mappings in deep residual networks,” in *European Conference on Computer Vision*. Springer, 2016, pp. 630–645.
- [32] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, “Aggregated residual transformations for deep neural networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1492–1500.
- [33] H. Robbins and S. Monro, “A Stochastic Approximation Method,” *Ann. Math. Stat.*, 1951. [Online]. Available: <http://www.jstor.org/stable/2236626>
- [34] L. Bottou, F. E. Curtis, and J. Nocedal, “Optimization methods for large-scale machine learning,” *arXiv preprint arXiv:1606.04838*, 2016.
- [35] E. Haber and L. Ruthotto, “Stable architectures for deep neural networks,” *Inverse Problems*, vol. 34, no. 1, 2017.
- [36] B. Chang, L. Meng, E. Haber, L. Ruthotto, D. Begert, and E. Holtham, “Reversible architectures for arbitrarily deep residual neural networks,” in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [37] E. Weinan, “A Proposal on Machine Learning via Dynamical Systems,” *Communications in Mathematics and Statistics*, vol. 5, no. 1, pp. 1–11, Mar. 2017.
- [38] P. Chaudhari, A. Oberman, S. Osher, S. Soatto, and G. Carlier, “Deep Relaxation: Partial Differential Equations for Optimizing Deep Neural Networks,” pp. 1–22, 2017.
- [39] Y. Lu, A. Zhong, Q. Li, and B. Dong, “Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations,” in *International Conference on Machine Learning (ICML)*, 2018.
- [40] T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, “Neural ordinary differential equations,” in *Advances in Neural Information Processing Systems*, 2018, pp. 6571–6583.
- [41] E. Haber, K. Lensink, E. Treister, and L. Ruthotto, “Imexnet: A forward stable deep neural network,” *International Conference on Machine Learning (ICML)*, 2019.
- [42] L. Ruthotto and E. Haber, “Deep neural networks motivated by partial differential equations,” *arXiv preprint arXiv:1804.04272*, 2018.
- [43] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, “Inception-v4, inception-resnet and the impact of residual connections on learning,” in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [44] <https://devblogs.nvidia.com/cutlass-linear-algebra-cuda/>, [Online; accessed May-2019].
- [45] <http://icl.cs.utk.edu/magma/index.html>, [Online; accessed May-2019].
- [46] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, “Automatic differentiation in pytorch,” in *Advances in Neural Information Processing Systems*, 2017.
- [47] A. Krizhevsky and G. Hinton, “Learning multiple layers of features from tiny images,” Citeseer, Tech. Rep., 2009.
- [48] A. Coates, A. Ng, and H. Lee, “An analysis of single-layer networks in unsupervised feature learning,” in *Proceedings of the 14th AISTATS*, 2011, pp. 215–223.
- [49] <http://www.image-net.org/>, [Online; accessed May-2019].
- [50] <https://tiny-imagenet.herokuapp.com/>, [Online; accessed May-2019].
- [51] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [52] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, “Rethinking atrous convolution for semantic image segmentation,” *arXiv preprint arXiv:1706.05587*, 2017.
- [53] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, “The cityscapes dataset for semantic urban scene understanding,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3213–3223.
- [54] T. Pohlen, A. Hermans, M. Mathias, and B. Leibe, “Full-resolution residual networks for semantic segmentation in street scenes,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4151–4160.
- [55] B. Hariharan, P. Arbeláez, L. Bourdev, S. Maji, and J. Malik, “Semantic contours from inverse detectors,” in *2011 International Conference on Computer Vision*. IEEE, 2011, pp. 991–998.
- [56] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [57] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.
- [58] D. Marculescu, D. Stamoulis, and E. Cai, “Hardware-aware machine learning: modeling and optimization,” in *Proceedings of the International Conference on Computer-Aided Design*. ACM, 2018, p. 137.
- [59] H. Fassold, S. Wechtitsch, M. Thaler, K. Kozłowski, and W. Bailer, “Real-time video quality analysis on mobile devices,” in *Proceedings of the 7th ACM International Workshop on Mobile Video*. ACM, 2015, pp. 23–24.