arXiv:1906.03808v1 [cs.LG] 10 Jun 2019

A Closed-Form Learned Pooling for Deep Classification Networks

Vighnesh Birodkar *
vighneshb@google.com

Hossein Mobahi
hmobahi@google.com

Dilip Krishnan
dilipkay@google.com

Samy Bengio
bengio@google.com

Abstract

In modern computer vision tasks, convolutional neural networks (CNNs) are indispensable for image classification tasks due to their efficiency and effectiveness. Part of their superiority compared to other architectures, comes from the fact that a single, local filter is shared across the entire image. However, there are scenarios where we may need to treat spatial locations in non-uniform manner. We see this in nature when considering how humans have evolved foveation to process different areas in their field of vision with varying levels of detail. In this paper we propose a way to enable CNNs to learn different pooling weights for each pixel location. We do so by introducing an extended definition of a pooling operator. This operator can learn a strict super-set of what can be learned by average pooling or convolutions. It has the benefit of being shared across feature maps and can be encouraged to be local or diffuse depending on the data. We show that for fixed network weights, our pooling operator can be computed in closed-form by spectral decomposition of matrices associated with class separability. Through experiments, we show that this operator benefits generalization for ResNets and CNNs on the CIFAR-10, CIFAR-100 and SVHN datasets and improves robustness to geometric corruptions and perturbations on the CIFAR-10-C and CIFAR-10-P test sets.

1 Introduction

Convolutional Neural Networks (CNNs) have revolutionized the field of computer vision (Krizhevsky et al., 2012a; He et al., 2016b, 2017). Their success (compared to fully connected networks) is often attributed to their weight sharing in form of a convolution, which reduces the number of learnable parameters (Krizhevsky et al., 2012b). In addition, the "shift invariance" property of convolution has been believed to be crucial for improved generalization in vision tasks (Fukushima and Miyake, 1982) (although some modifications may be required (Azulay and Weiss, 2018; Zhang, 2019)). Shift invariance, while crucial for handling translation in images, is a very limited form of real-world geometric transformations. For instance, convolutional representations are not invariant or equivariant to other basic transforms such as image rotation and scaling (Azulay and Weiss, 2018).

There have been recent attempts to incorporate additional forms of invariances, such as rotation, reflection, and scaling. (Sifre and Mallat, 2013; Bruna and Mallat, 2013; Esteves et al., 2018a; Kanazawa

---

[*] Work done as a Google AI Resident

et al., 2014; Worrall et al., 2017). However, these methods engineer the invariance into networks; requiring a-priori identification of invariance types of interest. In this work, our goal is to achieve invariance using a data-driven approach to handle geometric transforms that may not fall into the categories mentioned above.

One way to achieve invariance or equivariance for transforms beyond translation is non-uniform sub-sampling of the image (to be fed as input to a convolutional layer). For example, log-polar sampling of an image results in a new image, where rotation and scaling in the original image become equivalent to a translation in the resulted image (Esteves et al., 2018b). This suggests that by ***adapting the pooling operator***, one can build representations that are better suited for variety of geometric transforms.

Non-uniform sampling is also the chosen scheme by nature; foveal vision implements a spatially-varying sampling similar to log-polar transform (Larson and Loschky, 2009). Central and peripheral regions are sampled at different frequencies and both contribute to efficient and effective human vision. In addition, it is known that non-uniform sampling of image can facilitate image registration when geometric transforms are beyond translation (Mobahi et al., 2012). Our results with learned pooling operator confirms the advantage of spatially varying pooling. For example, Figure 1a shows the response map of a learned pooling operator on SVHN dataset. The operator places more weights on the center pixel to take advantage of the fact that SVHN digits are mostly in the center. Contrast this with the operator learned for a CIFAR-100 model which places weights all across the spatial field (see Supplementary). The form of the learned pooling operator also affects the pooled feature maps. In Figure 1b, the pooled feature maps are more clustered around the mean feature map of each class, compared to the feature maps produced by a regular CNN. This results in better separability of classes and better generalization as seen in Table 1.

In addition to adapting to the geometric transforms present in the data, and hence improving generalization, our learned pooling operator helps with robustness of the model. It has been observed that small geometric transforms in the image can result in prediction errors in existing deep models, which can be traced to the pooling operator (Azulay and Weiss, 2018; Zhang, 2019); max pooling results in aliasing effects in the representation. While average pooling can prevent the aliasing issues (because it acts as a low-pass filter), the blurring causes loss of information and hence inferior classification performance to max pooling. However, by adapting the pooling operator to the data, in way that it can provide more class separability, relevant information are automatically picked up. In fact, our experiments show that the learned pooling can outperform the naive uniform down-sampling scheme that is used in most state-of-the-art models (strided pooling) and yet is robust to geometric perturbations on robustness benchmark datasets (CIFAR-10-C and CIFAR-10-P (Hendrycks and Dietterich, 2019)).

2 Related Work

Convolutional Neural Networks (CNNs) rely on pooling or sub-sampling to reduce the size of the hidden representation. This is known to have important implications towards the kinds of invariances and generalization abilities of a network (Cohen and Shashua, 2016). Earlier architectures have relied on average-pooling (LeCun et al., 1998) and max-pooling (Krizhevsky et al., 2012a), whereas modern ones learn parameters of pooling through strided convolutions (He et al., 2016a).

The history of pooling in computer vision however goes past CNNs popularity. For instance, (Boureau et al., 2011) combines SIFT with pooling separately over learned clusters of features. (Malinowski and Fritz, 2013) learns pooling parameters of all spatial lower-level features through a fully connected network, whereas (Gong et al., 2014) learns pooling separately at each scale using VLAD (Jégou et al.,

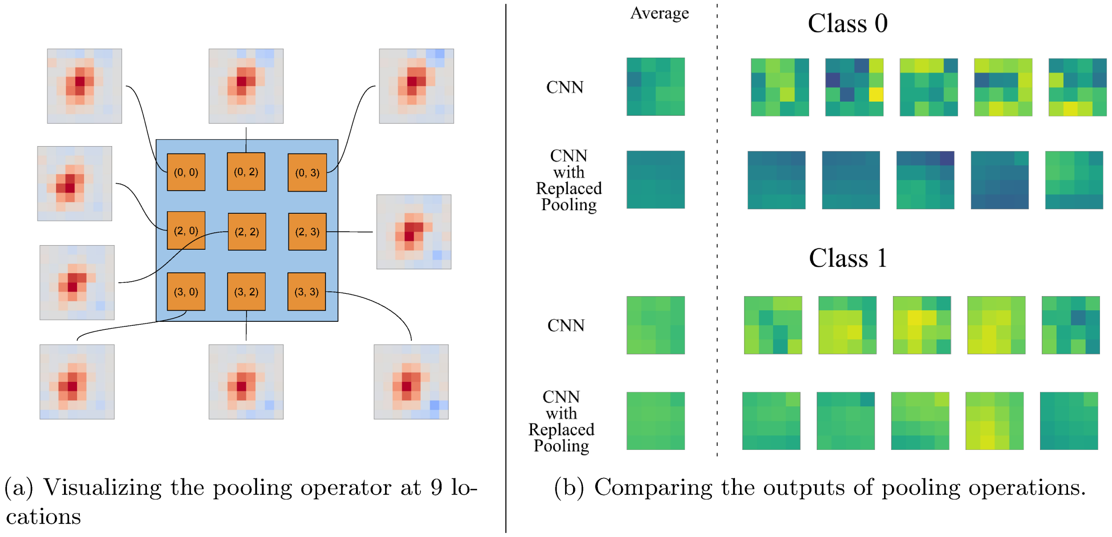

(a) Visualizing the pooling operator at 9 locations

(b) Comparing the outputs of pooling operations.

Figure 1: One the left we visualize the learned pooling map for a CNN model on the reduced SVHN dataset. Each heat map corresponds to one of the 9 evenly spaced locations in the output feature map. For each pixel in the output, we can see where the algorithm chooses to put a positive(red) or negative(blue) weight over the input. On the right, we compare the output of the default pooling used (strided convolution) against the output of our pooling operator. We randomly select a channel and display the average value of that channel per class along with 5 other values closest to the average. We note that with our pooling operator, the per-class feature maps are closer to the per-class average feature maps as compared to the regular CNN model.

2010). (Girshick et al., 2015) defines distance transform pooling as part of deformable part models using a quadratic function of distance from the center and use a latent SVM (Felzenszwalb et al., 2009) to learn it on top of a pre-trained CNN. (Li et al., 2015) uses pooling at different scales in spatial dimensions and also performs pooling on the color channels while aggregating using the max operator.

Pooling has also been used to aggregate varying input sizes into a fixed size representation. (Passalis and Tefas, 2017) uses a fixed number of RBF neurons on top of a regular CNN to output a fixed size representation irrespective of the input image size. (Zhou et al., 2017) uses a specially designed pooling function for a multi-instance learning setting to output tags for a video from tags predicted for each frame. (Miech et al., 2017) uses NetVLAD defined in (Arandjelovic et al., 2016) and approximations of Bag of Words and Fischer vector encoding to aggregate features across time for learning video classification.

Recently, there have been some attempts to to learn parameters of local pooling operations of CNNs in an end-to-end fashion though gradient descent. (Sun et al., 2017) propose learning a local pooling operator, one per each channel and train a deep-neural network for classification. (Saeedan et al., 2018) try to preserve small details in the input while pooling and introduce two new parameters per input feature map to control which details are preserved. (Lee et al., 2016) experiment with learned and fixed combinations of average/max pooling and also suggest organizing the outputs of multiple local filters in the form of a binary tree to learn the parameters of mixing them. Although these approaches learn pooling parameters from data, the pooling operator is limited to spatially uniform; the same sampling scheme is used to pool each output pixel. As we discuss in Section 1, spatially varying pooling is necessary to learn efficient and robust representation for transformations other than translation.

3 Method

3.1 Notation

We first formalize the definition of linear pooling. Let $\mathcal{N}_k \triangleq \{1, 2, 3, \dots, k\}$. Given a spatial domain $\Omega \triangleq \mathcal{N}_I \times \mathcal{N}_J$ and set of intensity values Γ, a feature map of depth C is a map $f : \Omega \times \mathcal{N}_C \to \Gamma$. We can represent a feature map f in matrix form:

$$
\begin{aligned}
\boldsymbol{F}_{c\,I\times J} &\triangleq \begin{pmatrix} f(1,1,c) & f(1,2,c) & \dots & f(1,J,c) \\ f(2,1,c) & f(2,2,c) & \dots & f(2,J,c) \\ \vdots & \vdots & \vdots & \vdots \\ f(I,1,c) & f(I,2,c) & \dots & f(I,J,c) \end{pmatrix} && (1)\\
\boldsymbol{f}_{c\,N\times 1} &\triangleq \operatorname{vec}(\boldsymbol{F}_c) && (2)\\
\boldsymbol{X}_{N\times C} &\triangleq \big(\boldsymbol{f}_1 \;\mid\; \boldsymbol{f}_2 \;\mid\; \dots \;\mid\; \boldsymbol{f}_C \big)\,. && (3)
\end{aligned}
$$

where N denotes the domain size $N \triangleq |\Omega| = I \times J$ and vec converts a matrix into a column vector by concatenating the columns of the matrix.

We define ***linear pooling*** as the operator $\boldsymbol{P}_{M\times N}$ which maps $\boldsymbol{X}_{N\times C}$ into another feature map $\boldsymbol{Z}_{M\times C}$ where $1 \le M \le N$. That is, the operator may shrink the input spatially but maintains the number of channels. Formally, the output of the linear operator is an element in $\omega \times \mathcal{N}_C$, where $\omega \triangleq \mathcal{N}_{I'} \times \mathcal{N}_{J'}$ and $M \triangleq |\omega| = I' \times J'$. Note that the operator is applied to each channel (column) of the input matrix to generate the corresponding channel (column) in the output matrix.

$$\boldsymbol{Z}_{M\times C} \triangleq \boldsymbol{P}_{M\times N}\, \boldsymbol{X}_{N\times C} \qquad (4)$$

3.2 Formulation

Obviously, ***average pooling*** is seen as a special case for the linear operator $\boldsymbol{P}$ when the entries of the operator are set in a specific way. However, one may wonder if, within the space of all linear operators, there could be better choices than average pooling. Trivially, the answer is data dependent and hence we let the data itself discover the operator that suits the task in hand. In the classification setting, a good operator should help with classification. One possible way to quantify helping with classification is to improve separability of the classes, as explained below.

To simplify the formulation, we focus on finding each row of $\boldsymbol{P}$ separately. Let $\boldsymbol{z}_m$ be the m'th row of $\boldsymbol{Z}$, arranged as a column vector. Similarly, let $\boldsymbol{p}_m$ be the m'th row of $\boldsymbol{P}$, also arranged as a column vector. Then the pooling identity in (4) can be equivalently expressed using vectors as:

$$\boldsymbol{z}_{m\,C\times 1} \triangleq {\boldsymbol{X}^T}_{C\times N}\, \boldsymbol{p}_{m\,N\times 1} \quad , \quad m = 1, \dots, M\,. \qquad (5)$$

To reduce mathematical clutter, we drop the index m from $\boldsymbol{z}_m$ and $\boldsymbol{p}_m$. The reader should remember that the following result needs to be applied for each choice of $m = 1, \dots, M$ separately. Hence with abuse of notation we proceed as:

$$\boldsymbol{z}_{C\times 1} \triangleq {\boldsymbol{X}^T}_{C\times N}\, \boldsymbol{p}_{N\times 1}\,. \qquad (6)$$

To define separability, we require a training set. Consider a set of K feature maps $\boldsymbol{X}_1, \dots, \boldsymbol{X}_K$, whose elements are associated with a label $y_k \in \{1, \dots, Q\}$, with Q being the number of classes in the dataset.

We define the following total and per class average quantities:

$$\overline{\boldsymbol{z}}_{C\times 1} \triangleq \frac{1}{K}\sum_{k} \boldsymbol{z}_{k\,C\times 1} \quad , \quad \overline{\boldsymbol{z}}_{q\,C\times 1} \triangleq \frac{1}{K_q}\sum_{k|y_k=q} \boldsymbol{z}_{k\,C\times 1} \quad , \quad K_q \triangleq \sum_{k|y_k=q} 1 \,. \tag{7}$$

Inspired by Linear Discriminant Analysis (LDA), we quantify separability of the classification as the ratio of between-class scatter S_b and within-class scatter S_w.

$$S_b \triangleq \frac{1}{Q}\sum_{q} \|\overline{\boldsymbol{z}}_{q\,C\times 1} - \overline{\boldsymbol{z}}_{C\times 1}\|^2 \quad , \quad S_w \triangleq \frac{1}{Q}\sum_{q}\frac{1}{K_q}\sum_{k|y_k=q} \|\boldsymbol{z}_{k\,C\times 1} - \overline{\boldsymbol{z}}_{q\,C\times 1}\|^2 \,. \tag{8}$$

To achieve a good representation for classification, we aim to improve separability of the data points by maximizing the ratio:

$$\boldsymbol{p}^* \triangleq \arg\max_{\boldsymbol{p}} \frac{S_b}{S_w} \,. \tag{9}$$

Plugging definitions from (7) and (6) into the above objective function yields:

$$\frac{S_b}{S_w} = \frac{\boldsymbol{p}^T{}_{1\times N} \sum_q \big((\overline{\boldsymbol{X}}_q - \overline{\boldsymbol{X}})_{N\times C}(\overline{\boldsymbol{X}}_q - \overline{\boldsymbol{X}})^T{}_{C\times N}\big)\boldsymbol{p}_{N\times 1}}{\boldsymbol{p}^T{}_{1\times N}\big(\sum_q \frac{1}{K_q}\sum_{k|y_k=q}(\boldsymbol{X}_k - \overline{\boldsymbol{X}}_q)_{N\times C}(\boldsymbol{X}_k - \overline{\boldsymbol{X}}_q)^T{}_{C\times N}\big)\boldsymbol{p}_{N\times 1}} \,, \tag{10}$$

where $\overline{\boldsymbol{X}}, \overline{\boldsymbol{X}_{\boldsymbol{q}}}$ are defined based on $\boldsymbol{X}_k$'s in a similar way done for $\boldsymbol{z}$ in (7). For brevity, define:

$$\boldsymbol{A}_{N\times N} \quad \triangleq \quad \sum_{q}(\overline{\boldsymbol{X}}_q - \overline{\boldsymbol{X}})_{N\times C}(\overline{\boldsymbol{X}}_q - \overline{\boldsymbol{X}})^T{}_{C\times N} \tag{11}$$

$$\boldsymbol{B}_{N\times N} \quad \triangleq \quad \sum_{q}\frac{1}{K_q}\sum_{k|y_k=q}(\boldsymbol{X}_k - \overline{\boldsymbol{X}}_q)_{N\times C}(\boldsymbol{X}_k - \overline{\boldsymbol{X}}_q)^T{}_{C\times N} \,. \tag{12}$$

This way our goal is to maximize separability:

$$\boldsymbol{p}^*_{N\times 1} = \arg\max_{\boldsymbol{p}} \frac{\boldsymbol{p}^T{}_{1\times N}\boldsymbol{A}_{N\times N}\boldsymbol{p}_{N\times 1}}{\boldsymbol{p}^T{}_{1\times N}\boldsymbol{B}_{N\times N}\boldsymbol{p}_{N\times 1}} \,. \tag{13}$$

3.3 Closed-Form Solution

The solution to (13) is ill-posed; if some $\boldsymbol{p}^\dagger$ is a solution, then so is $a\,\boldsymbol{p}^\dagger$ for any $a \neq 0$. To avoid such freedom of scale, we anchor S_w to a fixed value, e.g. 1 and then solve:

$$\boldsymbol{p}^* = \arg\max_{\boldsymbol{p}} \boldsymbol{p}^T{}_{1\times N}\boldsymbol{A}_{N\times N}\boldsymbol{p}_{N\times 1} \quad \text{s.t.} \quad \boldsymbol{p}^T{}_{1\times N}\boldsymbol{B}_{N\times N}\boldsymbol{p}_{N\times 1} = 1 \,. \tag{14}$$

In addition, we wish to keep $\boldsymbol{p}$ ***localized*** so that the operator $\boldsymbol{P}$ respects the topology of the space Ω. This is important when the pooled feature maps are to be processed by convolution operators (in the future layers). We can encourage localization by introducing a penalty term of the form $\boldsymbol{p}^T{}_{1\times N}\boldsymbol{C}_{N\times N}\boldsymbol{p}_{N\times 1}$, where $\boldsymbol{C}$ is a diagonal matrix with positive components. How the elements of $\boldsymbol{C}$ are chosen is described in Section 3.4.

Applying localization penalty and anchoring of S_w results in the following optimization:

$$\boldsymbol{p}^* = \arg\max_{\boldsymbol{p}} \boldsymbol{p}^T_{1\times N}\boldsymbol{A}_{N\times N}\boldsymbol{p}_{N\times 1} \; - \; \rho\,\boldsymbol{p}^T\boldsymbol{C}\boldsymbol{p} \quad , \quad \text{s.t. } \boldsymbol{p}^T_{1\times N}\boldsymbol{B}_{N\times N}\boldsymbol{p}_{N\times 1} = 1 \,, \tag{15}$$

where $\rho > 0$ is the penalty coefficient. It can be shown (see the supplementary appendix for proof) that the solution $\boldsymbol{p}^*$ must satisfy the following ***generalized eigenvalue problem*** for some[1] $\alpha > 0$.

$$\boldsymbol{A}_{N\times N}\boldsymbol{p}_{N\times 1} = -\lambda\Big(\boldsymbol{B}_{N\times N} + \alpha\,\boldsymbol{C}_{N\times N}\Big)\boldsymbol{p}_{N\times 1}\,. \tag{16}$$

One way to solve the generalized eigenvalue problem is by matrix inversion. If the matrix on the r.h.s. is invertible, then we have:

$$\Big(\boldsymbol{B}_{N\times N} + \alpha\,\boldsymbol{C}_{N\times N}\Big)^{-1}\boldsymbol{A}_{N\times N}\boldsymbol{p}_{N\times 1} = -\lambda\boldsymbol{p}_{N\times 1}\,, \tag{17}$$

which implies that the optimal $\boldsymbol{p}$ must be an (in fact the leading) eigevector of the following matrix:

$$\boldsymbol{p}^*_{N\times 1} = \texttt{top\_eigvec}\Big((\boldsymbol{B}_{N\times N} + \alpha\,\boldsymbol{C}_{N\times N})^{-1}\boldsymbol{A}_{N\times N}\Big)\,. \tag{18}$$

Since the matrix $\boldsymbol{C}$ is diagonal with positive entries, and the matrix $\boldsymbol{B}$ is positive semi-definite, the matrix $\boldsymbol{C}$ has a regularization effect when computing the inverse of $\boldsymbol{B} + \alpha, \boldsymbol{C}$. Thus we refer to $\boldsymbol{C}$ as a regularization matrix.

3.4 Regularization Matrix $\boldsymbol{C}$

We now explain how the components of $\boldsymbol{C}$ are chosen. For clarity, we temporarily (throughout this subsection) switch from the brief notation $\boldsymbol{p}$ to the full notation $\boldsymbol{p}_m$. We also need to switch from $\boldsymbol{C}$ to $\boldsymbol{C}_m$ accordingly. Note that each component of $\boldsymbol{p}_m$ corresponds to a coordinate $(i,j) \in \Omega$. To show this relationship we use the notation $\texttt{coord}_\Omega(n) = (i,j)$. Similarly, for the space ω, each index m is associated with a coordinate $(i,j) \in \omega$, and the relationship is shown via $\texttt{coord}_\omega(m) = (i,j)$. We penalize the n'th component of $\boldsymbol{p}_m$ (recall each $\boldsymbol{p}_m$ is a vector of size N, thus $1 \leq n \leq N$) by its coordinate distance from that of m. Since in general $N \neq M$, a scale correction needs to be done. This way, the amount of penalty for the n'th component of $\boldsymbol{p}_m$, which is encoded in the diagonal element $c^m_{n,n}$, is set to $c^m_{n,n} \triangleq \|\texttt{coord}_\Omega(n) - s\,\texttt{coord}_\omega(m)\|^2$, where s is the scale factor $s = \frac{I}{I'} = \frac{J}{J'}$. Here $c^m_{i,j}$ refers to the (i,j)'th component of the matrix $\boldsymbol{C}_m$.

In words, this penalty scheme means that if a point in the source feature maps contributes to a point in the destination feature map, where the latter is far from the source point, then that contribution is penalized.

3.5 Algorithm

The resulted procedure is shown in Algorithm 1. Note that the matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ are the same for any m. From the beginning of the algorithm up to line 16 is to compute these matrices. However, the matrix $\boldsymbol{C}$ and the vector $\boldsymbol{p}^*_m$ both depend on m and thus Line 20 to the end loops over m to compute each $\boldsymbol{C}$ and its resulted $\boldsymbol{p}^*_m$.

3.6 Implementation Details

Choice of Norm. The last line of algorithm returns $\frac{\boldsymbol{p}^*_1}{\|\boldsymbol{p}^*_1\|}, \dots, \frac{\boldsymbol{p}^*_M}{\|\boldsymbol{p}^*_M\|}$. We will explore ℓ_1 and ℓ_2 norms in the experiments.

---

[1] It turns out α is proportional to ρ and hence still serves as some penalty coefficient. See the supplementary appendix for details.

Algorithm 1 Learning Pooling Operator.

```
Input: Training pairs $\cup_{k=1}^{K}\{(\boldsymbol{X}_k, y_k)\}$ where $\boldsymbol{X}_k \in \mathbb{R}^{N\times C}$ and $y_k \in \{1, \dots, Q\}$ with $Q$ being number of classes, penalty coefficient $\alpha > 0$, scaling factor $s$.
$\bar{\boldsymbol{S}} \leftarrow \boldsymbol{O}_{N\times C}$, $\bar{\boldsymbol{T}} \leftarrow \boldsymbol{O}_{N\times N}$, $\boldsymbol{B} \leftarrow \boldsymbol{O}_{N\times N}$
for $q \leftarrow 1$ to $Q$ do
    $\boldsymbol{S} \leftarrow \boldsymbol{O}_{N\times C}$, $\boldsymbol{T} \leftarrow \boldsymbol{O}_{N\times N}$
    for $k \leftarrow 1$ to $K_q$ do
        if $y_k = q$ then
            $\boldsymbol{S} \leftarrow \boldsymbol{S} + \boldsymbol{X}_k$
            $\boldsymbol{T} \leftarrow \boldsymbol{T} + \boldsymbol{X}_k \boldsymbol{X}_k^T$
        end if
    end for
    $\bar{\boldsymbol{X}}_q \leftarrow \frac{1}{K_q}\boldsymbol{S}$
    $\bar{\boldsymbol{S}} \leftarrow \bar{\boldsymbol{S}} + \bar{\boldsymbol{X}}_q$
    $\bar{\boldsymbol{T}} \leftarrow \bar{\boldsymbol{T}} + \bar{\boldsymbol{X}}_q \bar{\boldsymbol{X}}_q^T$
    $\boldsymbol{B} \leftarrow \boldsymbol{B} + \frac{1}{K_q}\boldsymbol{T} - \bar{\boldsymbol{X}}_q \bar{\boldsymbol{X}}_q^T$
end for
$\boldsymbol{A} \leftarrow \bar{\boldsymbol{T}} - \frac{1}{Q}\bar{\boldsymbol{S}}\bar{\boldsymbol{S}}^T$
$\boldsymbol{d}_{N\times 2} \leftarrow [\texttt{coord}_\Omega(1), \dots, \texttt{coord}_\Omega(N)]$
$\boldsymbol{e}_{N\times 1} \leftarrow [\|\texttt{coord}_\Omega(1)\|^2, \dots, \|\texttt{coord}_\Omega(N)\|^2]$
$M \leftarrow \frac{N}{s^2}$
for $m \leftarrow 1$ to $M$ do
    $\boldsymbol{c}_{N\times 1} \leftarrow \boldsymbol{e}_{N\times 1} - 2s\boldsymbol{d}_{N\times 2}\texttt{coord}_\omega(m)_{2\times 1} + s^2\|\texttt{coord}_\omega(m)\|^2\mathbf{1}_{N\times 1}$
    $\boldsymbol{C}_{N\times N} \leftarrow \mathrm{diag}(\boldsymbol{c}_{N\times 1})$
    $\boldsymbol{p}^*_{m\,N\times 1} \leftarrow \texttt{top\_eigvec}\Big((\boldsymbol{B}_{N\times N} + \alpha\,\boldsymbol{C}_{N\times N})^{-1}\boldsymbol{A}_{N\times N}\Big)$
end for
return $\frac{\boldsymbol{p}^*_1}{\|\boldsymbol{p}^*_1\|}, \dots, \frac{\boldsymbol{p}^*_M}{\|\boldsymbol{p}^*_M\|}$
```

Normalization of Feature Maps. The pooling operator is shared across all channels. However, the intensity values in each channel could potentially have a different center and scale, making it hard for the same pooling to provide similar effect on all channels. To fix this, we normalize feature maps before forming the matrices $\boldsymbol{X}_k$ and applying Algorithm 1. More precisely, for a given feature map $f_k(i,j,c)$ ($k = 1, \ldots, K$, with K being size of the training set), the normalized feature map is defined as $g_k(i,j,c) \triangleq \frac{f_k(i,j,c)-\bar{f}(c)}{\sqrt{v(c)}}$, where $\bar{f}(c) \triangleq \frac{1}{IJK}\sum_{i,j,k} f_k(i,j,c)$ and $v(c) \triangleq \frac{1}{IJK}\sum_{i,j,k}(f_k(i,j,c) - \bar{f}(c))^2$. After the pooling operator is applied, we transform the feature map back to its original scale and center by multiplying by $v(c)$ and adding $\bar{f}(c)$. This helps the output of the pooling operation be consistent with rest of the network.

Use in Deep Networks Consider a trained deep network using some typical pooling operator. We can convert the pooling operator at any given layer to a learned one, by treating the hidden representation at that layer as the input feature maps to Algorithm 1 (after applying the channel normalization described above). We will then adapt the network weights for the learned pooling by retraining the network. This process can be repeated for multiple layers. In our experiments, however, we observe that sometimes learned pooling even at one layer can already give a boost in test accuracy.

Number of Eigenvectors. For simplicity of presentation, Algorithm 1 uses the top eigenvector. In principle, however, top few eigenvectors could be used instead. In fact, modern architectures often double the number of output channels while down-sampling via strided convolutions. To imitate that, we chose to select the top two eigenvectors from (18), which results in two feature maps per input feature map. This keeps the size of the hidden representation after pooling consistent between our method and the common practice.

Computing the Generalized Eigenvalue. For simpler exposition, in Section 3.3 and also Algorithm 1 we have used matrix inversion to solve the generalized eigenvalue problem. However, there are more efficient approaches for solving generalized eigenvalue problem without matrix inversion. In addition, we only need the top-1 or top-2 eigenvectors, which allows further efficiency in computation. There are numerical recipes that can leverage these two properties, such as `scipy.sparse.linalg.eigs` that we used for our Python implementation

Learning Pooling by SGD. One may wonder why not using gradient descent to optimize a total loss (sum of the usual cross-entropy and separability criterion), instead of Algorithm 1 and thus simultaneously learn network weights and pooling operator. The answer is that it is either impractical or leads to inferior performance. To learn the regularized pooling operator it is necessary to store the matrix $\boldsymbol{C}$ for each location in the output feature map. This would incur a memory cost of $O(MN^2)$ and would be extremely space inefficient. The performance with an un-regularized pooling map is reported in Table 4 in the appendix. In performs worse than our approach in almost all cases and in some, even worse than the baseline.

4 Experiments

We study the performance of our pooling operator on the SVHN (Netzer et al., 2011) and CIFAR-10/CIFAR-100 (Krizhevsky and Hinton, 2009) datasets. For the SVHN dataset we also experiment with a reduced 5% subset to measure the performance of our algorithm in presence of limited labelled data. We use 2 models for our experiments, a CNN model which is a 4-layer CovnNet and a 18 layer ResNet (He et al., 2016a) model. Both these models have 3 pooling layers in which they down sample

Table 1: Effect of replacing the pooling operator on generalization. We report the mean test error and standard deviation after averaging over 5 trials. When multiple pooling layers are replaced, it is indicated by separating the hyper-parameters by a comma. Experiments which result in improvements are highlighted in bold.

Model	Dataset	Baseline Error	Pooling Layer	α	Norm	Error with replaced pooling
CNN	Reduced SVHN	32.53 ± 0.284	3^{rd}	5	1	$\mathbf{21.56 \pm 0.525}$
	SVHN	10.68 ± 0.435	2^{nd}	25	2	$\mathbf{9.93 \pm 0.47}$
	CIFAR-10	14.59 ± 0.352	3^{rd}	45	2	16.08 ± 0.191
	CIFAR-100	45.04 ± 0.397	3^{rd}	5	2	45.13 ± 0.483
ResNet	Reduced SVHN	12.38 ± 0.33	2^{nd}	25	1	$\mathbf{10.46 \pm 0.534}$
	SVHN	4.07 ± 0.126	2^{nd}	10	1	$\mathbf{3.62 \pm 0.057}$
	CIFAR-10	4.57 ± 0.123	3^{rd}, 1^{st}	$15, 65$	$2, 1$	$\mathbf{4.3 \pm 0.151}$
	CIFAR-100	22.31 ± 0.2896	2^{nd}	5	1	$\mathbf{21.5 \pm 0.145}$

via strided convolutions. [2]

4.1 Effect on generalization

Table 1 shows the effect of our pooling operator on generalization. We are able to improve on the ResNet model in all settings and with the CNN model on both versions on the SVHN dataset. The largest gain is observed with the CNN model on the reduced SVHN dataset of over 10%. Even when the CNN model fails to improve on the CIFAR datasets, the performance is on par with the CNN.

4.2 Robustness to corruptions and perturbations

(Hendrycks and Dietterich, 2019) have developed a dataset of real-world corruptions to test model robustness. For these set of experiments, the model is trained on the original CIFAR-10 training set and evaluated on the modified test sets provided. We use the given CIFAR-10-C and CIFAR-10-P test sets and evaluate our approach by measuring the suggested quantities. For all of these measurements, we use the original ResNet architecture as the baseline.

In Table 2 we measure Corruption Error on the CIFAR-10-C dataset as suggested by (Hendrycks and Dietterich, 2019) . In the bottom most row we report the average corruption for each corruption type. We note that among others, the model with the replaced pooling operator is more robust in presence of geometric transformations for 4 out of 5 cases. We define geometric transformations as those which can move/displace pixels.

In Table 3, we measure how our algorithm responds to gradually applied perturbations with the CIFAR-10-P dataset. Each cell reports the Flip Probability, which indicates the probability of the predicted label changing in presence of a perturbation. The bottom row reports the Flip Rate which is the ratio of the flip probability of our model over the flip probability of the original ResNet model. We note that our model does better than the original network for 10 out of 14 perturbations and for 6 out 8 geometric perturbations.

---

[2] Additional details about the models, datasets and training can be found in the supplementary material, along with the source code. The implementation will be open-sourced with the camera ready version.

Table 2: Measuring robustness to corruptions as defined by (Hendrycks and Dietterich, 2019) on the CIFAR-10-C test set. Each cell lists the average test error (percentage) over 5 models for a particular corruption and severity. We use our CIFAR-10 model as described in Table 1. In the last row, we report the average corruption error (CE) across all 5 severities. In this metric, lower is better and the vanilla ResNet itself would score 100. We highlight corruptions for which we do better in bold.

		Corruption														
Model	Severity	Geometric					Non-Geometric									
		Defocus.	Frost.	Motion.	Zoom.	Elastic	Gauss.	Shot.	Impulse.	Snow	Frost	Fog	Bright.	Contr.	Pixel.	Jpeg.
ResNet	1	4.7	39.1	8.6	11.1	8.8	19.5	12.9	14.9	9.0	8.8	4.8	4.6	5.0	6.6	12.6
	2	5.8	37.7	15.0	14.1	9.0	38.3	21.3	27.4	17.2	13.6	5.6	4.9	7.7	11.6	18.6
	3	10.5	36.5	23.4	20.1	13.1	56.8	42.4	38.1	14.8	21.7	7.3	5.3	11.0	16.4	20.7
	4	10.5	36.5	23.4	20.1	13.1	56.8	42.4	38.1	14.8	21.7	7.3	5.3	11.0	16.4	20.7
	5	41.6	47.9	32.0	37.4	24.2	69.7	62.2	74.9	21.5	32.6	27.0	7.9	55.7	51.4	27.5
ResNet with pooling replaced	1	4.5	41.4	8.1	10.3	8.4	19.4	12.0	15.7	8.9	8.0	4.5	4.4	4.8	6.3	12.1
	2	5.6	39.9	14.5	13.0	8.5	40.8	21.3	28.5	17.4	12.6	5.5	4.8	7.8	11.4	17.6
	3	9.6	37.8	23.2	18.7	12.2	61.9	45.4	39.2	15.1	20.9	7.1	5.2	11.1	16.1	19.7
	4	18.4	52.3	23.2	24.5	19.7	69.4	54.6	60.9	17.4	22.3	9.6	5.8	18.7	34.6	22.6
	5	41.8	50.3	32.4	34.8	24.6	75.0	67.3	78.1	21.5	32.2	23.3	7.8	57.7	54.2	26.8
CE		**97.0**	105.1	**99.0**	**93.0**	**98.0**	107.2	105.9	103.8	100.7	**96.0**	**91.0**	**98.0**	102.5	103.2	**96.0**

Table 3: Measuring robustness to perturbations as defined by (Hendrycks and Dietterich, 2019) on the CIFAR-10-P dataset by using our best model from Table 1. In each cell we report Flip Probability (FP) averaged over 5 models. In the last row we report the Flip Rate when using default ResNet model as a baseline. When there are multiple severities we report the average. Flip probability and rate are reported out of 100, with lower being better. Improvements are highlighted in bold.

		Corruption													
Model	Severity	Geometric								Non-Geometric					
		Scale	Rot.	Tilt	Tran.	Shear	Motion.	Zoom.	Ga. B	Bright.	Spatter	Snow	Shot.	Speckle	Ga. N
ResNet	1	3.5	2.92	0.83	2.1	1.91	8.43	0.45	1.26	0.51	1.56	2.2	6.86	6.34	5.01
	2												10.56	10.39	14.26
	3												17.6	17.01	25.05
ResNet with pooling replaced	1	3.25	2.83	0.76	2.15	1.91	8.39	0.38	1.03	0.5	1.6	2.29	6.27	5.82	4.41
	2												9.96	9.84	13.51
	3												16.69	16.19	22.95
Flip Rate		**92.9**	**96.9**	**91.6**	102.4	100.0	**99.5**	**84.4**	**81.7**	**98.0**	102.6	104.1	**94.0**	**94.4**	**92.22**

5 Conclusion

We propose a more general pooling operator than currently being used in literature. We also present an algorithm to learn the pooling operator in closed form given the distribution of its inputs. Compared to pooling operations that are shared throughout spatial dimensions, ours allows more flexibility by being spatially varying. We replace the standard pooling operations in a CNN and a ResNet model and see benefits in generalizations on the CIFAR-10/CIFAR-100 and SVHN datasets. The operator is demonstrably more robust to unseen geometric transformations, which we show by evaluating on the CIFAR-10-C and CIFAR-10-P test sets.

References

Arandjelovic, R., Gronat, P., Torii, A., Pajdla, T., and Sivic, J. (2016). Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5297–5307.

Azulay, A. and Weiss, Y. (2018). Why do deep convolutional networks generalize so poorly to small image transformations? *arXiv preprint arXiv:1805.12177*.

Boureau, Y.-L., Le Roux, N., Bach, F., Ponce, J., and LeCun, Y. (2011). Ask the locals: multi-way local pooling for image recognition. In *ICCV'11-The 13th International Conference on Computer Vision*.

Bruna, J. and Mallat, S. (2013). Invariant scattering convolution networks. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1872–1886.

Cohen, N. and Shashua, A. (2016). Inductive bias of deep convolutional networks through pooling geometry. *arXiv preprint arXiv:1605.06743*.

Esteves, C., Allen-Blanchette, C., Zhou, X., and Daniilidis, K. (2018a). Polar transformer networks. In *International Conference on Learning Representations*.

Esteves, C., Allen-Blanchette, C., Zhou, X., and Daniilidis, K. (2018b). Polar transformer networks. In *International Conference on Learning Representations*.

Felzenszwalb, P. F., Girshick, R. B., McAllester, D., and Ramanan, D. (2009). Object detection with discriminatively trained part-based models. *IEEE transactions on pattern analysis and machine intelligence*, 32(9):1627–1645.

Fukushima, K. and Miyake, S. (1982). Neocognitron: A self-organizing neural network model for a mechanism of visual pattern recognition. In *Competition and cooperation in neural nets*, pages 267–285. Springer.

Girshick, R., Iandola, F., Darrell, T., and Malik, J. (2015). Deformable part models are convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 437–446.

Gong, Y., Wang, L., Guo, R., and Lazebnik, S. (2014). Multi-scale orderless pooling of deep convolutional activation features. In *European conference on computer vision*, pages 392–407. Springer.

He, K., Gkioxari, G., Dollár, P., and Girshick, R. (2017). Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969.

He, K., Zhang, X., Ren, S., and Sun, J. (2016a). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.

He, K., Zhang, X., Ren, S., and Sun, J. (2016b). Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer.

Hendrycks, D. and Dietterich, T. (2019). Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*.

Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.

Jégou, H., Douze, M., Schmid, C., and Pérez, P. (2010). Aggregating local descriptors into a compact image representation. In *CVPR 2010-23rd IEEE Conference on Computer Vision & Pattern Recognition*, pages 3304–3311. IEEE Computer Society.

Kanazawa, A., Sharma, A., and Jacobs, D. W. (2014). Locally scale-invariant convolutional neural networks. *CoRR*, abs/1412.5104.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Krizhevsky, A. and Hinton, G. (2009). Learning multiple layers of features from tiny images. Technical report, Citeseer.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012a). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012b). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105.

Larson, A. M. and Loschky, L. C. (2009). The contributions of central versus peripheral vision to scene gist recognition. *Journal of Vision*, 9(10):6–6.

LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.

Lee, C.-Y., Gallagher, P. W., and Tu, Z. (2016). Generalizing pooling functions in convolutional neural networks: Mixed, gated, and tree. In *Artificial Intelligence and Statistics*, pages 464–472.

Lee, C.-Y., Xie, S., Gallagher, P., Zhang, Z., and Tu, Z. (2014). Deeply-supervised nets. *arXiv preprint arXiv:1409.5185*.

Li, C., Reiter, A., and Hager, G. D. (2015). Beyond spatial pooling: fine-grained representation learning in multiple domains. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4913–4922.

Liu, K. (2018). pytorch-cifar. `https://github.com/kuangliu/pytorch-cifar`.

Malinowski, M. and Fritz, M. (2013). Learnable pooling regions for image classification. *arXiv preprint arXiv:1301.3516*.

Miech, A., Laptev, I., and Sivic, J. (2017). Learnable pooling with context gating for video classification. *arXiv preprint arXiv:1706.06905*.

Mobahi, H., Zitnick, C. L., and Ma, Y. (2012). Seeing through the blur. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1736–1743. IEEE.

Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. (2011). Reading digits in natural images with unsupervised feature learning.

Passalis, N. and Tefas, A. (2017). Learning bag-of-features pooling for deep convolutional neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5755–5763.

Saeedan, F., Weber, N., Goesele, M., and Roth, S. (2018). Detail-preserving pooling in deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9108–9116.

Sifre, L. and Mallat, S. (2013). Rotation, scaling and deformation invariant scattering for texture discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1233–1240.

Sun, M., Song, Z., Jiang, X., Pan, J., and Pang, Y. (2017). Learning pooling for convolutional neural network. *Neurocomputing*, 224:96–104.

Worrall, D. E., Garbin, S. J., Turmukhambetov, D., and Brostow, G. J. (2017). Harmonic networks: Deep translation and rotation equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5028–5037.

Zhang, R. (2019). Making convolutional networks shift-invariant again. *arXiv preprint arXiv:1904.11486*.

Zhou, Y., Sun, X., Liu, D., Zha, Z., and Zeng, W. (2017). Adaptive pooling in multi-instance learning for web video annotation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 318–327.

6 Supplementary Appendix

6.1 Derivation of Closed Form

The goal is to solve the following optimization:

$$\boldsymbol{p}^* = \arg\max_{\boldsymbol{p}} \boldsymbol{p}^T_{1\times N} \boldsymbol{A}_{N\times N} \boldsymbol{p}_{N\times 1} \; - \; \rho \boldsymbol{p}^T \boldsymbol{C} \boldsymbol{p} \quad , \quad \text{s.t. } \boldsymbol{p}^T_{1\times N} \boldsymbol{B}_{N\times N} \boldsymbol{p}_{N\times 1} = 1 \,, \tag{19}$$

Using Lagrange multiplier λ, the optimization has the following Lagrangian:

$$L \; = \; \boldsymbol{p}^T_{1\times N} \boldsymbol{A}_{N\times N} \boldsymbol{p}_{N\times 1} \; + \; \lambda(\boldsymbol{p}^T_{1\times N} \boldsymbol{B}_{N\times N} \boldsymbol{p}_{N\times 1} - 1) \; - \; \rho \, \boldsymbol{p}^T_{1\times N} \boldsymbol{C}_{N\times N} \boldsymbol{p}_{N\times 1} \,. \tag{20}$$

The derivative of L w.r.t. $\boldsymbol{p}$ is:

$$\frac{\partial L}{\partial \boldsymbol{p}} = \boldsymbol{A}_{N\times N} \boldsymbol{p}_{N\times 1} + \lambda \Big(\boldsymbol{B}_{N\times N} - \frac{\rho}{\lambda} \boldsymbol{C}_{N\times N} \Big) \boldsymbol{p}_{N\times 1} \,. \tag{21}$$

It is not difficult to verify that $\lambda \leq 0$. Hence, defining $\alpha \triangleq -\frac{\rho}{\lambda}$, we learn that $\alpha \geq 0$ (because $\rho > 0$).

$$\frac{\partial L}{\partial \boldsymbol{p}} = \boldsymbol{A}_{N\times N} \boldsymbol{p}_{N\times 1} + \lambda \Big(\boldsymbol{B}_{N\times N} + \alpha \, \boldsymbol{C}_{N\times N} \Big) \boldsymbol{p}_{N\times 1} \,. \tag{22}$$

6.2 Dataset details

In our descriptions, an epoch indicates the number of steps necessary to perform one full pass over the training data. Whenever reduced datasets are used, the number of steps in each epoch is scaled accordingly. To choose hyper-parameters we use cross validation performed by evaluating performance on 5 distinct random held-out subsets of the training data.

- **CIFAR-10/CIFAR-100** For both the CIFAR datasets, we normalize the images by subtracting the mean and dividing by the standard deviation over the entire training set. During training, images are augmented using the technique described in (Lee et al., 2014) which consists of padding images by 4 pixels and randomly cropping a 32×32 piece along with adding horizontal flips. All models on the CIFAR datasets are trained for 350 epochs with the learning for the ResNet model decayed at epochs 150 and 250. For cross-validation we use 10 % subsets containing 5000 samples.
- **SVHN** For the SVHN dataset each image is normalized separately with no additional data-augmentation applied. All models on this dataset are trained for 150 epochs with the learning rate for the ResNet model decayed at epochs 50 and 100. Cross validation is done by holding out 5 % of the training data.
- **Reduced SVHN** We use this dataset to measure the performance of our algorithm in presence of less labelled samples. This is a reduced version of the SVHN dataset in which we only train with 5 % of the training data. For cross validation, 20 % of the reduced dataset is held out.

6.3 Model Details

We use our pooling operator within two models. These models perform spatial pooling by using stride 2 convolutions, and we experiment with replacing the 3 different layers in which they reduce spatial dimensions, with the exception of the final average pooling layer:

- **CNN** This is a 4-layer ConvNet with convolution kernels of size 3 × 3. The first convolution uses 64 channels and is followed by a ReLU non-linearity. The 3 convolution layers which follow are strided convolutions with a stride of 2 (to reduce spatial dimensions) and double the number of channels from the previous layer, each of them followed by batch-norm (Ioffe and Szegedy, 2015) and ReLU. Towards the end, the feature map is aggregated via global average pooling and fed into a linear layer which outputs logits. The CNN model is trained using the Adam optimizer (Kingma and Ba, 2014) with a fixed learning rate of 0.001.
- **ResNet** The second architecture we use is a 18 layer ResNet described (He et al., 2016b) with its hyper parameters chosen form the implementation by (Liu, 2018). The network is trained with SGD and momentum coefficient of 0.9 and a starting learning rate of 0.1, decayed a factor of 10 after fixed number of epochs for each dataset.

6.4 Training Procedure

As the first step of our algorithm we train our models till convergence using the default pooling operation in each model. This is followed by using Algorithm 1 (in main paper) to compute our pooling operator along with the normalization parameters f and g. While estimating matrices $\boldsymbol{A}$ and $\boldsymbol{B}$, for both models, we use at most 10000 samples per class. We then replace each pooling layer, one at a time, with our own pooling operator with various values of α and choice of norm, and re-train the network from a random initialization. We choose the setting that leads to the best average cross-validation error. Using this setting, we train on the full training dataset and report numbers on the test set.

It is possible to use this procedure multiple times to replace more than one pooling layer. In our experiments, we tried replacing multiple pooling layers while using the CIFAR-10 and CIFAR-100 datasets. Only on CIFAR-10, replacing with 3$^{\text{rd}}$ and the 1$^{\text{st}}$ pooling layer respectively in a ResNet led to a non-trivial reduction in cross-validation error. For all other models and datasets we report performance after replacing only a single pooling layer inside the model.

6.5 Additional visualizations

6.5.1 ResNet on reduced SVHN

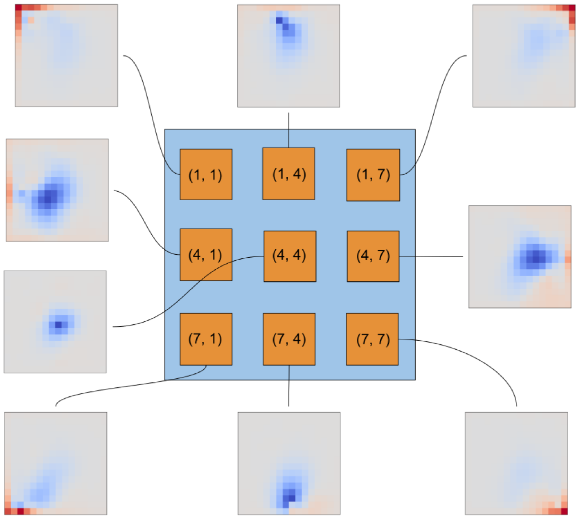
(1, 1)
(1, 4)
(1, 7)
(4, 1)
(4, 4)
(4, 7)
(7, 1)
(7, 4)
(7, 7)

Figure 2: Visualization of best performing pooling map on ResNet mode with Reduced SVHN dataset.

6.5.2 ResNet on CIFAR-100

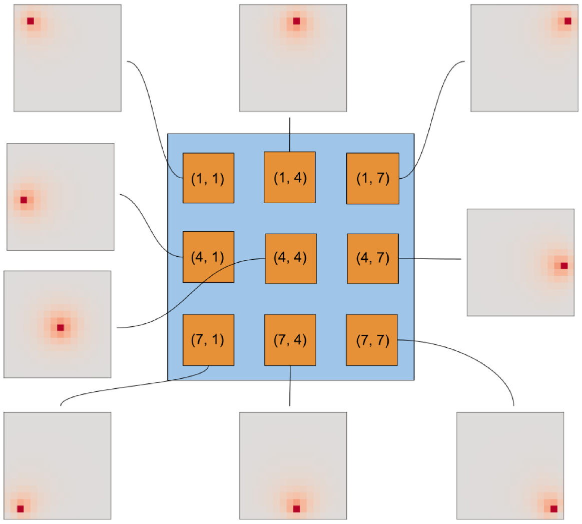
(1, 1)
(1, 4)
(1, 7)
(4, 1)
(4, 4)
(4, 7)
(7, 1)
(7, 4)
(7, 7)

Figure 3: Visualization of best performing pooling map on ResNet with the CIFAR-100 dataset.

6.6 Comparing operator learned through SGD

Table 4: Effect of learning a pooling map through SGD. To keep the results comparable with Table 1 we learned 2 distinct pooling maps which has the effect of doubling the number of input channels while down-sampling. We chose the pooling layer by using cross-validation as described in Section 6.4

Model	Dataset	Baseline Error	Pooling Layer	Error with SGD
CNN	Reduced SVHN	32.53 ± 0.284	3rd	22.49 ± 0.49
	SVHN	10.68 ± 0.435	3rd	9.19 ± 0.218
	CIFAR-10	14.59 ± 0.352	3rd	16.77 ± 0.796
	CIFAR-100	45.04 ± 0.397	3rd	48.17 ± 0.363
ResNet	Reduced SVHN	12.38 ± 0.33	3rd	13.78 ± 0.357
	SVHN	4.07 ± 0.126	3rd	3.97 ± 0.09
	CIFAR-10	4.57 ± 0.123	1st	18.33 ± 2.612
	CIFAR-100	22.31 ± 0.2896	1st	41.09 ± 2.346