LEAN: GRAPH-BASED PRUNING FOR CONVOLUTIONAL NEURAL NETWORKS BY EXTRACTING LONGEST CHAINS

Anonymous authors
Paper under double-blind review

ABSTRACT

Neural network pruning techniques can substantially reduce the computational cost of applying convolutional neural networks (CNNs). Common pruning methods determine which convolutional filters to remove by ranking the filters individually, i.e., without taking into account their interdependence. In this paper, we advocate the viewpoint that pruning should consider the interdependence between series of consecutive operators. We propose the *LongEst-chAiN* (LEAN) method that prunes CNNs by using graph-based algorithms to select relevant chains of convolutions. A CNN is interpreted as a graph, with the operator norm of each operator as distance metric for the edges. LEAN pruning iteratively extracts the highest value path from the graph to keep. In our experiments, we test LEAN pruning on several image-to-image tasks, including the well-known CamVid dataset, and a real-world X-ray CT dataset. Results indicate that LEAN pruning can result in networks with similar accuracy but 3–20x fewer convolutional filters than networks pruned with methods that rank filters individually.

1 INTRODUCTION

In recent years, convolutional neural networks (CNNs) have become state-of-the-art for many image-to-image translation tasks (LeCun et al., 2015), including image segmentation (Ronneberger et al., 2015), and denoising (Tian et al., 2020). They are increasingly used as a subcomponent of a larger system, e.g., visual odometry (Yang et al., 2020), as well as in energy-limited and real-time applications (Yang et al., 2017). In these situation, the applicability of high-accuracy CNNs may be limited by large computational resource requirements. Small networks may be more applicable in such settings, but may lack accuracy.

Neural network pruning (Mozer & Smolensky, 1989; Karnin, 1990) has recently gained popularity as a technique to reduce the size of neural networks (Blalock et al., 2020). Neural networks consist of learnable parameters, including the scalar components of the convolutional filters. When pruning, the neural network is reduced in size by removing such scalar parameters while trying to maintain high accuracy. We distinguish between individual parameter pruning (Han et al., 2016), where each parameter of an operation is ranked and pruned separately, and structured pruning (Li et al., 2017; Luo et al., 2017), where entire convolutional filters are ranked and pruned. As convolution operators can only be removed once all scalar parameters of the filter kernel have been pruned, structured pruning is favored over individual pruning when aiming to improve computational performance (Park et al., 2017). In the remainder of this paper, we focus on structured pruning.

Although structured pruning methods take into account the division of a neural network into operations, they do not take into account the fact that the output of the network is formed by a *sequence* of such operations. This has two drawbacks. First, ignoring the effect of a chain of operations may result in sub-optimal filter selection. Second, to significantly reduce evaluation time, a severe pruning regime must be considered, i.e., a pruning ratio (percentage of remaining parameters after pruning) of 1–10%. In this regime, pruning can result in network disjointness, i.e., the network contains sections that are not part of some path from the input to the network output. Some existing pruning methods take into account network structure to a limited degree by setting separate pruning

criteria for each layer in a CNN (Molchanov et al., 2017; Liu et al., 2017). In practice, however, these methods do not contain safeguards to avoid network disjointness.

In this paper, we present a novel pruning method called LongEst-chAiN (LEAN) pruning, which as opposed to conventional pruning approaches uses graph-based algorithms to keep or prune chains of operations collectively. In LEAN, a CNN is represented as a graph that contains all the CNN operators, with the operator norm of each operator as edge weights. We argue that strong subnetworks in a CNN can be discovered by extracting the longest (multiplicative) paths, using computationally efficient graph algorithms. The main focus of this work is to highlight how LEAN pruning can significantly improve the computation speed of CNNs for image-to-image tasks, and obtain high accuracy in the severe pruning regime that is difficult to achieve by pruning the operators individually.

This paper is structured as follows. In Section 2, we explore existing pruning approaches. In Section 3, we outline the preliminaries on CNNs, pruning filters, and the operator norm. Next, in Section 4, we introduce LEAN pruning. We discuss the setup of our experiments in Section 5. In Section 6, we demonstrate the results of the proposed pruning approach on a series of image segmentation problems and report practically realized wall time speedup. Our final conclusions are presented in Section 7.

2 Related work

Reducing the size of neural networks by removing parameters has been studied for decades (Mozer & Smolensky, 1989; Karnin, 1990; Hassibi et al., 1993). Recently, several works have taken into account the structure of the network to some degree. In Lin et al. (2017) filters are pruned at runtime based on the feature maps (Lin et al., 2017). Alternatively, one can prune entire channels (He et al., 2017), or decide which channels to keep so that the feature maps approximate the output of the unpruned network over several training examples (Luo et al., 2017). Apart from pruning the convolutional filters, it is possible to sparsify the parameters of other common components of CNNs, such as the batch normalization layers (Ye et al., 2018; Liu et al., 2017; Huang & Wang, 2018).

Many pruning approaches are aimed at reducing neural network size with little accuracy drop (Dong & Yang, 2019; He et al., 2019; Molchanov et al., 2019; Zhao et al., 2019), as opposed to sacrificing accuracy in favor of computation speed. These approaches rarely exceed a pruning ratio of 12–50% (Blalock et al., 2020; Luo et al., 2017; Lin et al., 2019). When a high pruning ratio is used, e.g., a range of 5–10% (Lin et al., 2017; Liu et al., 2019), a significant drop in accuracy is observed. Pruning ratios of 2–10% can be achieved with an accuracy drop of 1-3% by learning-rate rewinding (Renda et al., 2020). However, the reduction in FLOPs was less substantial (1.5–4.8 times), as individual parameters were pruned. A more extensive survey can be found in Blalock et al. (2020).

Criteria for deciding which elements of a neural network to prune have been extensively studied. A parameter's importance is commonly scored using its absolute value. Whether this is a reasonable metric has been questioned (LeCun et al., 1990). Singular values (which determine certain operator norms) have been used to compress network layers (Denton et al., 2014) and to prune feed-forward networks (Abid et al., 2002). Efficient methods for the computation of singular values have been developed for convolutional layers (Sedghi et al., 2019). Furthermore, a definition of ReLU singular values was proposed recently with an accompanying upper bound (Dittmer et al., 2019).

3 Preliminaries

3.1 CNNs for segmentation

A common image to image translation task is semantic image segmentation. The goal of semantic image segmentation is to assign a class label to each pixel in an image. A segmentation CNN computes a function $f : \mathbb{R}^{m \times n} \to [0, 1]^{k \times m \times n}$, which specifies the probability of each pixel being in one of the k classes for an $m \times n$ image.

CNNs are composed of layers of operations which pass images from one layer to the next. Every operation, e.g., convolution, has an input $\mathbf{x}$ and output $\mathbf{y}$. The input and output consist of one or more images, called channels. For clarity, we distinguish throughout this paper between an **operation**, which may have several input and output channels, and an **operator**, which computes the

relation between a single input channel and a single output channel. For instance, in a convolutional operation with input channels $\mathbf{x}_1, \ldots, \mathbf{x}_N$, an output channel $\mathbf{y}_j$ is computed by convolving input images with learned filters

$$\mathbf{y}_j = \left(\sum_{i=1}^{N} h_{ij} * \mathbf{x}_i \right) + b_j. \tag{1}$$

Here h_{ij} is the filter related to the convolution operator that acts between channel $\mathbf{x_i}$ and $\mathbf{y_j}$, and b_j is an additive bias parameter. In a similar way, every CNN operation produces an output which consists of a number of channels. The exact arrangement of operations, and connections between them, depends on the architecture.

A common operator to downsample images is the strided convolution. The stride defines the step size of the convolution kernel. A convolution with stride s defines a map $h : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{\frac{m}{s} \times \frac{n}{s}}$. Upsampling images can be done by transposed convolutions. Transposed convolutions intersperse the input image pixels with zeroes so that the output image has larger dimensions.

In addition to convolution operators, other common operators such as pooling and batch normalization are often used. A pooling operator downsamples an image by combining pixel values into one. Usually, the average or maximum value of a 2×2 patch is used. A batch normalization operator (Ioffe & Szegedy, 2015) is used to normalize the input images for convolutional layers. A batch normalization operator takes an image $\mathbf{x}_i$ and scales and shifts it by

$$\mathbf{y}_i = \gamma \frac{\mathbf{x}_i - \mu_B}{\sqrt{\sigma_B^2 + \epsilon}} + \beta. \tag{2}$$

Here, γ and β are scaling and bias parameters which are learned during training, and μ_B and σ_B^2 are the running mean and variance of the mini-batch, i.e., the set of images used for the current training step. For an overview of CNN components we refer to Goodfellow et al. (2016).

3.2 Pruning convolution filters

Pruning techniques aim to remove extraneous parameters from a neural network. Several schemes exist to prune parameters from a network, but retraining the network after pruning is critical to avoid significantly impacting accuracy (Han et al., 2015). Pruning a network once after training is called one-shot pruning. Alternatively, a network can be fine-tuned, where the network is repeatedly pruned by a certain percentage and is retrained for a few epochs after every pruning step. Fine-tuning typically gives better results than one-shot pruning (Renda et al., 2020).

Generic pruning algorithm: All pruning methods used in this work make use of the fine-tuning pruning algorithm outlined in Algorithm 1. The selection criteria for determining which filters to keep for each step define the different pruning methods. The pruning ratio `pRatio` is the fraction of remaining convolutions we ultimately want to keep, and `stepRatio` is the fraction of convolutions that is pruned at each step.

Algorithm 1 Fine-tuning pruning algorithm

```
procedure PRUNE(MODEL, PRATIO, NSTEPS, EPOCHS)
    stepRatio ← e^(ln(pRatio)/nSteps)
    for step ← 0 to nSteps do
        pruneParams ← selectPrunePars(model, stepRatio)
        model ← removePars(model, pruneParams)
        for k ← 0 to epochs do
            model ← trainOneEpoch(model, trainData)
    return model
```

As mentioned, we focus on structured pruning in this work. In structured pruning, a common approach to decide which filters to remove is *structured magnitude pruning*. When using structured magnitude pruning, a convolution filter $\mathbf{h} \in \mathbb{R}^{k \times k}$ is scored by its L_1 vector norm $||\mathbf{h}||_1$. Note that the norm of a filter equals the sum of the absolute values of its parameters. Filters with norms below a threshold are pruned. The threshold is determined by sorting a group of filters, removing a percentage based on the pruning ratio. The threshold can be chosen per layer or globally. Pruning a network globally can give higher accuracy than setting thresholds per layer (Blalock et al., 2020).

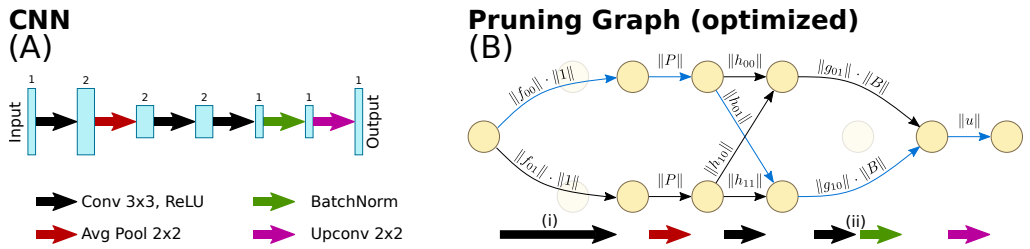

Figure 1: (A) Example CNN architecture with the number of channels indicated above each layer. (B) Associated pruning graph (after optimization). Every channel is a node, and every operator is an edge connecting input and output nodes. The edge weights are the corresponding operator norms. In the optimized pruning graph, ReLU edges (which have a weight of 1) and convolution edges are combined (see (i)), and batch normalization and convolution edges are combined (see (ii)). A possible longest path (blue) is shown.

3.3 Operator norm

As an alternative to the L_1-norm, one can interpret a convolution h as a linear operator which acts on the input image, and score it according to an operator norm. The (induced) operator norm $\|\cdot\|_p$ is defined as

$$\|h\|_p := \sup \left\{ \|h * \mathbf{x}\|_p \;\middle|\; \mathbf{x} \in \mathbb{R}^n, \|\mathbf{x}\|_p = 1 \right\}.$$

A common operator norm is the *spectral norm*, which is induced by the Euclidean norm ($p = 2$). The spectral norm can be obtained by calculating the largest singular value of the matrix H associated with h, $\|h\|_2 = \sigma_{max}(H)$ (Meyer, 2000). A property that we will use is that the spectral norm is submultiplicative, i.e., we have

$$\|AB\| \leq \|A\| \cdot \|B\|, \quad \text{for all} \quad A, B \in \mathbb{R}^{n \times n}. \tag{3}$$

4 Method

The idea behind LEAN is to construct a weighted graph structure formed by the operators of the CNN and having their respective norms as edge weights, such that the longest paths in this graph are selected as important subnetworks. The remaining unselected operators will then be pruned.

4.1 LEAN: creating the pruning graph

In this section, we define the pruning graph that is the basis of the LEAN algorithm. We first describe the graph structure, then the edge weights, and finally some practical limitations and optimizations. As discussed in Section 3.1, we say that every CNN operation has an input $\mathbf{x}$ and an output $\mathbf{y}$, consisting of channels $\mathbf{x}_i$ and $\mathbf{y}_i$. For each channel, we add a single node in the pruning graph. An edge connects two nodes corresponding to input channel $\mathbf{x}_i$ and output channel $\mathbf{y}_i$ if channel $\mathbf{x}_i$ is used in the computation of channel $\mathbf{y}_i$. In the terminology of Section 3.1, each edge corresponds to an operator. For instance, a convolution operation is converted by adding an edge from each input channel $\mathbf{x}_i$ to every output channel $\mathbf{y}_j$, each corresponding to exactly one filter h_{ij}. Certain CNN operations may be performed in-place in practice, but we consider them as separate nodes in the pruning graph. A combined convolution and ReLU operation, for instance, results in a node for the output of the convolution and a separate node for the output of the ReLU, as shown in Figure 1.

To each edge, we assign a weight that equals the operator norm of its corresponding operator. That is, we calculate the maximal scaling that an input could undergo as a result of applying the operator. In this calculation, we ignore any additive bias terms. For instance, applying a batch normalization results in a scaling of $|\gamma| / \sqrt{\sigma_B^2 + \epsilon}$ (see Equation (2)). The scaling of non-linear operators in neural networks is sometimes bounded, as in the case of the ReLU for instance, to which we assign a weight of 1. The calculation of the operator norm of convolutions can be performed efficiently using the Fourier transform (Sedghi et al., 2019). We describe the calculation of these operator norms in

Appendix A, where we extend this approach to down- and upsampling (strided and transposed) convolutions.

The length of a path in the graph is determined by multiplying the edge weights. The aim is to model the norm of the composition of the operators corresponding to the edges. Inspired by Equation (3), we assume that $\|AB\| \approx \|A\| \, \|B\|$ is a reasonably proxy for the combined operator norm.

The outlined procedure can lead to large pruning graphs, but the size of the graph can be reduced. First, the combination of a convolution followed by a ReLU can be combined into a single edge whose weight equals the norm of the convolution. In addition, batch normalization often precedes a convolution. The corresponding two edges can be combined into a single edge whose weight is the product of the two edge weights, preserving the path length. An optimized version of the pruning graph is illustrated in Figure 1.

There are some edge cases to consider. First, some CNNs contain operations that are meant to distribute features throughout the network, but are not implemented with learnable parameters, e.g., residual connections in ResNet (He et al., 2016). We include residual connections in the pruning graph with an edge weight of 1, but label them as unprunable to prevent the residual connections from being removed from the network. Second, we do not consider CNNs with recurrent connections. Therefore, the pruning graph is a Directed Acyclic Graph (DAG).

4.2 LEAN: EXTRACTING CHAINS FROM THE GRAPH

The LEAN method prunes chains of convolutions based on paths in the graph; we **keep** the longest paths (with the highest multiplicative operator norm). We refer to this as LongEst-chAiN (LEAN) pruning. When we perform LEAN pruning, we iteratively extract such paths from the graph until we have reached the pruning ratio. This means that the edges that are not extracted are pruned.

Algorithm 2 LEAN

1: **procedure** LEAN(MODEL, PRUNERATIO)
2: graph ← createPruningGraph(model)
3: retainedConvs ← []
4: **while** fractionRemainingConvs $< 1 - pruneRatio$ **do**
5: bestChain ← $longestPath(graph)$
6: retainedConvs ← retainedConvs + $bestChain$
7: graph ← $removeFromGraph$(graph, $bestChain$)
8: convsToPrune ← convsInModel − retainedConvs
9: **return** convsToPrune

LEAN pruning is incorporated in the fine-tuning procedure. For each pruning step in the fine-tuning procedure, Algorithm 2 is used to select the filters to prune (line 4 in Algorithm 1). For DAGs, the longest path in a graph can be found in $\mathcal{O}(|V| + |E|)$ time, where $|V|$ is the number of nodes, and $|E|$ is the number of edges (Sedgewick & Wayne, 2011). For a CNN with m channels (nodes), and k operators (edges), we can extract a longest path in $\mathcal{O}(m + k)$ time. As we extract at least one operator with every execution of line 5 in Algorithm 2, the longest path algorithm is run at most $k(1 - p)$ times. Therefore, that the worst-case complexity of Algorithm 2 is $\mathcal{O}(k(1 - p)(m + k))$ for a pruning ratio p. Note that unprunable edges can be part of a longest path to extract operator chains, but they do not count towards the pruning ratio.

After every LEAN pruning step is concluded, some post-processing is performed. In some cases there are channels which receive no input data at all, or which are equal to a homogeneous constant image for all input data. We therefore remove nodes without incoming edges as well as nodes where the succeeding batch normalization has running variance below some threshold (10^{-40} by default). A low running variance can occur when the output of a convolution is always zero after applying the ReLU activation function, for instance. Second, bias terms are removed from the CNN when all associated convolution or batch normalization operations are pruned.

5 EXPERIMENTAL SETUP

In our experiments, we compare LEAN pruning to two structured pruning methods across three image segmentation datasets and three CNN architectures: MS-D, U-Net4, and ResNet50. We assess the performance of pruned neural networks across 5 independent runs of fine-tuning (Algorithm 1). For each dataset, we have trained a single model as a starting point for pruning. In every experimental run, the same trained model was pruned.

For both structured pruning methods and LEAN pruning, we measure the pruning ratio as the fraction of convolutions remaining, rather than the fraction of parameters that is pruned when performing the experiments. This means that other parameters, such as batch normalization and bias parameters, are pruned when the associated convolutions are pruned, but do not count towards the pruning ratio. In the experimental results, when a network is pruned to a $k\%$ pruning ratio, this means that $k\%$ of the convolutions remain active, and the other pruned parameters are not counted in this ratio.

The MS-D networks were pruned to a pruning ratio of 1% in 45 steps. The U-Net4, and ResNet50 networks were pruned to a ratio of 0.1% in 30 steps. All were retrained for 5 epochs after each step. We chose relatively severe pruning ratios because we are interested in significant computational speedup. U-Net4 and ResNet50 are pruned to a lower ratio as they have orders of magnitude more convolutions than the MS-D network.

All experiments were run on an AMD Ryzen 3800X processor and NVIDIA GeForce RTX 2070 Super GPU, or on two 8-core Xeon CPUs and four Titan X (12GB) GPUs.

5.1 STRUCTURED PRUNING METHODS

The first method we compare to is **structured magnitude pruning**, i.e., pruning entire filters by their L_1-norm (see Section 3.2), and the second method we compare to is **operator norm pruning**, i.e., pruning entire filters using the operator norm. Here, we choose to use the spectral norm. The difference between LEAN and structured magnitude pruning is 1) the operator norm; 2) the consideration of network structure. The comparison of LEAN to structured operator norm pruning measures the effect of incorporating the network structure.

As explained in Section 3.2, the pruning thresholds in structured pruning can be set globally, or per layer. In the severe pruning regime that we consider in these experiments, setting pruning thresholds per layer can be detrimental as entire layers will be pruned at severe pruning ratios, creating a disjoint network. Therefore, we set pruning thresholds globally.

5.2 CNN ARCHITECTURES

In this section, we describe three fully-convolutional CNN architectures that are used in the experiments: the Mixed-Scale Dense convolutional neural network (MS-D) network (Pelt & Sethian, 2018), U-Net (Ronneberger et al., 2015), and ResNet (He et al., 2016). Table 1 outlines which operators are present in the networks. The networks were trained using ADAM (Kingma & Ba, 2014) with $lr = 0.001$, minimizing the negative log likelihood function.

	Convolution				Batch	#	Edges in
CNN	Strided	Transposed	Dilated	Pooling	normalization	Parameters	pruning graph
MS-D	No	No	Yes	No	No	$4.57 \cdot 10^4$	$5.05 \cdot 10^3$
ResNet50	Yes	No	Yes	Yes	Yes	$3.29 \cdot 10^7$	$1.44 \cdot 10^7$
U-Net4	No	Yes	No	Yes	Yes	$1.48 \cdot 10^7$	$1.84 \cdot 10^6$

Table 1: Operators present in MS-D, ResNet50, U-Net4 architectures.

In our experiments we use MS-D networks as described in Pelt & Sethian (2018) and implemented in Hendriksen (2020). We chose a depth of either 100 or 150 layers, each with 1 channel per layer. All convolutions have a dilated filter of size 3×3, except the final 1×1-layer. The dilations for layer i were set to $1 + (i \mod 10)$. The final layer is excluded from pruning. In preliminary experiments, for both structured pruning methods we observed that almost the entire final layer is immediately pruned, which would create an unfair comparison to LEAN. This may be due to the lack of batch

normalization layers, or zero-initialization of the output layer (Pelt & Sethian, 2018). As the output layer of the MS-D networks uses few FLOPs, it has little impact on computational efficiency.

As U-Net architecture we use a fully-convolutional (FCN) U-Net4 network, i.e., a U-Net with 4 scaling operations. We used a U-Net4 architecture from the PyTorch-UNet repository Milesi (2020). As ResNet architecture, we use an FCN-ResNet50 network. This is a ResNet with 50 layers, and a bottleneck block-structure (He et al., 2016). The ResNet50 model is adapted from PyTorch's model zoo code. We replace the max pooling layers of ResNet and U-Net with average pooling layers as average pooling is a linear operator which can be modeled as a strided convolution. This means that we can efficiently determine the operator norm for average pooling for different kernel sizes and strides. In some cases U-Net with average pooling can perform better than with max pooling (Astono et al., 2020).

5.2.1 DATASETS

For our experiments, we consider three datasets: a high-noise, but relatively simple, segmentation dataset (Pelt & Sethian, 2018); the well-known CamVid dataset (Brostow et al., 2008; 2009); and a real-world CT dataset (Coban & Lucka, 2019; Coban et al., 2020) to test the methods in practice.

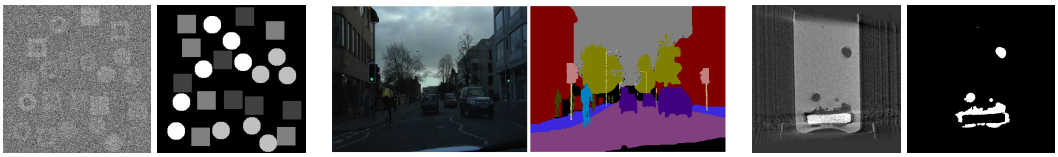

Figure 2: Example input and target images of the (Left) Circle-Square (CS), (middle) CamVid, (right) real-world dynamic CT datasets.

Simulated Circle-Square (CS) dataset: We used a simulated high-noise 5-class segmentation dataset containing 256×256 images of randomly placed squares and circles (CS dataset) (Pelt & Sethian, 2018) (see Figure 2). The objects were assigned a random grey value and Gaussian noise was added to the images. In total, we generated 1000 training, 250 validation, and 100 test images. Experimental results on the CS dataset are quantified using global accuracy, i.e., the ratio of correctly classified pixels, regardless of class, to the total number of pixels. As a starting point for pruning, we trained a 100-layer MS-D network which obtained a global accuracy of 97.5%. The trained ResNet50 and U-Net4 networks obtained 95.8%, and 97.6% respectively.

CamVid: The Cambridge-driving Labeled Video Database (CamVid) (Brostow et al., 2008; 2009) is a collection of videos with labels, captured from the perspective of a driving automobile. In total, 700 labeled frames are split into 367 training, 100 validation, and 233 test images. As there are few training images, we combined the training and validation datasets and trained for a fixed 500 epochs. Similar to other papers that apply CNNs to CamVid (Badrinarayanan et al., 2017; Paszke et al., 2017), we use 11 classes, and a single class representing unlabeled pixels (see Figure 2).

We used median frequency balancing (Eigen & Fergus, 2015) to balance classes for training, and set the unlabeled class weights to zero. During training, we used data augmentation by cropping and (horizontally) flipping input images. As in Paszke et al. (2017), experimental results on the CamVid dataset are quantified using mean Intersection-over-Union (mIoU). As a starting point for pruning, we trained an MS-D network with 150 layers with an mIoU of 0.52, a ResNet50 with an mIoU of 0.71, and a U-Net4 with an mIoU of 0.65.

Real-world dynamic CT dataset: The real-time dynamic X-ray CT dataset contains images of a dissolving tablet suspended in gel (Coban & Lucka, 2019; Coban et al., 2020). The bubbles are to be segmented within a glass container filled with gel (Schoonhoven et al., 2020) (see Figure 2). The dataset consists of 512×512 images, split into 9216 training images, 2048 validation images, and 1536 test images. As in Schoonhoven et al. (2020), we use the F1-score because the large amount of background pixels make global accuracy an unsuitable metric. As a starting point for pruning, we trained an MS-D network with 100 layers with an F1-score of 0.88, a ResNet50 with an F1-score of 0.85, and a U-Net4 with an F1-score of 0.88.

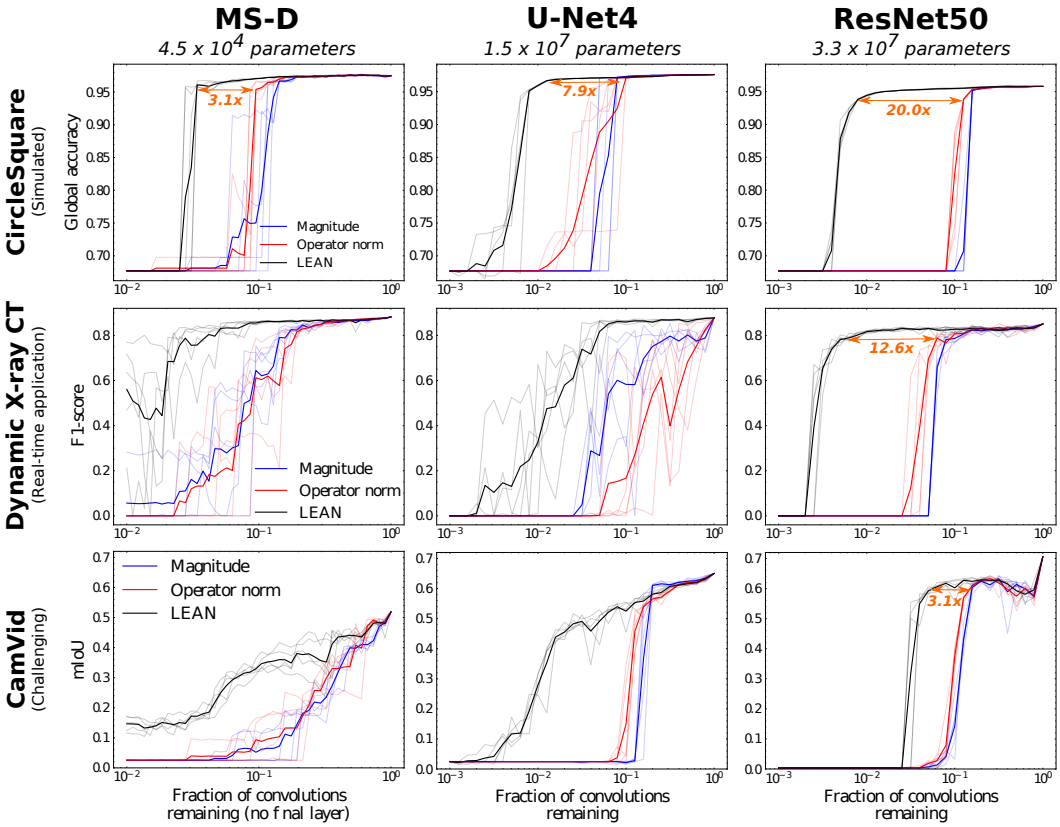

Figure 3: Comparison of structured magnitude, operator norm, and LEAN pruning on three datasets (rows), ordered by difficulty. The pruning methods are applied to MS-D, U-Net4, ResNet50 network architectures, ordered by number of parameters. The base model is pruned to a ratio of 1% (MS-D) or 0.1% (ResNet50, U-Net4) for all datasets. This is repeated five times (translucent lines) and the average is taken (solid lines). In selected panels, the relative number of remaining convolutions is shown at equal accuracy, comparing LEAN to structured pruning (orange).

6 Results

6.1 Experimental results for severe pruning

The results of the pruning experiments are displayed in Figure 3, showing that LEAN pruning at similar accuracy obtains networks with a lower pruning ratio than structured pruning. The pruning ratio that can be achieved without significant loss of accuracy depends on the network architecture and the complexity of the dataset.

In the CS dataset results, we notice a drop-off point where performance decreased significantly for all networks (highlighted by orange arrows). On average over 5 runs of pruning, LEAN achieved a 3.4%, 0.79%, and 0.79% pruning ratio for MS-D, U-Net4, and ResNet50, with an average accuracy reduction of 1.4%, 2.5%, and 2.1% respectively. In all cases, LEAN performs significantly better than the structured pruning methods.

On the dynamic X-ray CT dataset, we notice large fluctuations in F1 for MS-D and U-Net4. This may be due to the F1-score which is defined as a reciprocal. In addition, on U-Net4, LEAN performs better than the structured pruning methods right from the start. For ResNet50, a drop-off point is again observed, which occurs at a significantly lower pruning ratio for LEAN than for the structured pruning methods. On average over 5 runs, LEAN achieved a 5.1%, 6.3%, and 1.0% pruning ratio for MS-D, U-Net4, and ResNet50, with an average F1 drop of 5.7%, 3.3%, and 4.1% respectively.

On the CamVid dataset, we observe a declining mIoU for MS-D and U-Net4 as pruning progresses for this challenging dataset. On ResNet50 we notice an initial drop in mIoU, but subsequent pruning

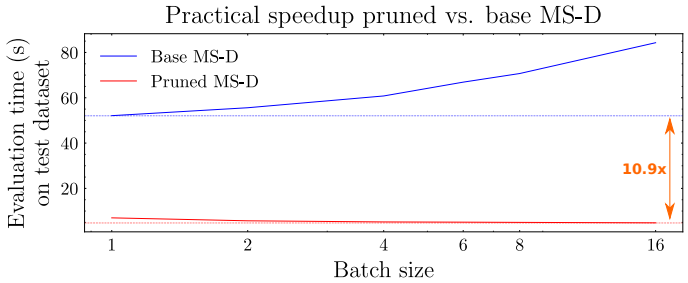

Figure 4: Practically realized speedup of pruned MS-D networks evaluated on the test dataset.

steps increase the performance initially. These observations could indicate that 5 epochs of retraining are not sufficient for the CamVid dataset to recover performance. Interestingly for U-Net4, we notice that for a few pruning ratios, structured magnitude pruning appears to perform slightly better than LEAN. After the initial drop in mIoU, we notice again a later drop-off pruning ratio on ResNet50. LEAN achieves a 6.3% pruning ratio with an average mIoU reduction of 14.3% on ResNet50, whereas structured operator pruning dropped-off at a 20.0% pruning ratio.

6.2 Speedup real-world dynamic X-ray CT segmentation

In addition to measuring the achievable pruning ratios, we measure the practically realized wall-time speedup. We tested this on the dynamic X-ray CT dataset for which a speedup has immediate benefits in practice. Existing pruning support in PyTorch only masks pruned filters, thereby not creating a faster network. Therefore, we implemented a custom MS-D model which loads only the unpruned filters. During the experiments, an MS-D network pruned to a ratio of 2.5% (40-fold reduction) with LEAN achieved an F1-score of 0.83 (drop of 5.4%). This network was 10.9 times faster than the unpruned network in practice. The speed of evaluating the entire test set is impacted by the batch size, which we take into account as shown in Figure 4.

7 Conclusion

In this paper, we have introduced a novel pruning method (LEAN) for CNN pruning by extracting the highest value paths of operators. We incorporate existing graph algorithms and computationally efficient methods for determining the operator norm. We show that LEAN pruning permits removing significantly more operators while retaining better network accuracy than structured pruning across different norms. Our results show that LEAN pruning can increase the speed of the network, both in theoretical speedup (FLOPs reduction) and in practice. In conclusion, LEAN enables severe pruning of CNNs while maintaining a high accuracy, by effectively exploiting the interdependency of network operations.

Future work could be split along two lines. On one hand, there are more CNN operators for which methods to compute their operator norms could be developed. Notably, we have mostly disregarded non-linear operators in this work. Perhaps LEAN could be further improved if we incorporate more operators, such as max-pooling. In combination with taking into account the network structure of multi-layer perceptrons, this could open up the possibility of extending LEAN to other image analysis tasks, such as classification. On the other hand, new graph theoretic approaches for extracting meaningful paths from the graph could be explored. Such algorithms are already abundant in the field of graph theory, and could quite readily be carried over to neural network pruning research.

8 Reproducibility Statement

In order to aid reproducability the authors have published (Python) code for LEAN pruning. There are installation instructions that aim to help users install and run the code on their machines, as well as explanations and documentation for the code and scripts. The scripts are intended to be usable on relatively simple machines (with GPU and CUDA), and to be limited to several minutes run time. First, the code contains a script to generate the CircleSquare (CS) dataset. Second, the

code contains a script to train an MS-D network (a pre-trained network is also supplied) on the CS dataset. The script uses the MS-D network as it is the least computationally expensive to run, but the U-Net4 and ResNet50 models and pruning methods are also supplied. Third, the code contains a script to run LEAN pruning, and the two structured pruning methods, on the trained MS-D network for the CS dataset (example results are also supplied). This way, users can check the experimental results, and experiment themselves with the LEAN method. In addition, the code contains several testing procedures that aim to verify the correctness of the function in the code. For example, the Fourier-based method of computing the operator norm is tested against the power method for random convolutions and strides, and against explicitly computing the SVD of the matrix that is associated with a convolution. Another example is that these tests check whether the pruning methods actually prune the correct number of convolutions. The code is submitted as .zip file in the supplementary materials section, with README instructions.

REFERENCES

S. Abid, F. Fnaiech, and M. Najim. A new neural network pruning method based on the singular value decomposition and the weight initialisation. In *2002 11th European Signal Processing Conference*, pp. 1–4, 2002. 2

Indriani P. Astono, James S. Welsh, Stephan Chalup, and Peter Greer. Optimisation of 2D U-Net Model Components for Automatic Prostate Segmentation on MRI. *Applied Sciences*, 10 (7), 2020. ISSN 2076-3417. doi: 10.3390/app10072601. URL https://www.mdpi.com/2076-3417/10/7/2601. 7

Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. SegNet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017. 7

Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Guttag. What is the state of neural network pruning? *Proceedings of Machine Learning and Systems*, 2020. 1, 2, 3

Gabriel J Brostow, Jamie Shotton, Julien Fauqueur, and Roberto Cipolla. Segmentation and recognition using structure from motion point clouds. In *European conference on computer vision (ECCV)*, pp. 44–57. Springer, 2008. 7

Gabriel J Brostow, Julien Fauqueur, and Roberto Cipolla. Semantic object classes in video: A high-definition ground truth database. *Pattern Recognition Letters*, 30(2):88–97, 2009. 7

Sophia Bethany Coban and Felix Lucka. Dynamic 3D X-ray micro-CT data of a tablet dissolution in a water-based gel, 2019. URL https://doi.org/10.5281/zenodo.3610187. 7

Sophia Bethany Coban, Felix Lucka, Willem Jan Palenstijn, Denis Van Loo, and Kees Joost Batenburg. Explorative Imaging and Its Implementation at the FleX-ray Laboratory. *Journal of Imaging*, 6(4):18, 2020. URL "https://doi.org/10.3390/jimaging6040018". 7

Emily L Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *Advances in neural information processing systems*, pp. 1269–1277, 2014. 2

S. Dittmer, E. J., and P. Maass. Singular values for relu layers. *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–12, 2019. 2

Xuanyi Dong and Yi Yang. Network pruning via transformable architecture search. In *Advances in Neural Information Processing Systems*, pp. 760–771, 2019. 2

David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE international conference on computer vision*, pp. 2650–2658, 2015. 7

Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT Press, 2016. 3

Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, pp. 1135–1143, 2015. 3

Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *International Conference on Learning Representations (ICLR)*, 2016. 1

Babak Hassibi, David G Stork, and Gregory J Wolff. Optimal brain surgeon and general network pruning. In *IEEE international conference on neural networks*, pp. 293–299. IEEE, 1993. 2

K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016. 5, 6, 7

Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2

Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 1389–1397, 2017. 2

Allard A. Hendriksen. Mixed-scale dense networks for pytorch. https://github.com/ahendriksen/msd_pytorch, 2020. 6

Zehao Huang and Naiyan Wang. Data-driven sparse structure selection for deep neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 304–320, 2018. 2

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *International Conference on Machine Learning (ICML)*, 2015. 3

Anil K. Jain. *Fundamentals of Digital Image Processing*. Prentice-Hall, Inc., USA, 1989. ISBN 0133361659. 13

E. D. Karnin. A simple procedure for pruning back-propagation trained neural networks. *IEEE Transactions on Neural Networks*, 1(2):239–242, 1990. 1, 2

Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*, 2014. 6

Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Advances in neural information processing systems*, pp. 598–605, 1990. 2

Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015. 1

Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient ConvNets. In *5th International Conference on Learning Representations, ICLR 2017*, 2017. 1

Ji Lin, Yongming Rao, Jiwen Lu, and Jie Zhou. Runtime neural pruning. In *Advances in neural information processing systems*, pp. 2181–2191, 2017. 2

Shaohui Lin, Rongrong Ji, Chenqian Yan, Baochang Zhang, Liujuan Cao, Qixiang Ye, Feiyue Huang, and David Doermann. Towards optimal structured CNN pruning via generative adversarial learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2

Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2736–2744, 2017. 2

Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. *International Conference on Learning Representations (ICLR)*, 2019. 2

Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, 2015. 14

Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*, pp. 5058–5066, 2017. 1, 2

Carl D Meyer. *Matrix analysis and applied linear algebra*, volume 71. Siam, 2000. 4

Alexandre Milesi. Unet: semantic segmentation with pytorch. https://github.com/milesial/Pytorch-UNet, 2020. 7

Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. *International Conference on Learning Representations (ICLR)*, 2017. 2

Pavlo Molchanov, Arun Mallya, Stephen Tyree, Iuri Frosio, and Jan Kautz. Importance estimation for neural network pruning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2

Michael C Mozer and Paul Smolensky. Skeletonization: A technique for trimming the fat from a network via relevance assessment. In D. S. Touretzky (ed.), *Advances in Neural Information Processing Systems 1*, pp. 107–115. Morgan-Kaufmann, 1989. 1, 2

Jongsoo Park, Sheng Li, Wei Wen, Ping Tak Peter Tang, Hai Li, Yiran Chen, and Pradeep Dubey. Faster CNNs with direct sparse convolutions and guided pruning. *International Conference on Learning Representations (ICLR)*, 2017. 1

Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. ENet: A deep neural network architecture for real-time semantic segmentation. *International Conference on Learning Representations (ICLR)*, 2017. 7

Daniël M. Pelt and James A. Sethian. A mixed-scale dense convolutional neural network for image analysis. *Proceedings of the National Academy of Sciences*, 115(2):254–259, 2018. ISSN 0027-8424. doi: 10.1073/pnas.1715832114. URL http://www.pnas.org/content/115/2/254. 6, 7

Alex Renda, Jonathan Frankle, and Michael Carbin. Comparing rewinding and fine-tuning in neural network pruning. *International Conference on Learning Representations (ICLR)*, 2020. 2, 3

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241, 2015. 1, 6

R. Schoonhoven, J. W. Buurlage, D. M. Pelt, and K. J. Batenburg. Real-time segmentation for tomographic imaging. In *2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*, pp. 1–6, 2020. doi: 10.1109/MLSP49062.2020.9231642. 7

Robert Sedgewick and Kevin Wayne. *Algorithms*. Addison-wesley professional, 4 edition, 2011. 5

Hanie Sedghi, Vineet Gupta, and Philip M Long. The singular values of convolutional layers. *International Conference on Learning Representations (ICLR)*, 2019. 2, 4, 13, 14

Chunwei Tian, Yong Xu, and Wangmeng Zuo. Image denoising using deep CNN with batch renormalization. *Neural Networks*, 121:461–473, 2020. 1

Nan Yang, Lukas von Stumberg, Rui Wang, and Daniel Cremers. D3vo: Deep depth, deep pose and deep uncertainty for monocular visual odometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1281–1292, 2020. 1

T. Yang, Y. Chen, and V. Sze. Designing energy-efficient convolutional neural networks using energy-aware pruning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6071–6079, 2017. 1

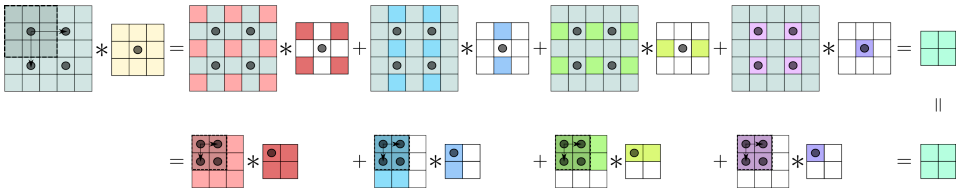

Figure 5: The output of a strided convolution of stride 2 can be obtained by summing the result of regular convolutions. First, split the image, and the original convolution filter, into 4 smaller parts (corresponding sections given in color) with white entries representing zeroes. Next, sum the outputs of 4 regular convolutions pixel by pixel. The diagram outlines the procedure for a 5×5 image and 3×3 kernel. The convolution is performed without padding and produces a 2×2 output image. The dots in the image represent the positions of the dot in the filter, during computation of the convolution.

Jianbo Ye, Xin Lu, Zhe Lin, and James Z Wang. Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. *International Conference on Learning Representations (ICLR)*, 2018. 2

Chenglong Zhao, Bingbing Ni, Jian Zhang, Qiwei Zhao, Wenjun Zhang, and Qi Tian. Variational convolutional neural network pruning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2

A Appendix: Operator norm calculation

In addition to the structure of nodes and edges, the pruning graph contains edge we assign edges in the pruning graph weights related to the operator norms of the CNN operators. In this Appendix we will discuss how the operator norms for various common CNN operators are computed.

A.1 Operator norm of convolutions

As mentioned in Section 3.3, the spectral norm of a convolution with filter h can be determined by the singular values of the associated matrix H. A substantially more efficient method than constructing H is to make use of Fourier transforms Sedghi et al. (2019). The matrix H is diagonalized by the two-dimensional unitary Discrete Fourier Transform (DFT), and the eigenvalues of H are given by the two-dimensional DFT of h^+ (see Section 5 of Jain (1989)). For an image of size $n \times n$, h^+ is the filter padded with zeroes to size $n \times n$. Practically, the singular values of H are the magnitudes of the complex entries of $F_{2D}(h^+)$. Here F_{2D} is the 2D Fourier transform.

$$\sigma_{max}(H) = \max\left\{\left|F_{2D}(h^+)\right|\right\}$$

Note that the Fourier-based method used in this work assumes a doubly-blocked circulant definition of H, which corresponds to periodic boundary conditions as opposed to zero padding. When n is chosen sufficiently large, the difference in operator norm is negligible.

A.2 Downsampling: operator norm for strided convolutions

The operator norm of a strided convolution can be computed by observing that it is equal to the sum of several regular convolutions on smaller input images. Consider a single parameter for a stride 2 convolution kernel. Because of the stride, it is multiplied only with every other pixel (both horizontally and vertically). Similarly, parameters in the kernel which are exactly 2 apart are multiplied with the same pixels. In Figure 5 we outline the procedure for a 3×3 kernel, with stride 2. Here, there are 4 different convolution positions before the pattern repeats itself; every combination of odd and even numbered indices for the horizontal and vertical axes (each pattern is marked by a different color). In general, the same properties hold for other values of s. We will use these properties to compute the operator norm.

Denote the partitioned convolution kernels as $h^{[i]}$, and the partitioned input images as $X^{[i]}$. Both $h^{[i]}$ and $X^{[i]}$ are zero-padded to the correct size (illustrated as white cells in Figure 5). Then we have that

$$h *_2 X = \sum_{i=1}^{4} h^{[i]} * X^{[i]} \tag{4}$$

where $*_2$ denotes a stride 2 convolution. As $h^{[i]} * X^{[i]}$ is a regular convolution, we can determine the singular values of $h^{[i]}$ using Fourier transforms as described above.

In general, Equation (4) is analogous to the equation of a convolutional layer with 4 input channels and 1 output channel. The operator norm of such a convolutional layer is described in Sedghi et al. (2019). There, they collect the singular values of the convolutional layer into a tensor $P \in \mathbb{R}^{4\times 1\times n\times n}$, with

$$P_{c_{in},c_{out},i,j} = F_{2D}(h^{[c_{in}]+})_{i,j}.$$

According to Theorem 6 of Sedghi et al. (2019), the spectral norm of the convolutional layer equals the maximum of the singular values of the 4×1 matrices $P_{:,:,i,j}$. Since the matrices are single-column, their singular value equals their L_2-norm. Therefore, the spectral norm of h equals

$$\|h\| = \max_{i,j} \left\{ \sum_{c_{in}} P^2_{c_{in},:,i,j} \right\}$$

In the case of 2×2 kernels, and stride 2, a relatively simple calculation can be performed. In this case, each $h^{[i]}$ contains just one non-zero parameter, and the convolution $h^{[i]} * X^{[i]}$ is in fact a scalar multiplication. So for a 2×2 kernel with parameters h_{ij}, we get that $\sigma_{max}(h) = \sqrt{\sum_{i,j} h^2_{i,j}}$.

A.3 Upsampling: operator norm for transposed convolutions

The operator norm of a transposed convolution can be computed in a similar fashion as the strided convolution. The matrix associated with a transposed convolution of stride s is exactly the transpose of the matrix for a stride s convolution Long et al. (2015). In general we have for any linear map A between Euclidean vector spaces that $\|A\| = \left\|A^T\right\|$. Therefore, for a given transposed convolution kernel h, the operator norm is equal to that as if h were applied as a strided convolution. As such, we can utilize the procedure outlined in Section A.2.

A.4 Operator norm for batch normalization and pooling operators

As operator norm for a batch normalization channel B we take the multiplicative scaling it applies to images (see Section 3.1)

$$\|B\| = \left| \frac{\gamma}{\sqrt{\sigma_B^2 + \epsilon}} \right|.$$

Note that we do not include the additive β term in the computation of the operator norm, but rather interpret it as a separate bias term.

In addition to convolutional operators, there are other CNN operators to which we can attribute an operator norm. Average pooling can be seen as a convolution operation. *Average pooling* with stride s is the same as a stride s convolution, with an $s \times s$ kernel with values $1/s^2$ everywhere. As such, we can calculate the operator norm efficiently as outlined in Section A.2.