

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Learning a smooth kernel regularizer for convolutional neural networks

Permalink

<https://escholarship.org/uc/item/63c6q8p9>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

Authors

Feinman, Reuben

Lake, Brenden M.

Publication Date

2019

Peer reviewed

Learning a smooth kernel regularizer for convolutional neural networks

Reuben Feinman (reuben.feinman@nyu.edu)

Center for Neural Science
New York University

Brenden M. Lake (brenden@nyu.edu)

Department of Psychology and Center for Data Science
New York University

Abstract

Modern deep neural networks require a tremendous amount of data to train, often needing hundreds or thousands of labeled examples to learn an effective representation. For these networks to work with less data, more structure must be built into their architectures or learned from previous experience. The learned weights of convolutional neural networks (CNNs) trained on large datasets for object recognition contain a substantial amount of structure. These representations have parallels to simple cells in the primary visual cortex, where receptive fields are smooth and contain many regularities. Incorporating smoothness constraints over the kernel weights of modern CNN architectures is a promising way to improve their sample complexity. We propose a smooth kernel regularizer that encourages spatial correlations in convolution kernel weights. The correlation parameters of this regularizer are learned from previous experience, yielding a method with a hierarchical Bayesian interpretation. We show that our correlated regularizer can help constrain models for visual recognition, improving over an L2 regularization baseline.


Keywords: convolutional neural networks; regularization; model priors; visual recognition

Introduction

Convolutional neural networks (CNNs) are powerful feed-forward architectures inspired by mammalian visual processing capable of learning complex visual representations from raw image data (LeCun et al., 2015). These networks achieve human-level performance in some visual recognition tasks; however, their performance often comes at the cost of hundreds or thousands of labelled examples. In contrast, children can learn to recognize new concepts from just one or a few examples (Bloom, 2000; Xu & Tenenbaum, 2007), evidencing the use of rich structural constraints (Lake et al., 2017). By enforcing structure on neural networks to account for the regularities of visual data, it may be possible to substantially reduce the number of training examples they need to generalize. In this paper, we introduce a soft architectural constraint for CNNs that encourages smooth, correlated structure on their convolution kernels through transfer learning.¹ We see this as an important step towards a general, off-the-shelf CNN regularizer that operates independently of previous experience.

The basis for our constraint is the idea that the weights of a convolutional kernel should in general be well-structured and smooth. The weight kernels of CNNs that have been trained on the large-scale ImageNet object recognition task contain a substantial amount of structure. These kernels have parallels to simple cells in primary visual cortex, where smooth receptive fields implement bandpass oriented filters of various scale (Jones & Palmer, 1987).

¹Experiments from this paper can be reproduced with the code found at <https://github.com/rfeinman/SK-regularization>.



(a) VGG16 layer-1 kernels



(b) i.i.d. Gaussian (L2-reg)



(c) correlated Gaussian (SK-reg)

Figure 1: Kernel priors for VGG16. The layer-1 convolution kernels of VGG16, shown in (a), possess considerable correlation structure. An i.i.d. Gaussian prior that has been fit to the VGG layer-1 kernels, samples from which are shown in (b), captures little of the structure in these kernels. A correlated multivariate Gaussian prior, samples from which are shown in (c), captures the correlation structure of these kernels well.

The consistencies of visual receptive fields are explained by the regularities of image data. Locations within the kernel window have parallels to locations in image space, and images are generally smooth (Li, 2009). Consequently, smooth, structured receptive fields are necessary to capture important visual features like edges. In landmark work, Hubel & Wiesel (1962) discovered edge-detecting features in the primary visual cortex of cat. Since then, the community has successfully modeled receptive fields in early areas of mammalian visual cortex using Gabor kernels (Jones & Palmer, 1987). These kernels are smooth and contain many spatial correlations. In later stages of visual processing, locations of kernel space continue to parallel image space; however, inputs to these kernels are visual features, such as edges. Like earlier layers, these layers also benefit from smooth, structured kernels that capture correlations across the input space. Geisler et al. (2001) showed that human contour perception—an important component of object recognition—is well-explained by a model of edge co-occurrences, suggesting that correlated receptive fields are useful in higher layers of processing as well.

Despite the clear advantages of structured receptive fields, constraints placed on the convolution kernels of CNNs are typically chosen to be as general as possible, with neglect of this structure. L2 regularization—the standard soft constraint applied to kernel weights, which is interpreted as a zero-mean, independent identically distributed (i.i.d.) Gaussian prior—treats each weight as an independent random variable, with no correlations between weights expected a priori. Fig. 1 shows the layer-1 convolutional kernels of VGG16, a ConvNet trained on the large-scale ImageNet object recognition task (Simonyan & Zisserman, 2015). Fig. 1b shows some samples from an i.i.d. Gaussian prior, the equivalent of L2 regularization. Clearly, this prior captures little of the correlation structure possessed by the kernels.

A simple and logical extension of the i.i.d. Gaussian prior is a correlated multivariate Gaussian prior, which is capable of capturing some of the covariance structure in the convolution kernels. Fig. 1c shows some samples from a correlated Gaussian prior that has been fit to the VGG16 kernels. This prior provides a much better model of the kernel distribution. In this paper, we perform a series of controlled CNN learning experiments using a smooth kernel regularizer—which we denote “SK-reg”—based on a correlated Gaussian prior. The correlation parameters of this prior are obtained by fitting a Gaussian to the learned kernels from previous experience. We compare SK-reg to standard L2 regularization in two object recognition use cases: one with simple silhouette images, and another with Tiny ImageNet natural images. In the condition of limited training data, SK-reg yields improved generalization performance.

Background

Our goal in this paper is to introduce new a priori structure into CNN receptive fields to account for the regularities of image data and help reduce the sample complexity of these models. Previous methods from this literature often require a fixed model architecture that cannot be adjusted from task to task. In contrast, our method enforces structure via a statistical prior over receptive field weights, allowing for flexible architecture adaption to the task at hand. Nevertheless, in this section we review the most common approaches to structured vision models.

A popular method to enforce structure on visual recognition models is to apply a fixed, pre-specified representation. In computational vision, models of image recognition consist of a hierarchy of transformations motivated by principles from neuroscience and signal processing (e.g., Serre et al., 2007; Bruna & Mallat, 2013). These models are effective at extracting important statistical features from natural images, and they have been shown to provide a useful image representation for SVMs, logistic regression and other “shallow” classifiers when applied to recognition tasks with limited training data. Unlike CNNs, the kernel parameters of these models are not learned by gradient descent. As result, these features may not be well-adapted to the specific task at hand.

In machine learning, it is commonplace to use the features from CNNs trained on large object recognition datasets as a generic image representation for novel vision tasks (Donahue et al., 2014; Razavian et al., 2014). Due to the large variety of training examples that these CNNs receive, the learned features of these networks provide an effective representation for a range of new recognition tasks. Some *meta-learning* algorithms use a similar form of feature transfer, where a feature representation is first learned via a series of classification episodes, each with a different support set of classes (e.g., Vinyals et al., 2016). As with pre-specified feature models, the representations of these feature transfer models are fixed for the new task; thus, performance on the new task may be sub-optimal.

Beyond fixed feature representations, other approaches use a pre-trained CNN as an initialization point for a new network, following with a fine-tuning phase where network weights are further optimized for a new task via gradient descent (e.g., Girshick et al., 2014; Girshick, 2015). By adapting the CNN representation to the new task, this approach often enables better performance than fixed feature methods; however, when the scale of the required adaptation is large and the training data is limited, fine-tuning can be difficult. Finn et al. (2017) proposed a modification of the pre-train/fine-tune paradigm called model-agnostic meta-learning (MAML) that enables flexible adaptation in the fine-tuning phase when the training data is limited. During pre-training (or *meta-learning*), MAML optimizes for a representation that can be easily adapted to a new learning task in a later phase. Although effective for many use cases, this approach is unlikely to generalize well when the type of adaptation required differs significantly from the adaptations seen in the meta-learning episodes. A shared concern for all pre-train/fine-tune methods is that they require a fixed model architecture between the pre-train and fine-tune phases.

The objective of our method is distinct from those of fixed feature representations and pre-train/fine-tune algorithms. In this paper, we study the structure in the learned parameters of vision models, with the aim of extracting general structural principles that can be incorporated into new models across a broad range of learning tasks. SK-reg serves as a parameter prior over the convolution kernels of CNNs and has a theoretical foundation in Bayesian parameter estimation. This approach facilitates a CNN architecture and representation that is adapted to the specific task at hand, yet that possesses adequate structure to account for the regularities of image data. The SK-reg prior is learned from previous experience, yielding an interpretation of our algorithm as a method for hierarchical Bayesian inference.

Independently of our work, Atanov et al. (2019) developed the *deep weight prior*, an algorithm to learn and apply a CNN kernel prior in a Bayesian framework. Unlike our prior, which is parameterized by a simple multivariate Gaussian, the deep weight prior uses a sophisticated density estimator parameterized by a neural network to model the learned kernels of previously-trained CNNs. The application of this prior to new learning tasks requires variational inference with a well-calibrated variational distribution. Our goal with SK-reg differs in that we aim to provide an interpretable, generalizable prior for CNN weight kernels that can be applied to existing CNN training algorithms with little modification.

Bayesian interpretation of regularization

From the perspective of Bayesian parameter estimation, the L2 regularization objective can be interpreted as performing *maximum a-posteriori* inference over CNN parameters with a zero-mean, i.i.d. Gaussian prior. Here, we review this connection, and we discuss the extension to SK-reg.




Figure 2: SK-reg workflow. A) First, a CNN is trained repeatedly (20x) on an object recognition task. B) Next, the learned parameters of each CNN are studied and statistics are extracted. For each convolution layer, kernels from the multiple CNNs are consolidated, yielding a kernel dataset for the layer. A multivariate Gaussian is fit to each kernel dataset. C) SK-reg is applied to a fresh CNN trained on a new learning task with limited training data (possibly with a different architecture or numbers of kernels), using the resulting Gaussians from each layer.

L2 regularization. Assume we have a dataset $X = \{x_1, \dots, x_N\}$ and $Y = \{y_1, \dots, y_N\}$ consisting of N images x_i and N class labels y_i . Let θ define the parameters of the CNN that we wish to estimate. The L2 regularization objective is stated as follows:

$$\tilde{\theta} = \arg \max_{\theta} \log p(Y | \theta; X) - \lambda * \theta^T \theta. \quad (1)$$

Here, the first term of our objective is our prediction accuracy (classification log-likelihood), and the second term is our L2 regularization penalty.

From a Bayesian perspective, this objective can be thought of as finding the *maximum a-posteriori* (MAP) estimate of the network parameter posterior $p(\theta | Y; X) \propto p(Y | \theta; X) * p(\theta)$, leading to the optimization problem

$$\tilde{\theta} = \arg \max_{\theta} \log p(Y | \theta; X) + \log p(\theta). \quad (2)$$

To make the connection with L2 regularization, we assume a zero-mean, i.i.d Gaussian prior over the parameters θ of a weight kernel, written as

$$p(\theta) = \frac{1}{Z} \exp\left(-\frac{1}{2\sigma^2} \theta^T \theta\right). \quad (3)$$

With this prior, Eq. 2 becomes

$$\tilde{\theta} = \arg \max_{\theta} \log p(Y | \theta; X) - \frac{1}{2\sigma^2} \theta^T \theta,$$

which is the L2 objective of Eq. 1, with $\lambda = \frac{1}{2\sigma^2}$.

SK regularization. The key idea behind SK-reg is to extend the L2 Gaussian prior to include a non-diagonal covariance matrix; i.e., to add correlation. In the case of SK-reg,

the prior over kernel weights θ of Eq. 3 becomes

$$p(\theta) = \frac{1}{Z} \exp\left(-\frac{1}{2} \theta^T \Sigma^{-1} \theta\right)$$

for some covariance matrix Σ , and the new objective is written

$$\tilde{\theta} = \arg \max_{\theta} \log p(Y | \theta; X) - \lambda * \theta^T \Sigma^{-1} \theta. \quad (4)$$

Hierarchical Bayes. When Σ is learned from previous experience, SK-reg can be interpreted as approximate inference in a hierarchical Bayesian model. The SK regularizer for a CNN with C layers, $\Sigma = \{\Sigma_1, \dots, \Sigma_C\}$, assumes a unique zero-mean Gaussian prior $\mathcal{N}(\theta_i; 0, \Sigma_i)$ over the weight kernels for each convolutional layer, $\theta = \{\theta_1, \dots, \theta_C\}$. Due to the regularities of the visual world, it is plausible that effective general priors exist for each layer of visual processing. In this paper, transfer learning is used to fit the prior covariances Σ from previous datasets $X^{1:M-1}$ and $Y^{1:M-1}$, which informs the solution for a new problem X^M and Y^M , yielding the hierarchical Bayesian interpretation depicted in Fig. 3. Task-specific




Figure 3: A hierarchical Bayesian interpretation of SK-reg. A point estimate of prior parameters Σ is first computed with MAP estimation. Next, this prior is applied to estimate CNN parameters θ^j in a new task.

CNN parameters $\theta^{1:M}$ are drawn from a common Σ , and Σ has a hyperprior specified by β . Ideal inference would compute $p(Y^M|Y^{1:M-1}; X^{1:M})$, marginalizing over $\theta^{1:M}$ and Σ .

We propose a very simple empirical Bayes procedure for learning the kernel regularizer in Eq. 4 from data. First, $M - 1$ CNNs are fit independently to the datasets $X^{1:M-1}$ and $Y^{1:M-1}$ using standard methods, in this case optimizing Eq. 1 to get point estimates $\tilde{\theta}^{1:M-1}$. Second, a point estimate $\tilde{\Sigma}$ is computed by maximizing $p(\Sigma|\tilde{\theta}^{1:M-1}; \beta)$, which is a simple regularized covariance estimator. Last, for a new task M with training data X^M and Y^M , a CNN with parameters θ^M is trained with the SK-reg objective (Eq. 4), with $\Sigma = \tilde{\Sigma}$.

This procedure can be compared with the hierarchical Bayesian interpretation of MAML (Grant et al., 2018). Unlike MAML, our method allows flexibility to use different architectures for different datasets/episodes, and the optimizer for θ^M is run to convergence rather than just a few steps.

Experiments

We evaluate our approach within a set of controlled visual learning environments. SK-reg parameters Σ_i for each convolution layer θ_i are determined by fitting a Gaussian to the kernels acquired from an earlier learning phase. We divide our learning tasks into two unique phases, applying the same CNN architecture in each case. We note that our approach does not require a fixed CNN architecture across these two phases; the number of feature maps in each layer may be easily adjusted. A depiction of the two learning phases is given in Fig. 2.

Phase 1. The goal of phase 1 is to extract general principles about the structure of learned convolution kernels by training an array of CNNs and collecting statistics about the resulting kernels. In this phase, we train a CNN architecture to classify objects using a sufficiently large training set with numerous examples per object class. Training is repeated multiple times with unique random seeds, and the learned convolution kernels are stored for each run. During this phase, standard L2 regularization is applied to enforce a minimal constraint on each layer’s weights (optimization problem of Eq. 1). After training, the convolution kernels from each run are consolidated, holding each layer separate. A multivariate Gaussian is fit to the centered kernel dataset of each layer, yielding a distribution $N(0, \Sigma_i)$ for each convolution layer i . To ensure the stability of the covariance estimators, we apply shrinkage to each covariance estimate, mixing the empirical covariance with an identity matrix of equal dimensionality. This can be interpreted as a hyperprior $p(\Sigma; \beta)$ (Fig. 3) that favors small correlations. The optimal mixing parameter is determined via cross-validation.

Phase 2. In phase 2, we test the aptitude of SK-reg on a new visual recognition task, applying the covariance matrices Σ_i obtained from phase 1 to regularize each convolution layer i in a freshly-trained CNN (optimization problem of Eq. 4). In order to adequately test the generalization capability of our



Figure 4: Exemplars of the phase 1 silhouette object classes.

Layer	Window	Stride	Features	λ
Input (200x200x3)				
Conv2D	5x5	2	5	0.05
MaxPooling2D	3x3	3		
Conv2D	5x5	1	10	0.05
MaxPooling2D	3x3	2		
Conv2D	5x5	1	8	0.05
MaxPooling2D	3x3	1		
FullyConnected			128	0.01
Softmax				

Table 1: CNN architecture. Layer hyperparameters include window size, stride, feature count, and regularization weight (λ). Dropout is applied after the last pooling layer and the fully-connected layer with rates 0.2 and 0.5, respectively.

algorithm, we use a new set of classes that differ from the phase 1 classes in substantial ways, and we provide just a few training examples from each class. Performance of SK-reg is compared against standard L2 regularization.

Silhouettes

As a preliminary use case, we train our network using the binary shape image dataset developed at Brown University², henceforth denoted “Silhouettes.” Silhouette images are binary masks that depict the structural form of various object classes. Simple shape-based stimuli such as these provide a controlled learning environment for studying the inductive biases of CNNs (Feinman & Lake, 2018). We select a set of 20 well-structured silhouette classes for phase 1, and a set of 10 unique, well-structured classes for phase 2 that differ from phase 1 in their consistency and form. The images are padded to a fixed size of 200×200 .

During phase 1, we train our network to perform 20-way object classification. Exemplars of the phase 1 classes are shown in Fig. 4. The number of examples varies for each class, ranging from 12 to 49 with a mean of 24. Class weighting is used to remedy class imbalances. To add complexity to the silhouette images, colors are assigned randomly to each silhouette before training. During training, random translations, rotations and horizontal flips are applied at each training epoch to improve generalization performance.

We use a CNN architecture with 3 convolution layers, each followed by a max pooling layer (see Table 1). Hyperparameters including convolution window size, pool size, and filter counts were selected via randomized grid-search, using a validation set with examples from each class to score candidate values. A rectified linear unit (ReLU) nonlinearity is applied to the output of each convolution layer, as well as to the

²The binary shape dataset is available in the “Databases” section at <http://vision.lcms.brown.edu>




Figure 5: Learned first-layer kernels vs. Gaussian samples. (a) depicts some of the learned first-layer kernels acquired from phase 1 silhouette training. For comparison, (b) shows a few samples from a multivariate Gaussian that was fit to the first-layer kernel dataset.

fully-connected layer. The network is trained 20 times using the Adam optimizer, each time with a unique random initialization. It achieves an average validation accuracy of 97.7% across the 20 trials, indicating substantial generalization.

Following the completion of phase 1 training, a kernel dataset is obtained for each convolution layer by consolidating the learned kernels for that layer from the 20 trials. Covariance matrices Σ_i for each layer i are obtained by fitting a multivariate Gaussian to the layer’s kernel dataset. For a first-layer convolution with window size $K \times K$, this Gaussian has dimensionality $3K^2$, equal to the window area times RGB depth. We model the input channels as separate variables in layer 1 because these channels have a consistent interpretation as the RGB color channels of the input image. For remaining convolution layers, where the interpretation of input channels may vary from case to case, we treat each input channel as an independent sample from a Gaussian with dimensionality K^2 . The kernel datasets for each layer are centered to ensure zero mean, typically requiring only a small perturbation vector.

To ensure that our multivariate Gaussians model the kernel data well, we computed the cross-validated log-likelihoods of this estimator on each layer’s kernel dataset and compared them to those of an i.i.d. Gaussian estimator fit to the same data. The multivariate Gaussian achieved an average score of 358.5, 413.3 and 828.1 for convolution layers 1, 2 and 3, respectively. In comparison, the i.i.d. Gaussian achieved an average score of 144.4, 289.6 and 621.9 for the same layers. These results confirm that our multivariate Gaussian provides an improved model of the kernel data. Some examples of the first-layer convolution kernels are shown in Fig. 5 alongside samples from our multivariate Gaussian that was fit to the first-layer kernel dataset. The samples appear structurally consistent with our phase 1 kernels.

In phase 2, we train our CNN on a new 10-way classification task, providing the network with just 3 examples per class for gradient descent training and 3 examples per class for validation. Colors are again added at random to each silhouette in the dataset. The network is initialized randomly, and we apply SK-reg to the convolution kernels of each layer during training using the covariance matrices obtained in phase 1. Our validation set is used to track and save the best model over the course of the training epochs (early

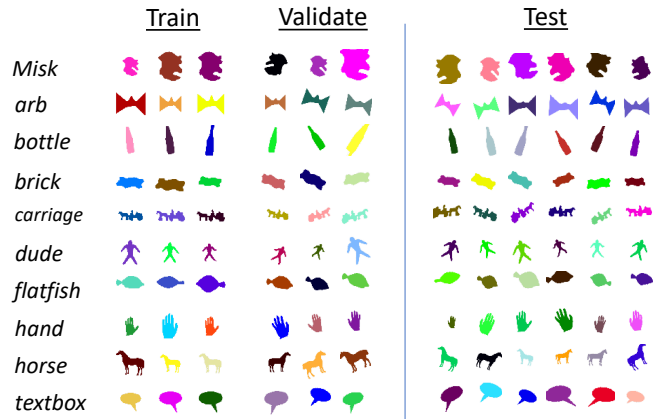


Figure 6: Silhouettes phase 2 datasets. 3 examples per class are provided in both the train and validation sets. A holdout test set with 6 examples per class is used to evaluate final model performance.

Method	λ	Cross-entropy	Accuracy
L2	0.214	2.000 (+/- 0.033)	0.530 (+/- 0.013)
SK	0.129	0.597 (+/- 0.172)	0.821 (+/- 0.056)

Table 2: Silhouettes phase 2 results. For each regularization method, the optimal regularization weight λ was selected via grid-search. Results show the average cross-entropy and classification accuracy achieved on the holdout test set over 10 phase 2 training runs.

stopping). A holdout set with 6 examples per class is used to assess the final performance of the model. A depiction of the train, validation and test sets used for phase 2 is given in Fig. 6. The validation and test images have been shifted, translated and flipped to make for a more challenging generalization test. Similar to phase 1, random shifts, rotations and horizontal flips are applied to the training images at each training epoch. As a baseline, we also train our CNN using standard L2 regularization.

The regularization weight λ is an important hyperparameter of both SK and L2 regularization. Before performing the phase 2 training assessment, we use a validated grid search to select the optimal λ for each regularization method, applying our train/validate sets.³ The same weight λ is applied to each convolution layer, as done in phase 1.

Results. With our optimal λ values selected, we trained our CNN on the 10-way phase 2 classification task of Fig. 6, comparing SK regularization to a baseline L2 regularization model. Average results for the two models collected over 10 training runs are presented in Table 2. Average test accuracy is improved by roughly 55% with the addition of SK reg, a substantial performance boost from 53.0% correct to 82.1% correct. Clearly, a priori structure is beneficial to generalization in this use case. An inspection of the learned kernels confirms that SK-reg encourages the structure we expect; these

³To yield interpretable λ values that can be compared between the SK and L2 cases, we normalize each covariance matrix to unit determinant by applying a scaling factor c , such that $\det(c\Sigma) = \det(I)$.

kernels look visually similar to samples from the Gaussian (e.g. Fig. 5).

Tiny ImageNet

Our silhouette experiment demonstrates the effectiveness of SK-reg when the parameters of the regularizer are determined from the structure of CNNs trained on a similar image domain. However, it remains unclear whether these regularization parameters can generalize to novel image domains. Due to the nature of the silhouette images, the silhouette recognition task encourages representations with properties that are desirable for object recognition tasks in general. Categorizing silhouettes requires forming a rich representation of shape, and shape perception is critical to object recognition. Therefore, this family of representation may be useful in a variety of object recognition tasks.

To test whether our kernel priors obtained from silhouette training generalize to a novel domain, we applied SK-reg to a simplified version of the Tiny ImageNet visual recognition challenge, using covariance parameters fitted to silhouette-trained CNNs. Tiny ImageNet images were up-sampled with bilinear interpolation from their original size of 64×64 to mirror the Silhouette size 200×200 . We selected 10 well-structured ImageNet classes that contain properties consistent with the silhouette images.⁴ We performed 10-way image classification with these classes, using the same CNN architecture from Table 1 and applying the SK-reg soft constraint. The network is provided 10 images per class for training and 10 per class for validation. Because of the increased complexity of the Tiny ImageNet data, a larger number of examples per class is merited to achieve good generalization performance. A holdout test set with 20 images per class is used to evaluate performance. Fig. 7 shows a breakdown of the train, validate and test sets.

A few modifications were made to account for the new image data. First, we modified the phase 1 silhouette training used to acquire our covariance parameters, this time applying random colors to both the foreground and background of each silhouette. Previously, each silhouette overlaid a strictly white background. Consequently, the edge detectors of the learned CNNs would be unlikely to generalize to novel color gradients. Second, we added additional regularization to our covariance estimators to avoid over-fitting and help improve the generalization capability of the resulting kernel priors. Due to the nature of the phase 2 task in this experiment, and the extent to which the images differ from phase 1, additional regularization was necessary to ensure that our kernel priors could generalize. Specifically, we applied L1-regularized inverse covariance estimation (Friedman et al., 2008) to estimate each Σ_i , which can be interpreted as a hyperprior $p(\Sigma; \beta)$ (Fig. 3) that favors a sparse inverse covariance (Lake & Tenenbaum, 2010).

Similar to the silhouettes experiment, the validation set is

⁴Desirable classes have a uniform, centralized object with consistent shape properties and a distinct background.




Figure 7: Tiny ImageNet datasets. 10 classes were selected to form a 10-way classification task. The train and validate sets each contain 10 examples per class. The holdout test set contains 20 examples per class.

Method	λ	Cross-entropy	Accuracy
L2	0.450	1.073 (+/- 0.102)	0.700 (+/- 0.030)
SK	0.450	0.956 (+/- 0.180)	0.776 (+/- 0.035)

Table 3: Tiny ImageNet SK-reg and L2 results. Table shows the average cross-entropy and classification accuracy achieved on the holdout test set over 10 training runs.

used to select weighting hyperparameter λ and to track the best model over the course of learning. As a baseline, we again compared SK-reg to a λ -optimized L2 regularizer.

Results. SK-reg improved the average holdout performance received from 10 training runs as compared to an L2 baseline, both in accuracy and cross-entropy. Results for each regularization method, as well as their optimal λ values, are reported in Table 3. An improvement of 8% in test accuracy suggests that some of the structure captured by our kernel prior is useful even in a very distinct image domain. The complexity of natural images like ImageNet is vast in comparison to simple binary shape masks; nonetheless, our prior from phase 1 silhouette training is able to influence ImageNet learning in a manner that is beneficial to generalization.

Discussion

Using a set of controlled visual learning experiments, our work in this paper demonstrates the potential of structured receptive field priors in CNN learning tasks. Due to the properties of image data, smooth, structured receptive fields have many desirable properties for visual recognition models. In our experiments, we have shown that a simple multivariate Gaussian model can effectively capture some of the structure in the learned receptive fields of CNNs trained on simple object recognition tasks. Samples from the fitted Gaussians are visually consistent with learned receptive fields, and when applied as a model prior for new learning tasks, these Gaussians can help a CNN generalize in conditions of limited training data. We demonstrated our new regularization method in two simple use cases. Our silhouettes experiment shows that,

when the parameters of SK-reg are determined from CNNs trained on a similar image domain to that of the new task, the performance increase that results in the new task can be quite substantial—as large as 55% over an L2 baseline. Our Tiny ImageNet experiment demonstrates that SK-reg is capable of encoding generalizable structural principles about the correlations in receptive fields; the statistics of learned parameters in one domain can be useful in a completely new domain with substantial differences.

The Gaussians that we fit to kernel data in phase 1 of our experiments could be overfit to the CNN training runs. We have discussed the application of sparse inverse covariance (precision) estimation as one approach to reduce over-fitting. In future work, we would like to explore a Gaussian model with graphical connectivity that is specified by a 2D grid MRF. Model fitting would consist of optimizing the non-zero precision matrix values subject to this pre-specified sparsity. The grid MRF model is enticing for its potential to serve as a general “smoothness” prior for CNN receptive fields. Ultimately, we hope to develop a general-purpose kernel regularizer that does not depend on transfer learning.

Although a Gaussian can model some kernel families sufficiently, other families would give it a difficult time. The first-layer kernels of AlexNet—which are 11×11 and are visually similar to Gabor wavelets and derivative kernels—are not well-modeled by a multivariate Gaussian. A more sophisticated prior is needed to model kernels of this size effectively. In future work, we hope to investigate more complex families of priors that can capture the regularities of filters such as Gabors and derivatives. Nevertheless, a simple Gaussian estimator works well for smaller kernels, and in the literature, it has been shown that architectures with a hierarchy of smaller convolutions followed by nonlinearities can achieve equal (and often better) performance as those with fewer, larger kernels (Simonyan & Zisserman, 2015). Thus, the ready-made Gaussian regularizer we introduced here can be used in many applications.

Acknowledgements

We thank Nikhil Parthasarathy, Emin Orhan and Brian McFee for their valuable comments. Reuben Feinman is supported by a Google PhD Fellowship in Computational Neuroscience.

References

Atanov, A., Ashukha, A., Struminsky, K., Vetrov, D., & Welling, M. (2019). The deep weight prior. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*.

Bloom, P. (2000). *How children learn the meanings of words*. Cambridge, MA: MIT Press.

Bruna, J., & Mallat, S. (2013). Invariant scattering convolution networks. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 35(8), 1872–1886.

Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., & Darrell, T. (2014). A deep convolutional activation feature for generic visual recognition. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*.

Feinman, R., & Lake, B. M. (2018). Learning inductive biases with simple neural networks. In *Proceedings of the 40th Annual Conference of the Cognitive Science Society (CogSci)*.

Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*.

Friedman, J., Hastie, T., & Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3), 432–441.

Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41(6), 711–724.

Girshick, R. (2015). Fast R-CNN. In *Proceedings of the International Conference on Computer Vision (ICCV)*.

Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Grant, E., Finn, C., Levine, S., Darrell, T., & Griffiths, T. (2018). Recasting gradient-based meta-learning as hierarchical Bayes. In *Proceedings of the 6th International Conference on Learning Representations (ICLR)*.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J. Physiology*, 160, 106–154.

Jones, J. P., & Palmer, L. A. (1987). An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *J. Neurophysiology*, 58, 1233–1258.

Lake, B. M., & Tenenbaum, J. B. (2010). Discovering structure by learning sparse graphs. In *Proceedings of the 32nd Annual Conference of the Cognitive Science Society (CogSci)*.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, E253.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.

Li, S. Z. (2009). *Markov random field modeling in image analysis*. New York, NY: Springer-Verlag.

Razavian, A. S., Azizpour, H., Sullivan, J., & Carlsson, S. (2014). CNN features off-the-shelf: An astounding baseline for recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Serre, T., Wolf, L., Bileschi, S., Riesenhuber, M., & Poggio, T. (2007). Robust object recognition with cortex-like mechanisms. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 29(3), 411–426.

Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.

Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., & Wierstra, D. (2016). Matching networks for one shot learning. In *Advances in Neural Information Processing Systems (NIPS)*.

Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, 114(2), 245–272.