

Blockout: Dynamic Model Selection for Hierarchical Deep Networks

Calvin Murdock
Carnegie Mellon University
cmurdock@cs.cmu.edu

Zhen Li, Howard Zhou, Tom Duerig
Google Research
{zhenli, howardzhou, tduerig}@google.com

Abstract

Most deep architectures for image classification—even those that are trained to classify a large number of diverse categories—learn shared image representations with a single model. Intuitively, however, categories that are more similar should share more information than those that are very different. While hierarchical deep networks address this problem by learning separate features for subsets of related categories, current implementations require simplified models using fixed architectures specified via heuristic clustering methods. Instead, we propose *Blockout*, a method for regularization and model selection that simultaneously learns both the model architecture and parameters. A generalization of *Dropout*, our approach gives a novel parametrization of hierarchical architectures that allows for structure learning via back-propagation. To demonstrate its utility, we evaluate *Blockout* on the CIFAR and ImageNet datasets, demonstrating improved classification accuracy, better regularization performance, faster training, and the clear emergence of hierarchical network structures.

1. Introduction

Multi-class classification is an important problem in visual understanding with applications ranging from image retrieval to robot navigation. Due to the vast space of variability, the seemingly simple task of identifying the subject of a photograph is extremely difficult. While once restricted to small label sets and constrained image domains, recent advances in deep neural networks have allowed image recognition to be applied to real-world collections of photographs. Effective image classification with thousands of labels and datasets with millions of images are now commonplace. However, as classification tasks become more involved, larger networks with more capacity are required, emphasizing the importance of careful model selection. While much manual engineering effort has been dedicated to the task of designing deep architectures that are able to effectively generalize from available training data, model selection is typically performed using subjective heuristics by experienced practitioners. Furthermore, an appropriate ar-




Figure 1: Example deep network architectures for multi-class classification that can be learned using Blockout. Input and output nodes are shown in green, groups of layers are shown in blue, and arrows indicate connections between them. (a) Traditional architectures make use of a combined model that computes a single feature vector for predicting a large number of diverse categories. (b) Hierarchical architectures instead partition the output categories into clusters and learn separate, high-level feature vectors for each of them. (c) Unlike previous approaches, Blockout allows for end-to-end learning of more complex hierarchical architectures.

chitecture is closely tied to the dataset on which it is trained, so this work often must be repeated for each new application. Ideally, model selection should be performed automatically, allowing the architecture to adapt to training data. As a step towards this goal, we propose an automated, end-to-end system for model selection within the class of hierarchical deep networks, which have demonstrated excellent performance on large-scale image classification tasks.

Deep neural networks are known to be organized such that specificity increases with depth [15, 29]. Lower layers tend to represent general, low-level image features like lines and edges while higher layers encode higher-level concepts like object parts or even objects themselves [4]. Most classification architectures make use of a single shared model with a flat logistic loss layer. Intuitively, however, categories that are more similar should share more information than those that are very different.

One solution is to train independent fine-grained models for these subsets of related labels. This results in specialized features that are tuned to differentiating between subtle visual differences between similar categories. However, this is often infeasible due to limited training examples. On the other hand, a combined model for classifying many categories is able to use information common to all training images to learn shared, low-level representations.

Ideally, then, model architectures should be hierarchical. Low-level representations should be shared while higher layers should be separated out and connected only to subsets of classes, allowing for efficient information sharing and reduced training data requirements. However, this raises an obvious question of model selection: which hierarchical architecture is best? Figure 1 visualizes some potential candidates. Design choices include: the number and locations of branches, the allocation of nodes to each branch, and the clustering of classes in the final layer. Previous approaches to hierarchical deep networks (e.g. [28, 27]) have simplified this question by fixing the base architecture and using heuristic clustering methods for separating the classes into groups. While class similarity may provide an effective heuristic for model selection, it is not guaranteed to actually improve performance and ignores important factors such as heterogeneous classification difficulty.

To achieve automatic model selection in hierarchical deep networks, we introduce Blockout, an approach for simultaneously learning both the model architecture and parameters. This allows for more complex hierarchical architectures specifically tuned to the data without requiring a separate procedure for model selection, which would likely be infeasible due to the vast search space of possible architectures. Inspired by Dropout [9], Blockout can be viewed as a technique for stochastic regularization that adheres to hierarchically-structured model architectures. Importantly, its hyper-parameters (analogous to node Dropout probabilities) are represented such that they can be learned using simple back-propagation. Thus, Blockout performs a relaxed form of model selection by effectively learning an ensemble of hierarchical networks, i.e. the distribution of hierarchical architectures over which to average during inference. Despite the additional parameters, the representational power of Blockout is exactly the same as a standard layer and can be parametrized as such during inference. Surprisingly, however, the resulting network is able to achieve improved performance, as demonstrated experimentally on standard image classification datasets.

In summary, we make the following contributions: (1) a novel parametrization of hierarchical deep networks, (2) stochastic regularization analogous to Dropout that effectively averages over all models within this class of hierarchical architectures, (3) an approach for learning the regularization parameters allowing for architectures that dynamically adapt to the data throughout training, and (4) quantitative and qualitative analyses, including substantial performance gains over baseline models.

2. Related Work

Despite the long history of deep neural networks in computer vision [14], the modern incarnation of “deep learning” is a relatively recent phenomenon that began with empirical

success in the task of image recognition [13] on the ImageNet dataset [17]. Since then, tactful architecture modifications have yielded a steady stream of further improvements [30, 23], even surpassing human performance [8].

In addition to general classification of arbitrary images, deep learning has also made a significant impact on fine-grained recognition within constrained domains [3, 11, 16]. In these cases, deep neural networks are trained (often alongside additional annotations or segmentations of parts) to recognize subtle differences between similar categories, e.g. bird species. However, these methods are often limited by the availability of training data as they typically require expert annotations for ground truth labels. Some approaches have alleviated this problem by pre-training on large collections of general images and then fine-tuning on smaller, domain-specific datasets [16].

Attempts have also been made to incorporate information from a known hierarchy to improve prediction performance without requiring architecture changes. For example, [5] replaced the flat softmax classification layer with a probabilistic graphical model that respects given relationships between labels. Other methods for incorporating label structure are summarized in [24]. However, they typically rely on fixed, manually-specified hierarchies, which could contain errors and result in biases that reduce performance.

Hierarchical deep networks [27, 28] attempt to address these issues by learning multi-task models with shared lower layers and parallel, domain-specific higher layers for predicting different subsets of categories. While these methods address one component of model selection by learning clusters of output categories, other architectural hyper-parameters such as the location of branches and the relative allocation of nodes between them must still be specified prior to training.

The most common approach for model selection in deep learning is simply searching over the space of hyper-parameters [2]. Unfortunately, because training and inference in deep networks are computationally expensive, this is often impractical. While costs can sometimes be reduced (e.g. by taking advantage of the behavior of some network architectures with random weights [18]), they still require training and evaluating a large number of models. Bayesian optimization approaches [20] attempt to perform this search more efficiently, but they are still typically applied only to smaller models with few hyper-parameters. Alternatively, [1] proposed a theoretically-justified approach to learning a deep network with a layer-wise strategy that automatically selects the appropriate number of nodes during training. However, it is unclear how it would perform on large-scale image classification benchmarks.

A parallel but related task to model selection is regularization. A network with too much capacity (e.g. with too many parameters) can easily overfit without sufficient

training data, resulting in poor generalization performance. While the size of the model could be reduced, an easier and often more effective approach is to use regularization. Common methods include imposing constraints on the weights (e.g. through convolution or weight decay), rescaling or whitening internal representations for better conditioning [6, 10], or randomly perturbing activations for improved robustness and better generalizability [9, 26, 25].

3. Deep Neural Networks

Deep neural networks are layered nonlinear functions $f : \mathbb{R}^d \rightarrow \mathbb{R}^p$ that take d -dimensional images as input and output p -dimensional predictions. They have been found to be very successful for image classification, most likely due to the complexity of the class of representable functions along with their ability to effectively and efficiently make use of very large sets of training data.

Most deep neural networks are simply compositions of alternating linear and nonlinear functions. More concretely, consider a deep network with m layers. Each layer consists of a linear transformation $g_j(x) = \mathbf{W}_j x$ parametrized by \mathbf{W}_j followed by a fixed nonlinear function $a_j(x)$, e.g. a nonlinear activation or a pooling operator. Altogether, the full neural network can be represented as:

$$f = a_m \circ g_m \circ a_{m-1} \circ g_{m-1} \circ \dots \circ a_1 \circ g_1 \quad (1)$$

Similarly, hierarchical deep networks can be expressed with a separate function f for each subset of outputs where some intermediate representations a_j are shared, i.e. they can be used as the inputs to multiple layers.

The set of all model parameters $\mathbf{W} = \{\mathbf{W}_j\}$ can be learned from a dataset of n training images x_i and corresponding ground-truth label vectors y_i using standard empirical risk minimization with a loss function L (e.g. softmax) that measures the discrepancy between y_i and the network predictions $f(x_i; \mathbf{W})$, as shown in Equation 2:

$$\arg \min_{\mathbf{W}} \frac{1}{n} \sum_{i=1}^n L(y_i, f(x_i; \mathbf{W})) \quad \text{s.t. } \{\mathbf{W}_j \in \mathcal{S}_j\} \quad (2)$$

Learning is typically accomplished through stochastic gradient descent, where the gradients of intermediate layers are computed using back-propagation.

Consistent with their name, deep neural networks typically consist of many layers that produce high-dimensional intermediate representations, resulting in an extremely large number of parameters to be learned. To prevent overfitting, regularization is typically employed through constraint sets \mathcal{S}_j on the parameter matrices. The most common and effective form of regularization is convolution, which takes advantage of the local correlations of images and essentially restricts that the weight matrices contain shared parameters with a specific Toeplitz structure, resulting in far




Figure 2: An illustration of the equivalence between single layers with block-structured parameter matrices (top) and parallel layers over subsets of nodes (bottom). Solid boxes indicate groups of nodes, dotted boxes represent the corresponding parameter matrices (where zero values are shown in black), and colors indicate cluster membership. (a) Independence between layers can be achieved when nodes only belong to a single cluster. When nodes belong to multiple clusters, hierarchical connections such as merging (b) and branching (c) can be achieved.

fewer free parameters to learn. Other examples of regularization include weight decay, which penalizes the norm of the weights, and Dropout, which has been shown (under certain assumptions) to indirectly impose a penalty function through stochastic perturbations of the internal network activations [25]. Blockout employs a similar form of stochastic regularization with the additional restriction that the parameter matrices be block-structured leading to hierarchical network architectures.

4. Hierarchical Network Parametrization

Blockout is based on the observation that parallel, independent layers can be equivalently expressed as a single combined layer with a block-structured weight matrix (up to a permutation of its rows and columns), as visualized in Figure 2. Thus, enforcing that the learned weight matrix have this type of structure during training automatically separates the input and output nodes into independent branches of a hierarchical architecture.

This can be parametrized by assigning each node to any number of k clusters and masking out parameters if their corresponding input and output nodes do not belong to the same cluster, thus restricting the information that can be shared between nodes. Here, k represents the maximum number of blocks in the parameter matrix or, equivalently, the maximum number of independent branches in the network. Though simple, this parametrization can encode a wide range of hierarchical structures, as shown in Figure 3.

More formally, we mask the parameter corresponding to s^{th} input node and the t^{th} output node as follows in Equation 3, where $\tilde{w}_{t,s}$ is the original, unconstrained parameter




Figure 3: A summary of the types of basic high-level architecture components that can be represented with Blockout. For each (a-e), a single layer is shown where groups nodes are shown as solid boxes, cluster assignments as colors, and connections within clusters as arrows. These connections allow for a rich space of potential model architectures.

value and $\mathbb{I}(s \in C_l)$ equals one if node s belongs to cluster l and zero otherwise:

$$w_{t,s} = \frac{1}{k} \sum_{l=1}^k \mathbb{I}(s \in C_l) \mathbb{I}(t \in C_l) \tilde{w}_{s,t} \quad (3)$$

This encodes the desired behavior that a parameter be nonzero only if its corresponding input and output nodes belong to the same class while restricting that the mask be between zero and one. Let $\mathbf{C}_j \in \{0, 1\}^{d_j \times k}$ be a binary indicator matrix containing these cluster membership assignments for each of the d_j nodes in the output of the j^{th} layer. In other words, $\mathbf{C}_j(s, l) = \mathbb{I}(s \in C_l)$. A full mask can then be constructed as $\frac{1}{k} \mathbf{C}_j \mathbf{C}_{j-1}^\top$ where the block-structured parameter matrix is the element-wise product of an unconstrained parameter matrix $\tilde{\mathbf{W}}_j$ and this mask. This class of hierarchical architectures can be summarized by the constraint set in Equation 4, where \odot indicates the element-wise Hadamard product.

$$\mathcal{S}_j = \left\{ \mathbf{W}_j : \mathbf{W}_j = \frac{1}{k} \tilde{\mathbf{W}}_j \odot \mathbf{C}_j \mathbf{C}_{j-1}^\top \right\} \quad (4)$$

These constraints act as a regularizer that enforces the parameter matrices to be block-structured with potentially many parameters set explicitly to zero. Ideally, we seek to learn the hierarchical structure during training, which is equivalent to learning the cluster membership assignments \mathbf{C}_j . However, because they are binary variables, learning them directly would be difficult. To address this problem, we instead take an approach akin to stochastic regularization approaches like Dropout: we treat cluster membership assignments as Bernoulli random variables and draw a different hierarchical architecture at each training iteration.

5. Stochastic Regularization

Stochastic regularization techniques are simple but effective approaches for reducing overfitting in deep networks by injecting noise into the intermediate activations or parameters during training. Examples include Dropout [9], which randomly sets activations to zero, and DropCon-




Figure 4: Example parameter masks that can be achieved with (a) Dropout, (b) Blockout, and (c) DropConnect. Note that Dropout and Blockout give block-structured, low-rank masks up to a permutation of the rows and columns while DropConnect is structureless, masking each parameter value independently.

nect [26], which randomly sets parameter values to zero.

Dropout [9] works by setting node activations to zero with a certain probability at each training iteration. Inference is accomplished by replacing each activation with its expected value, which amounts to rescaling by the Dropout probability. This procedure approximates an ensemble of different models from the class of network architectures containing all possible subsets of nodes, where the Dropout probability determines the weight given to each architecture in this implicit model average. For example, with a high Dropout probability, models with fewer nodes are more likely to be selected during training. In general, Dropout results in improved generalization performance by preventing the coadaptation of features.

Similarly, DropConnect [26] randomly sets parameter values to zero, which drops connections between nodes instead of the node activations themselves. During inference, a moment-matching procedure is used to better approximate an average over model architectures. Again, the success of this approach can be explained through its approximation of an ensemble within a much larger class of architectures: those that contain all possible combinations of connections between nodes in the network.

Blockout can be seen as another example of stochastic regularization that approximates an ensemble of models from the class of hierarchical architectures introduced in Section 4. Structured noise is introduced by randomly selecting cluster assignments \mathbf{C}_j corresponding to different hierarchical architectures at each iteration during training. We first consider the case of a single fixed probability p that each node belongs to each of the clusters, but in Section 6 we show how separate cluster probabilities can be learned for each node.

During inference, we take an approach similar to Dropout and approximate an ensemble of hierarchical architectures using an implicit average with weights determined by the cluster probabilities. This again amounts to simply rescaling the parameter values by the expected value of the parameter mask: p^2 .

Also note that Dropout can be interpreted as implicitly applying a random mask \mathbf{M} that sets parameters corresponding to the dropped inputs and outputs to zero. If we reorder the input and output dimensions and permute

the rows and columns of the weight matrix accordingly, the result is a single block of non-zero parameters, as shown in Figure 4a. This is very similar to the block-structured masks that can be explicitly represented with Blockout, as shown in Figure 4b. In fact, Dropout is equivalent to Blockout with $k = 1$ where dropped nodes correspond to those that do not belong to the single cluster. In this case, the resulting mask is a rank-one matrix. Similarly, the explicit parameter masks in DropConnect (shown in Figure 4c) are full-rank and can be equivalently represented by Blockout when the number of clusters is equal to the number of nodes in a layer. This allows each node to potentially belong to its own independent cluster resulting in a full-rank mask.

The full intuition behind why stochastic regularization approaches work and how best to select regularization hyper-parameters (e.g. Dropout probability) is lacking. On the other hand, Blockout gives a much clearer motivation: we assume that the output categories are hierarchically related, and so we approximate an ensemble only over hierarchical architectures. Furthermore, in Section 6 we show how the cluster probabilities for each node can be learned from data allowing for the interpretation of Blockout as model selection within this class of architectures.

6. Learning Hierarchies via Back-Propagation

The key difference between Blockout and other stochastic regularization techniques is that its hyper-parameters can be learned from data using simple back-propagation. To accomplish this, we replace the fixed, shared cluster probability p with learnable parameters $\mathbf{P}_j \in [0, 1]^{d_j \times k}$ whose elements represent the probability that each node belongs to each cluster. Essentially, they are relaxations of the binary cluster assignments \mathbf{C}_j that can take on any value between zero and one and are implemented as real-valued variables followed by element-wise logistic activations. At each iteration of training, hard binary cluster assignments are drawn from Bernoulli distributions parametrized by these probabilities, i.e. $\mathbf{C}_j \sim \mathbf{B}(1, \mathbf{P}_j)$.

During training, the forward computations are performed using random masked weight matrices from the set in Equation 4 for a different hierarchical architecture at each iteration. During inference, we again average over the cluster assignments to approximate an ensemble of hierarchical architectures. Since the cluster probabilities \mathbf{P}_j are now different for each node, we must rescale each parameter accordingly. Specifically, the adjusted weight matrix used during inference is:

$$\mathbb{E} \left[\frac{1}{k} \widetilde{\mathbf{W}}_j \odot \mathbf{C}_j \mathbf{C}_{j-1}^\top \right] = \frac{1}{k} \widetilde{\mathbf{W}}_j \odot \mathbf{P}_j \mathbf{P}_{j-1}^\top \quad (5)$$

Note that this leads to the same computation as that of the training forward pass except with \mathbf{P}_j instead of \mathbf{C}_j . Thus, during inference, we simply skip the random cluster assign-

ment step and use the soft clustering probabilities directly.

The masked parameter matrix is represented as a function of three variables: the unconstrained weight matrix $\widetilde{\mathbf{W}}_j$, the input cluster assignments \mathbf{C}_{j-1} , and the output cluster assignments \mathbf{C}_j . As such, gradients can be passed to all of them following the typical back-propagation algorithm. Specifically, updating \mathbf{W}_j (e.g. using stochastic gradient descent) requires computation of the gradient of the loss function with respect to those parameters. Using the expression of a deep neural network as a composition of functions from Equation 1, this can be expressed using the chain rule as follows:

$$\frac{\partial L}{\partial \mathbf{W}_j} = \frac{\partial L}{\partial \mathbf{a}_m} \frac{\partial \mathbf{a}_m}{\partial \mathbf{g}_m} \dots \frac{\partial \mathbf{g}_{j+1}}{\partial \mathbf{a}_j} \frac{\partial \mathbf{a}_j}{\partial \mathbf{g}_j} \frac{\partial \mathbf{g}_j}{\partial \mathbf{W}_j} = \delta_j \frac{\partial \mathbf{g}_j}{\partial \mathbf{W}_j} \quad (6)$$

where δ is the product of all gradients from the loss function backwards down to the j^{th} layer. Using simple linear algebra, the gradients with respect to a layer's input (i.e. the previous layers activations \mathbf{a}_{j-1}) are:

$$\frac{\partial \mathbf{g}_j}{\partial \mathbf{a}_{j-1}} = \mathbf{W}_j = \frac{1}{k} \widetilde{\mathbf{W}}_j \odot \mathbf{C}_j \mathbf{C}_{j-1}^\top \quad (7)$$

Similarly, the gradients with respect to all components of the weight matrix are computed as:

$$\begin{aligned} \frac{\partial L}{\partial \mathbf{W}_j} &= \delta_j \mathbf{a}_{j-1}^\top, \quad \frac{\partial L}{\partial \widetilde{\mathbf{W}}_j} = \frac{1}{k} \frac{\partial L}{\partial \mathbf{W}_j} \odot \mathbf{C}_j \mathbf{C}_{j-1}^\top, \quad (8) \\ \frac{\partial L}{\partial \mathbf{C}_j} &= \frac{1}{k} \left[\widetilde{\mathbf{W}}_j \odot \frac{\partial L}{\partial \mathbf{W}_j} \right]^\top \mathbf{C}_{j-1} + \frac{1}{k} \left[\widetilde{\mathbf{W}}_{j+1} \odot \frac{\partial L}{\partial \mathbf{W}_{j+1}} \right] \mathbf{C}_{j+1} \end{aligned}$$

Note that the cluster assignments for a set of nodes \mathbf{C}_j are shared between the two adjacent Blockout layers and hence its gradient contains components from each, acting as an additional form of regularization.

Recall that our goal is to learn the cluster probabilities \mathbf{P}_j that parametrize the cluster assignment random variables \mathbf{C}_j . Thus, to update the cluster probabilities, we simply use the cluster assignment gradients after masking them so that the gradients of unselected clusters are zero:

$$\frac{\partial L}{\partial \mathbf{P}_j} = \frac{\partial L}{\partial \mathbf{C}_j} \odot \mathbf{C}_j \quad (9)$$

This is similar to the technique used when back-propagating gradients through a Dropout layer. Finally, to update the real-valued cluster parameters, these gradients are then back-propagated through the logistic activation layer. The full training process is summarized in Algorithm 1.

Modifying Equation 2, our final optimization problem can thus be written as follows:

$$\begin{aligned} \arg \min_{\widetilde{\mathbf{W}}, \mathbf{P}} & \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{C} \sim \mathbf{B}(1, \mathbf{P})} L(\mathbf{y}_i, f(\mathbf{x}_i; \mathbf{W})) \\ \text{s.t. } & \mathbf{W}_j = \frac{1}{k} \widetilde{\mathbf{W}}_j \odot \mathbf{C}_j \mathbf{C}_{j-1}^\top \end{aligned} \quad (10)$$

Algorithm 1 Blockout Training Iteration

Input: Mini-batch of training images $\{\mathbf{x}_i\}_{i=1}^B$, parameters from previous iteration $\widetilde{\mathbf{W}}_j^{(t-1)}, \mathbf{P}_j^{(t-1)}$

Output: Updated parameters $\widetilde{\mathbf{W}}_j^{(t)}, \mathbf{P}_j^{(t)}$

Forward Pass:

- Draw cluster assignments: $\mathbf{C}_j \sim B(1, \mathbf{P}_j^{(t-1)})$
- Mask parameters: $\mathbf{W}_j = \frac{1}{k} \widetilde{\mathbf{W}}_j^{(t-1)} \odot \mathbf{C}_j \mathbf{C}_j^\top$
- Compute predictions: $\hat{\mathbf{y}}_i = f(\mathbf{x}_i; \mathbf{W}_j)$
- Evaluate empirical risk: $\frac{1}{B} \sum_{i=1}^B L(\mathbf{y}_i, \hat{\mathbf{y}}_i)$

Backward Pass:

- Compute gradients according to Equations 8 and 9.
 - Update parameters $\widetilde{\mathbf{W}}_j^{(t)}, \mathbf{P}_j^{(t)}$ accordingly.
-

For our implementation, the cluster probabilities \mathbf{P}_j are initialized to 0.5. Throughout training, there are a number of possible outcomes: (1) The probabilities could diverge, some towards one and others towards zero. This would result in a fixed clustering of nodes, giving high confidence to a particular learned hierarchical structure. (2) Alternatively, the gradients could be uninformative, averaging to zero and leading to unchanged probabilities. This could indicate that hierarchical architectures are helpful for regularization, but the particular grouping of nodes is arbitrary. (3) The probabilities could also all increase towards one, possibly demonstrating that hierarchical architectures are not beneficial and better performance could be achieved with single, fully-connected layers.

7. Experimental Results

To evaluate our approach, we apply Blockout to the standard image classification datasets CIFAR [12] and ImageNet [17]. As baselines, we use variations of the Inception architecture [23]. Specifically, for ImageNet we use the same model described in [10], and for CIFAR we use a compacted version of this model with fewer layers and parameters. We also follow the same training details described in [10] with standard data augmentation. These models have been hand-engineered to achieve very good, near state-of-the-art performance by themselves. Thus, our intention is to show how the addition of Blockout layers can easily improve performance without involved hyper-parameter tuning. Furthermore, we show that Blockout does indeed learn hierarchical network structures resulting in higher prediction accuracy and faster training.

Inception architectures are composed of multiple layers of parallel convolutional operations directly followed by softmax classification. However, it has been shown that fully-connected layers before classification act as a form of orderless pooling [16] and have been used extensively [13, 30, 19] demonstrating improved model capacity leading to better performance. Thus, we add two fully-

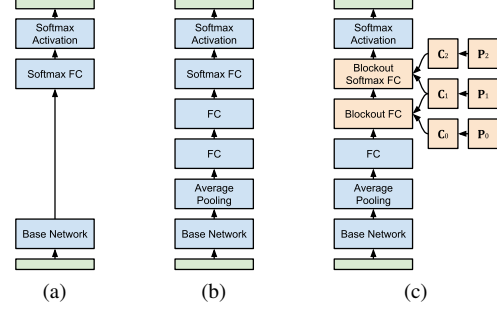


Figure 5: Block diagrams of the models compared in our experiments. (a) As baselines, we use variants of the Inception convolutional neural network architecture [23]. (b) For comparison, we add an average pooling layer to reduce the bottleneck size followed by two fully-connected layers (potentially with Dropout) before the softmax classifier. (c) Our model replaces the last two FC layers with Blockout layers of the same size.

connected layers after the convolutional layers of our base architectures. Because our baselines already have high network capacity, doing this naively can lead to extreme overfitting and reduced performance, even with standard regularization techniques such as Dropout. However, using Blockout prevents this overfitting and leads to substantial performance improvements with a wide range of hyper-parameter choices. The architectures compared in our experiments are shown in Figure 5.

To demonstrate the effectiveness of the different components of Blockout, we compare three variations of our proposed model. The first, indicated by (soft, learned) in the following experiments, omits random cluster selection by skipping the Bernoulli sampling step. This effectively removes the stochastic regularization effect of Blockout, instead using the relaxed soft clustering assignments directly by setting $\mathbf{C}_j = \mathbf{P}_j$. Without explicit zero-valued parameters, the same number of effective parameters are learned as with ordinary fully-connected layers, which could still potentially lead to over-fitting. However, the additional regularization provided by the shared cluster parameters often mitigates this, still resulting in improved performance. The second (hard, fixed) uses randomized hard cluster assignment during training, but uses fixed cluster probabilities of 0.5 instead of back-propagating gradients as described in Section 6. This shows the effects of stochastic regularization within the class of hierarchical architectures. Finally, the third (hard, learned) is our full proposed model.

7.1. CIFAR-100

CIFAR-100 is a challenging dataset comprising of 60k 32x32 color images (50k for training and 10k for testing) equally divided into 100 classes [12]. Table 1 shows the performance of our model with 6 clusters and 512 nodes in each fully-connected layer. We compare against the baseline Inception model, the baseline with fully-connected lay-

Table 1: **Cifar-100 Test Accuracy.** Left: Comparison with baseline methods. Right: Variable clusters with 512 nodes (top) and variable nodes with 6 clusters (bottom).

Method	Acc. (%)	Clusters	Acc. (%)
Baseline	61.56	2	64.54
Baseline + FC	62.66	4	65.93
Baseline + FC + Dropout	64.32	6	65.66
Blockout (soft, learned)	63.57		
Blockout (hard, fixed)	65.62		
Blockout (hard, learned)	65.66		

Nodes	Acc. (%)
512	65.66
1024	66.69
2048	66.71




Figure 6: The convergence of our models in comparison to the baselines on the CIFAR-100 dataset, showing accuracy on the (a) training and (b) testing sets throughout training. Note that Blockout converges in about half the time as Dropout while still achieving a higher final accuracy.

ers, and the baseline with fully-connected layers followed by 30% Dropout. Figure 6 shows the accuracy of these models throughout training, demonstrating faster convergence in comparison to Dropout. Table 1 also compares our full Blockout model (hard, learned) with a variety of hyperparameter selections, including the number of hidden nodes in each fully-connected layer and the number of clusters.

The best performance achieved by our method gave an accuracy of 66.71% with 6 clusters and 2048 nodes, showing a significant improvement over the baseline accuracy. Also note that, while other stochastic regularization methods like Dropout can still overfit if there are too many parameters or the Dropout probability is not set correctly, Blockout seems to adapt so that adding more nodes never reduces accuracy. Despite its minimal engineering effort, the results are comparable to state-of-the-art methods (e.g. 68.8% with [7], 67.76% with [22], 66.29% with [21], etc.)

7.2. ImageNet

ImageNet is the standard dataset for large-scale image classification [17]. We use the version of the dataset from the Imagenet Large Scale Visual Recognition Challenge (ILSVRC 2012), which has 1000 object categories, 1.2 million training images, and 50k validation images. Table 2 shows the top-1 prediction performance of our model with 6 clusters and 4096 nodes in each fully-connected layer. We compare to the baseline, the baseline with fully-connected layers, and the baseline with fully-connected layers followed by 50% Dropout. Because the baseline model was already carefully tuned to maximize performance on Ima-

Table 2: **ImageNet Evaluation Accuracy.** Left: Comparison with baseline methods. Right: Variable clusters with 4096 nodes (top) and variable nodes with 6 clusters (bottom).

Method	Acc. (%)	Clusters	Acc. (%)
Baseline	73.43 ¹	2	73.78
Baseline + FC	68.06	6	74.83
Baseline + FC + Dropout	73.88	15	74.19
Blockout (soft, learned)	72.43		
Blockout (hard, fixed)	74.44		
Blockout (hard, learned)	74.83		

Nodes	Acc. (%)
1024	74.16
2048	74.47
4096	74.83
8192	74.95

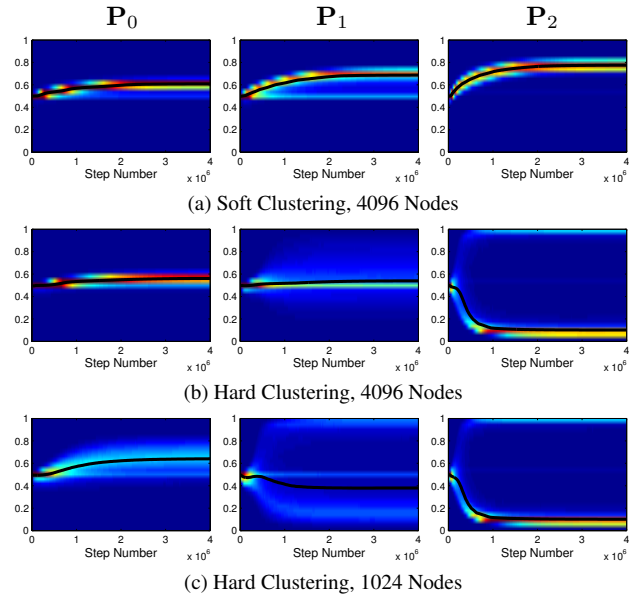


Figure 7: A visualization of the distributions of each layer’s cluster probabilities P_j throughout training on the ImageNet dataset. The iteration number varies along the x-axis with probability along the y-axis. Warmer colors indicate a higher density of cluster probabilities at a given iteration while the black line shows their median. With hard clustering, there is a clear separation towards higher confidence cluster assignments, especially in later layers.

geNet, adding fully-connected layers resulted in significant overfitting that could not be overcome with Dropout. However, Blockout was able to effectively remove these effects giving an improved final maximum performance of 74.95%.

Figure 7 shows the distribution of the learned cluster probabilities throughout training. Without random cluster selection, soft clustering causes all probabilities to increase towards one, which could indicate overfitting to the training data. On the other hand, stochastic regularization with hard clustering results in diverging probabilities giving higher confidence cluster membership assignments. This effect is

¹This is the accuracy of our implementation of the model in [10], which reported a maximum accuracy of 74.8%. This discrepancy is most likely due to a different learning rate decay schedule.

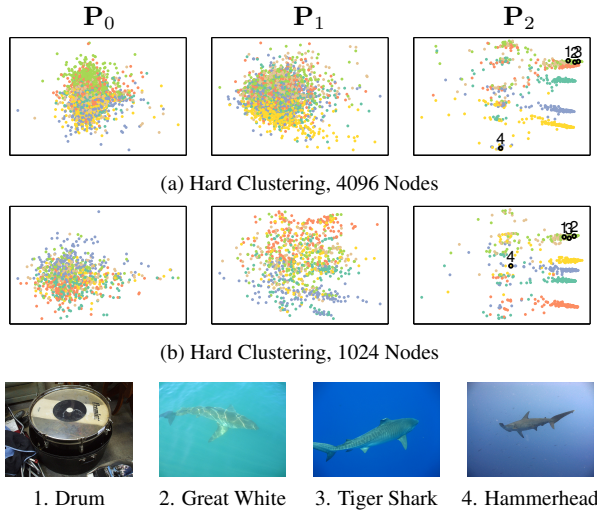


Figure 8: (Top) A visualization of the node cluster probabilities P_j projected to two dimensions using PCA for models with (a) 4096 nodes and (b) 1024 nodes. Dots indicate nodes while color indicates the cluster with the highest probability. Some example output categories (1-4) are also shown. (Bottom) The associated categories along with sample images. Despite the somewhat non-intuitive structure, there are clear, consistent groupings of nodes, especially in later layers and with fewer nodes.

also more prevalent in higher layers, agreeing with our intuition that more information should be shared in lower layers.

Figure 8 visualizes the k -dimensional cluster probability vectors for each node by projecting them to two dimensions using PCA. Because nodes can belong to multiple clusters with varying relative frequencies, the cluster probabilities can be interpreted as embeddings where nodes with similar probabilities indicate computations following similar paths in the hierarchical architecture. Again note that earlier layers tend to be less separated, especially with higher network capacity, allowing for more information to be shared between clusters. In addition, because this implicit node embedding is a side effect of maximizing prediction accuracy, the resulting clusters are less interpretable than an explicit clustering based on category or image similarity. For example, while nearly indistinguishable classes such as “great white shark” and “tiger shark” do share very similar cluster probabilities, so does the visually dissimilar and seemingly unrelated class “drum.” Furthermore, while one might expect “hammerhead shark” to be close to the other sharks, it actually belongs to a completely different set of clusters. Despite this, the final cluster probabilities do seem to be fairly consistent across different choices of hyper-parameters. This could indicate that the node clusters are indeed a function of the training data, but incorporate more information than just visual similarity.

Figure 9 shows the expected number of clusters assigned to each output category. Again, notice the consistency across different hyper-parameter selections. While the me-

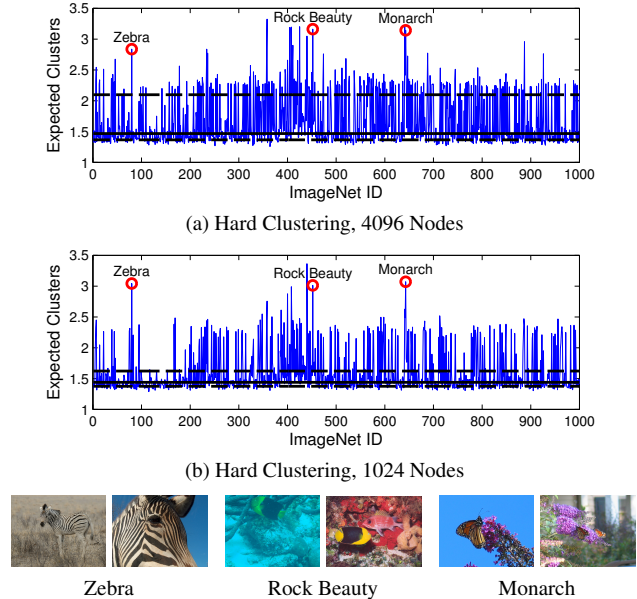


Figure 9: (Top) The expected number of clusters assigned to each of the ImageNet output categories for (a) 4096 nodes and (b) 1024 nodes. The solid black line shows the median number of clusters while the dotted black lines show the 25th and 75th percentiles. Also shown are 3 example categories with a relatively high expected number of clusters. (Bottom) Sample images from the indicated categories, showing classification challenges such as camouflage, varied background appearance, and small relative size.

dian number of clusters is around 1.5, some categories belong to close to 3 with many more parameters used in their predictions. These could correspond to categories that are more difficult to predict, perhaps due to natural camouflage (e.g. “zebra”), large variations in background appearance (e.g. “rock beauty”), or the small relative size of the subject (e.g. “monarch butterfly”).

8. Conclusion

Blockout is a novel generalization of stochastic regularization with parameters that can be learned during training, essentially allowing for automatic model selection within a class of hierarchical network structures. While our approach is not guaranteed to learn exact block-structured weight matrices, we demonstrated experimentally that Blockout consistently converges to an implicit clustering of the output categories with branches sharing similar representations. While further work is required to completely understand and interpret the learned clusters, Blockout results in substantial improvements in prediction accuracy and faster convergence in comparison to baseline methods. As a first step towards fully-automatic model selection, Blockout emphasizes the importance of the careful parametrization of deep network architectures and should inspire a family of similar approaches adapted to other application domains.

References

- [1] S. Arora, A. Bhaskara, R. Ge, and T. Ma. Provable bounds for learning some deep representations. In *International Conference on Machine Learning (ICML)*, 2014. 2
- [2] J. Bergstra and Y. Bengio. Random search for hyperparameter optimization. *Journal of Machine Learning Research (JMLR)*, 13(1):281–305, 2012. 2
- [3] S. Branson, G. Van Horn, S. Belongie, and P. Perona. Bird species categorization using pose normalized deep convolutional nets. In *British Machine Vision Conference (BMVC)*, 2014. 2
- [4] A. Coates, A. Karpathy, and A. Y. Ng. Emergence of object-selective features in unsupervised feature learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2012. 1
- [5] J. Deng, N. Ding, Y. Jia, A. Frome, K. Murphy, S. Bengio, Y. Li, H. Neven, and H. Adam. Large-scale object classification using label relation graphs. In *European Conference on Computer Vision (ECCV)*, 2014. 2
- [6] G. Desjardins, K. Simonyan, R. Pascanu, and K. Kavukcuoglu. Natural neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2015. 3
- [7] B. Graham. Fractional max-pooling. *arXiv preprint arXiv:1412.6071*, 2014. 7
- [8] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *International Conference on Computer Vision (ICCV)*, 2015. 2
- [9] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012. 2, 3, 4
- [10] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, 2015. 3, 6, 7
- [11] J. Krause, H. Jin, J. Yang, and L. Fei-Fei. Fine-grained recognition without part annotations. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [12] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images, 2009. 6
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2012. 2, 6
- [14] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 2
- [15] H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *International Conference on Machine Learning (ICML)*, 2009. 1
- [16] T.-Y. Lin, A. RoyChowdhury, and S. Maji. Bilinear cnn models for fine-grained visual recognition. In *International Conference on Computer Vision (ICCV)*, 2015. 2, 6
- [17] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, pages 1–42, April 2015. 2, 6, 7
- [18] A. Saxe, P. W. Koh, Z. Chen, M. Bhand, B. Suresh, and A. Y. Ng. On random weights and unsupervised feature learning. In *International Conference on Machine Learning (ICML)*, 2011. 2
- [19] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2014. 6
- [20] J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems (NIPS)*, 2012. 2
- [21] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. Striving for simplicity: The all convolutional net. In *International Conference on Learning Representations (ICLR)*, 2014. 7
- [22] R. K. Srivastava, K. Greff, and J. Schmidhuber. Training very deep networks. *arXiv preprint arXiv:1507.06228*, 2015. 7
- [23] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2, 6
- [24] A.-M. Tousch, S. Herbin, and J.-Y. Audibert. Semantic hierarchies for image annotation: A survey. *Pattern Recognition*, 45(1):333–345, 2012. 2
- [25] S. Wager, S. Wang, and P. S. Liang. Dropout training as adaptive regularization. In *Advances in Neural Information Processing Systems (NIPS)*, 2013. 3
- [26] L. Wan, M. Zeiler, S. Zhang, Y. L. Cun, and R. Fergus. Regularization of neural networks using dropconnect. In *International Conference on Machine Learning (ICML)*, 2013. 3, 4
- [27] D. Warde-Farley, A. Rabinovich, and D. Anguelov. Self-informed neural network structure learning. In *International Conference on Learning Representations (ICLR)*, 2015. 2
- [28] Z. Yan, V. Jagadeesh, D. DeCoste, W. Di, and R. Piramuthu. Hd-cnn: Hierarchical deep convolutional neural network for image classification. In *International Conference on Computer Vision (ICCV)*, 2015. 2
- [29] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson. How transferable are features in deep neural networks? In *Advances in Neural Information Processing Systems (NIPS)*, 2014. 1
- [30] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *European Conference on Computer Vision (ECCV)*, 2014. 2, 6