

Dataless Model Selection with the Deep Frame Potential

Calvin Murdock¹ Simon Lucey^{1,2}
¹Carnegie Mellon University ²Argo AI
 {cmurdock, slucey}@cs.cmu.edu

Abstract

Choosing a deep neural network architecture is a fundamental problem in applications that require balancing performance and parameter efficiency. Standard approaches rely on ad-hoc engineering or computationally expensive validation on a specific dataset. We instead attempt to quantify networks by their intrinsic capacity for unique and robust representations, enabling efficient architecture comparisons without requiring any data. Building upon theoretical connections between deep learning and sparse approximation, we propose the deep frame potential: a measure of coherence that is approximately related to representation stability but has minimizers that depend only on network structure. This provides a framework for jointly quantifying the contributions of architectural hyper-parameters such as depth, width, and skip connections. We validate its use as a criterion for model selection and demonstrate correlation with generalization error on a variety of common residual and densely connected network architectures.

1. Introduction

Deep neural networks have dominated nearly every benchmark within the field of computer vision. While this modern influx of deep learning originally began with the task of large-scale image recognition [18], new datasets, loss functions, and network configurations have quickly expanded its scope to include a much wider range of applications. Despite this, the underlying architectures used to learn effective image representations are generally consistent across all of them. This can be seen through the community’s quick adoption of the newest state-of-the-art deep networks from AlexNet [18] to VGGNet [28], ResNets [13], DenseNets [15], and so on. But this begs the question: why do some deep network architectures work better than others? Despite years of groundbreaking empirical results, an answer to this question still remains elusive.

Fundamentally, the difficulty in comparing network architectures arises from the lack of a theoretical foundation for characterizing their generalization capacities. Shal-

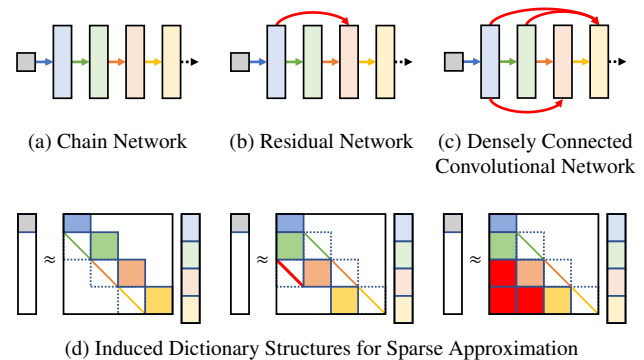


Figure 1: Why are some deep neural network architectures better than others? In comparison to (a) standard chain connections, skip connections like those in (b) ResNets [13] and (c) DenseNets [15] have demonstrated significant improvements in training effectiveness, parameter efficiency, and generalization performance. We provide one possible explanation for this phenomenon by approximating network activations as (d) solutions to sparse approximation problems with different induced dictionary structures.

low machine learning techniques like support vector machines [6] were aided by theoretical tools like the VC-dimension [31] for determining when their predictions could be trusted to avoid overfitting. Deep neural networks, on the other hand, have eschewed similar analyses due to their complexity. Theoretical explorations of deep network generalization [24] are often disconnected from practical applications and rarely provide actionable insight into how architectural hyper-parameters contribute to performance.

Building upon recent connections between deep learning and sparse approximation [26, 23], we instead interpret feed-forward deep networks as algorithms for approximate inference in related sparse coding problems. These problems aim to optimally reconstruct zero-padded input images as sparse, nonnegative linear combinations of atoms from architecture-dependent dictionaries, as shown in Fig. 1. We propose to indirectly analyze practical deep network architectures with complicated skip connections, like residual networks (ResNets) [13] and densely connected convolu-

tional networks (DenseNets) [15], simply through the dictionary structures that they induce.

To accomplish this, we introduce the deep frame potential for summarizing the interactions between parameters in feed-forward deep networks. As a lower bound on mutual coherence—the maximum magnitude of the normalized inner products between all pairs of dictionary atoms [9]—it is theoretically tied to generalization properties of the related sparse coding problems. However, its minimizers depend only on the dictionary structures induced by the corresponding network architectures. This enables dataless model comparison by jointly quantifying contributions of depth, width, and connectivity.

Our approach is motivated by sparse approximation theory [11], a field that encompasses properties like uniqueness and robustness of shallow, overcomplete representations. In sparse coding, capacity is controlled by the number of dictionary atoms used in sparse data reconstructions. While more parameters allow for more accurate representations, they may also increase input sensitivity for worse generalization performance. Conceptually, this is comparable to overfitting in nearest-neighbor classification, where representations are sparse, one-hot indicator vectors corresponding to nearest training examples. As the number of training data increases, the distance between them decreases, so they are more likely to be confused with one another. Similarly, nearby dictionary atoms may introduce instability that causes representations of similar data points to become very far apart leading to poor generalization performance. Thus, there is a fundamental tradeoff between the capacity and robustness of shallow representations due to the proximity of dictionary atoms as measured by mutual coherence.

However, deep representations have not shown the same correlation between model size and sensitivity [34]. While adding more layers to a deep neural network increases its capacity, it also simultaneously introduces implicit regularization to reduce overfitting. This can be explained through the proposed connection to sparse coding, where additional layers increase both capacity and effective input dimensionality. In a higher-dimensional space, dictionary atoms can be spaced further apart for more robust representations. Furthermore, architectures with denser skip connections induce dictionary structures with more nonzero elements, which provides additional freedom to further reduce mutual coherence with fewer parameters as shown in Fig. 2.

We propose to use the minimum deep frame potential as a cue for model selection. Instead of requiring expensive validation on a specific dataset to approximate generalization performance, architectures are chosen based on how efficiently they can reduce the minimum achievable mutual coherence with respect to the number of model parameters. In this paper, we provide an efficient frame potential minimization method for a general class of convolu-

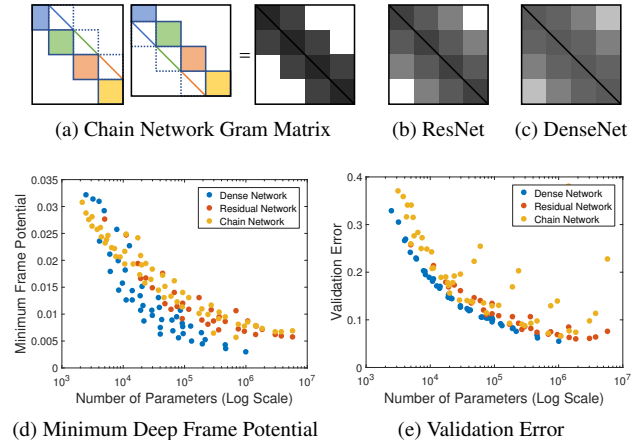


Figure 2: Parameter count is not a good indicator of generalization performance for deep networks. Instead, we compare different network architectures via the minimum deep frame potential, the average nonzero magnitude of inner products between atoms of architecture-induced dictionaries. In comparison to (a) chain networks, skip connections in (b) residual networks and (c) densely connected networks produce Gram matrix structures with more nonzero elements allowing for (d) lower deep frame potentials across network sizes. This correlates with improved parameter efficiency giving (e) lower validation error with fewer parameters.

tional networks with skip connections, of which ResNets and DenseNets are shown to be special cases. Furthermore, we derive an analytic expression for the minimum value in the case of fully-connected chain networks. Experimentally, we demonstrate correlation with validation error across a variety of network architectures.

2. Background and Related Work

Due to the vast space of possible deep network architectures and the computational difficulty in training them, deep model selection has largely been guided by ad-hoc engineering and human ingenuity. While progress slowed in the years following early breakthroughs [20], recent interest in deep learning architectures began anew due to empirical successes largely attributed to computational advances like efficient training using GPUs and rectified linear unit (ReLU) activation functions [18]. Since then, numerous architectural changes have been proposed. For example, much deeper networks with residual connections were shown to achieve consistently better performance with fewer parameters [13]. Building upon this, densely connected convolutional networks with skip connections between more layers yielded even better performance [15]. While theoretical explanations for these improvements were lacking, consistent experimentation on standardized benchmark datasets continued to drive empirical success.

However, due to slowing progress and the need for increased accessibility of deep learning techniques to a wider range of practitioners, more principled approaches to architecture search have recently gained traction. Motivated by observations of extreme redundancy in the parameters of trained networks [7], techniques have been proposed to systematically reduce the number of parameters without adversely affecting performance. Examples include sparsity-inducing regularizers during training [1] or through post-processing to prune the parameters of trained networks [14]. Constructive approaches to model selection like neural architecture search [12] instead attempt to compose architectures from basic building blocks through tools like reinforcement learning. Efficient model scaling has also been proposed to enable more effective grid search for selecting architectures subject to resource constraints [30]. While automated techniques can match or even surpass manually engineered alternatives, they require a validation dataset and rarely provide insights transferable to other settings.

To better understand the implicit benefits of different network architectures, there have been adjacent theoretical explorations of deep network generalization. These works are often motivated by the surprising observation that good performance can still be achieved using highly over-parametrized models with degrees of freedom that surpass the number of training data. This contradicts many commonly accepted ideas about generalization, spurning new experimental explorations that have demonstrated properties unique to deep learning. Examples include the ability of deep networks to express random data labels [34] with a tendency towards learning simple patterns first [2]. While exact theoretical explanations are lacking, empirical measurements of network sensitivity such as the Jacobian norm have been shown to correlate with generalization [25]. Similarly, Parseval regularization [22] encourages robustness by constraining the Lipschitz constants of individual layers.

Due to the difficulty in analyzing deep networks directly, other approaches have instead drawn connections to the rich field of sparse approximation theory. The relationship between feed-forward neural networks and principal component analysis has long been known for the case of linear activations [3]. More recently, nonlinear deep networks with ReLU activations have been linked to multilayer sparse coding to prove theoretical properties of deep representations [26]. This connection has been used to motivate new recurrent architecture designs that resist adversarial noise attacks [27], improve classification performance [29], or enforce prior knowledge through output constraints [23].

3. Deep Learning as Sparse Approximation

To derive our criterion for model selection, we build upon recent connections between deep neural networks and sparse approximation. Specifically, consider a feed-forward

network $f(\mathbf{x}) = \phi_l(\mathbf{B}_l^\top \cdots \phi_1(\mathbf{B}_1^\top(\mathbf{x})))$ constructed as the composition of linear transformations with parameters \mathbf{B}_j and nonlinear activation functions ϕ_j . Equivalently, $f(\mathbf{x}) = \mathbf{w}_l$ where $\mathbf{w}_j = \mathbf{B}_j^\top \mathbf{w}_{j-1}$ are the layer activations for layers $j = 1, \dots, l$ and $\mathbf{w}_0 = \mathbf{x}$. In many modern state-of-the-art networks, the ReLU activation function has been adopted due to its effectiveness and computational efficiency. It can also be interpreted as the nonnegative soft-thresholding proximal operator associated with the function Φ in Eq. 1, a nonnegativity constraint $\mathbb{I}(\mathbf{w} \geq \mathbf{0})$ and a sparsity-inducing ℓ_1 penalty with a weight determined by the scalar bias parameter λ .

$$\begin{aligned} \Phi(\mathbf{w}) &= \mathbb{I}(\mathbf{w} \geq \mathbf{0}) + \lambda \|\mathbf{w}\|_1 \quad (1) \\ \phi(\mathbf{x}) &= \text{ReLU}(\mathbf{x} - \lambda \mathbf{1}) = \arg \min_{\mathbf{w}} \frac{1}{2} \|\mathbf{w} - \mathbf{x}\|_2^2 + \Phi(\mathbf{w}) \end{aligned}$$

Thus, the forward pass of a deep network is equivalent to a layered thresholding pursuit algorithm for approximating the solution of a multi-layer sparse coding model [26]. Results from shallow sparse approximation theory can then be adapted to bound the accuracy of this approximation, which improves as the mutual coherence defined below in Eq. 2 decreases, and indirectly analyze other theoretical properties of deep networks like uniqueness and robustness.

3.1. Sparse Approximation Theory

Sparse approximation theory considers representations of data vectors $\mathbf{x} \in \mathbb{R}^d$ as sparse linear combinations $\mathbf{x} \approx \sum_j w_j \mathbf{b}_j = \mathbf{B}\mathbf{w}$ of atoms from an over-complete dictionary $\mathbf{B} \in \mathbb{R}^{d \times k}$. The number of atoms k is greater than the dimensionality d and the number of nonzero coefficients $\|\mathbf{w}\|_0$ in the representation $\mathbf{w} \in \mathbb{R}^k$ is small.

Through applications like compressed sensing [8], sparsity has been found to exhibit theoretical properties that enable data representation with efficiency far greater than what was previously thought possible. Central to these results is the requirement that the dictionary be “well-behaved,” essentially ensuring that its columns are not too similar. For undercomplete matrices with $k \leq d$, this is satisfied by enforcing orthogonality, but overcomplete dictionaries require other conditions. Specifically, we focus our attention on the mutual coherence μ , the maximum magnitude normalized inner product of all pairs of dictionary atoms. Equivalently, it is the maximum magnitude off-diagonal element in the Gram matrix $\mathbf{G} = \mathbf{B}^\top \mathbf{B}$ where the columns of \mathbf{B} are normalized to have unit norm:

$$\mu = \max_{i \neq j} \frac{|\mathbf{b}_i^\top \mathbf{b}_j|}{\|\mathbf{b}_i\| \|\mathbf{b}_j\|} = \max_{i,j} |(\mathbf{G} - \mathbf{I})_{ij}| \quad (2)$$

We are primarily motivated by the observation that a model’s capacity for low mutual coherence increases along with its capacity for both memorizing more training data through unique representations and generalizing to more validation data through robustness to input perturbations.

With an overcomplete dictionary, there is a space of coefficients \mathbf{w} that can all exactly reconstruct any data point as $\mathbf{x} = \mathbf{B}\mathbf{w}$, which would not support discriminative representation learning. However, if representations from a mutually incoherent dictionary are sufficiently sparse, then they are guaranteed to be optimally sparse and unique [9]. Specifically, if the number of nonzeros $\|\mathbf{w}\|_0 < \frac{1}{2}(1 + \mu^{-1})$, then \mathbf{w} is the unique, sparsest representation for \mathbf{x} . Furthermore, if $\|\mathbf{w}\|_0 < (\sqrt{2} - 0.5)\mu^{-1}$, then it can be found efficiently by convex optimization with ℓ_1 regularization. Thus, minimizing the mutual coherence of a dictionary increases its capacity for uniquely representing data points.

Sparse representations are also robust to input perturbations [10]. Specifically, given a noisy datapoint $\mathbf{x} = \mathbf{x}_0 + \mathbf{z}$ where \mathbf{x}_0 can be represented exactly as $\mathbf{x}_0 = \mathbf{B}\mathbf{w}_0$ with $\|\mathbf{w}_0\|_0 \leq \frac{1}{4}(1 + \mu^{-1})$ and the noise \mathbf{z} has bounded magnitude $\|\mathbf{z}\|_2 \leq \epsilon$, then \mathbf{w}_0 can be approximated by solving the ℓ_1 -penalized LASSO problem:

$$\arg \min_{\mathbf{w}} \|\mathbf{x} - \mathbf{B}\mathbf{w}\|_2^2 + \lambda \|\mathbf{w}\|_1 \quad (3)$$

Its solution is stable and the approximation error is bounded from above in Eq. 4, where $\delta(\mathbf{x}, \lambda)$ is a constant.

$$\|\mathbf{w} - \mathbf{w}_0\|_2^2 \leq \frac{(\epsilon + \delta(\mathbf{x}, \lambda))^2}{1 - \mu(4\|\mathbf{w}\|_0 - 1)} \quad (4)$$

Thus, minimizing the mutual coherence of a dictionary decreases the sensitivity of its sparse representations for improved robustness. This is similar to evaluating input sensitivity using the Jacobian norm [25]. However, instead of estimating the average perturbation error over validation data, it bounds the worst-case error over all possible data.

3.2. Deep Component Analysis

While deep representations can be analyzed by accumulating the effects of approximating individual layers in a chain network as shallow sparse coding problems [26], this strategy cannot be easily adapted to account for more complicated interactions between layers. Instead, we adapt the framework of Deep Component Analysis [23], which jointly represents all layers in a neural network as the single sparse coding problem in Eq. 5. The ReLU activations $\mathbf{w}_j \in \mathbb{R}^{k_j}$ of a feed-forward chain network approximate the solutions to a joint optimization problem where $\mathbf{w}_0 = \mathbf{x} \in \mathbb{R}^{k_0}$ and the regularization functions Φ_j are nonnegative sparsity-inducing penalties as defined in Eq. 1.

$$\mathbf{w}_j := \phi_j(\mathbf{B}_j^\top \mathbf{w}_{j-1}) \quad \forall j = 1, \dots, l \quad (5)$$

$$\approx \arg \min_{\{\mathbf{w}_j\}} \sum_{j=1}^l \|\mathbf{B}_j \mathbf{w}_j - \mathbf{w}_{j-1}\|_2^2 + \Phi_j(\mathbf{w}_j)$$

The compositional constraints between adjacent layers are relaxed and replaced by reconstruction error penalty terms, resulting in a convex, nonnegative sparse coding problem.

By combining the terms in the summation of Eq. 5 together into a single system, this problem can be equivalently represented as shown in Eq. 6. The latent variables \mathbf{w}_j for each layer are stacked in the vector \mathbf{w} , the regularizer $\Phi(\mathbf{w}) = \sum_j \Phi_j(\mathbf{w}_j)$, and the input \mathbf{x} is augmented with zeros.

$$\arg \min_{\mathbf{w}} \left\| \begin{array}{c} \mathbf{B} \\ \left[\begin{array}{cc} \mathbf{B}_1 & \mathbf{0} \\ -\mathbf{I} & \mathbf{B}_2 \quad \ddots \\ \ddots & \ddots & \mathbf{0} \\ -\mathbf{I} & \mathbf{B}_l \end{array} \right] \end{array} \right\|_2 \begin{array}{c} \mathbf{w} \\ \left[\begin{array}{c} \mathbf{w}_1 \\ \mathbf{w}_2 \\ \vdots \\ \mathbf{w}_l \end{array} \right] \end{array} - \begin{array}{c} \mathbf{x} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{array} \right\|_2^2 + \Phi(\mathbf{w}) \quad (6)$$

The layer parameters $\mathbf{B}_j \in \mathbb{R}^{k_{j-1} \times k_j}$ are blocks in the induced dictionary \mathbf{B} , which has $\sum_j k_{j-1}$ rows and $\sum_j k_j$ columns. It has a structure of nonzero elements that summarizes the corresponding feed-forward deep network architecture wherein the off-diagonal identity matrices are connections between adjacent layers.

Model capacity can be increased both by adding parameters to a layer or by adding layers, which implicitly pads the input data \mathbf{x} with more zeros. This can actually reduce the dictionary's mutual coherence because it increases the system's dimensionality. Thus, depth allows model complexity to scale jointly alongside effective input dimensionality so that the induced dictionary structures still have the capacity for low mutual coherence and improved capabilities for memorization and generalization.

3.3. Architecture-Induced Dictionary Structure

In this section, we extend this model formulation to incorporate more complicated network architectures. Because mutual coherence is dependent on normalized dictionary atoms, it can be reduced by increasing the number of nonzero elements, which reduces the magnitudes of the dictionary elements and their inner products. In Eq. 7, we replace the identity connections of Eq. 6 with blocks of nonzero parameters to allow for lower mutual coherence.

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_{11} & \mathbf{0} & & & \\ \mathbf{B}_{21}^\top & \mathbf{B}_{22} & & \ddots & \\ \vdots & \ddots & \ddots & \ddots & \mathbf{0} \\ \mathbf{B}_{l1}^\top & \cdots & \mathbf{B}_{l(l-1)}^\top & \mathbf{B}_{ll} & \end{bmatrix} \quad (7)$$

This structure is induced by the feed-forward activations in Eq. 8, which again approximate the solutions to a nonnegative sparse coding problem.

$$\mathbf{w}_j := \phi_j \left(-\mathbf{B}_{jj}^\top \sum_{k=1}^{j-1} \mathbf{B}_{jk}^\top \mathbf{w}_k \right) \quad \forall j = 1, \dots, l \quad (8)$$

$$\approx \arg \min_{\{\mathbf{w}_j\}} \sum_{j=1}^l \left\| \mathbf{B}_{jj} \mathbf{w}_j + \sum_{k=1}^{j-1} \mathbf{B}_{jk}^\top \mathbf{w}_k \right\|_F^2 + \Phi_j(\mathbf{w}_j)$$

In comparison to Eq. 5, additional parameters introduce skip connections between layers so that the activations \mathbf{w}_j of layer j now depend on those of all previous layers $k < j$.

These connections are similar to the identity mappings in residual networks [13], which introduce dependence between the activations of pairs of layers for even $j \in [1, l-1]$:

$$\mathbf{w}_j := \phi_j(\mathbf{B}_j^\top \mathbf{w}_{j-1}), \mathbf{w}_{j+1} := \phi_{j+1}(\mathbf{w}_{j-1} + \mathbf{B}_{j+1}^\top \mathbf{w}_j) \quad (9)$$

In comparison to chain networks, no additional parameters are required; the only difference is the addition of \mathbf{w}_{j-1} in the argument of ϕ_{j+1} . As a special case of Eq. 8, we interpret the activations in Eq. 9 as approximate solutions to the optimization problem:

$$\arg \min_{\{\mathbf{w}_j\}} \|\mathbf{x} - \mathbf{B}_1 \mathbf{w}_1\|_2^2 + \sum_{j=1}^l \Phi_j(\mathbf{w}_j) \quad (10)$$

$$+ \sum_{\text{even } j} \|\mathbf{w}_j - \mathbf{B}_j^\top \mathbf{w}_{j-1}\|_2^2 + \|\mathbf{w}_{j+1} - \mathbf{w}_{j-1} - \mathbf{B}_{j+1}^\top \mathbf{w}_j\|_2^2$$

This results in the induced dictionary structure of Eq. 7 with $\mathbf{B}_{jj} = \mathbf{I}$ for $j > 1$, $\mathbf{B}_{jk} = \mathbf{0}$ for $j > k + 1$, $\mathbf{B}_{jk} = \mathbf{0}$ for $j > k$ with odd k , and $\mathbf{B}_{jk} = \mathbf{I}$ for $j > k$ with even k .

Building upon the empirical successes of residual networks, densely connected convolution networks [15] incorporate skip connections between earlier layers as well. This is shown in Eq. 11 where the transformation \mathbf{B}_j of concatenated variables $[\mathbf{w}_k]_k$ for $k = 1, \dots, j-1$ is equivalently written as the summation of smaller transformations \mathbf{B}_{jk} .

$$\mathbf{w}_j := \phi_j\left(\mathbf{B}_j^\top [\mathbf{w}_k]_{k=1}^{j-1}\right) = \phi_j\left(\sum_{k=1}^{j-1} \mathbf{B}_{jk}^\top \mathbf{w}_k\right) \quad (11)$$

These activations again provide approximate solutions to the problem in Eq. 8 with the induced dictionary structure of Eq. 7 where $\mathbf{B}_{jj} = \mathbf{I}$ for $j > 1$ and the lower blocks \mathbf{B}_{jk} for $j > k$ are all filled with learned parameters.

Skip connections enable effective learning in much deeper networks than chain-structured alternatives. While originally motivated from the perspective of making optimization easier [13], adding more connections between layers was also shown to improve generalization performance [15]. As compared in Fig. 2, denser skip connections induce dictionary structures with denser Gram matrices allowing for lower mutual coherence. This suggests that architectures' capacities for low validation error can be quantified and compared based on their capacities for inducing dictionaries with low minimum mutual coherence.

4. The Deep Frame Potential

We propose to use lower bounds on the mutual coherence of induced structured dictionaries for the data-independent comparison of architecture capacities. Note that while one-sided coherence is better suited to nonnegativity constraints,

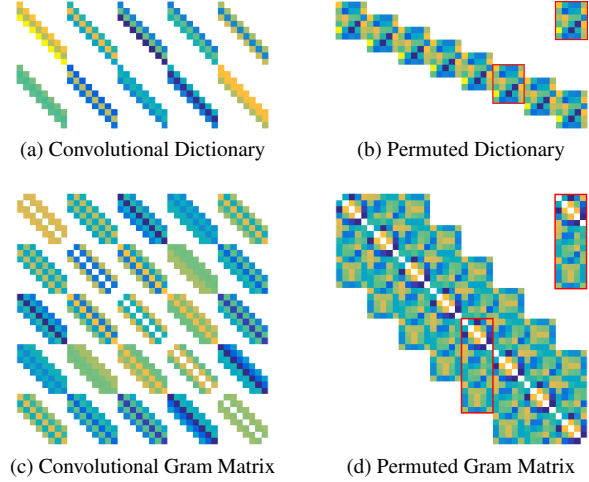


Figure 3: A visualization of a one-dimensional convolutional dictionary with two input channels, five output channels, and a filter size of three. (a) The filters are repeated over eight spatial dimensions resulting in a (b) block-Toeplitz structure that is revealed through row and column permutations. (c) The corresponding gram matrix can be efficiently computed by (d) repeating local filter interactions.

it has the same lower bound [5]. Directly optimizing mutual coherence from Eq. 2 is difficult due to its piecewise structure. Instead, we consider a tight lower by replacing the maximum off-diagonal element of the Gram matrix \mathbf{G} with the mean. This gives the averaged frame potential $F^2(\mathbf{B})$, a strongly-convex function that can be optimized more effectively [4]:

$$F^2(\mathbf{B}) = N^{-1}(\mathbf{G}) \left(\|\mathbf{G}\|_F^2 - \text{Tr } \mathbf{G} \right) \leq \mu^2(\mathbf{B}) \quad (12)$$

Here, $N(\mathbf{G})$ is the number of nonzero off-diagonal elements in the Gram matrix and $\text{Tr } \mathbf{G}$ equals the total number of dictionary atoms. Equality is met in the case of equiangular tight frames when the normalized inner products between all dictionary atoms are equivalent [17]. Due to the block-sparse structure of the induced dictionaries from Eq. 7, we evaluate the frame potential in terms of local blocks $\mathbf{G}_{jj'} \in \mathbb{R}^{k_j \times k_{j'}}$ that are nonzero only if layer j is connected to layer j' . In the case of convolutional layers with localized spatial support, there is also a repeated implicit structure of nonzero elements as visualized in Fig. 3.

To compute the Gram matrix, we first need to normalize the global induced dictionary \mathbf{B} from Eq. 7. By stacking the column magnitudes of layer j as the elements in the diagonal matrix $\mathbf{C}_j = \text{diag}(\mathbf{c}_j) \in \mathbb{R}^{k_j \times k_j}$, the normalized parameters can be represented as $\tilde{\mathbf{B}}_{ij} = \mathbf{B}_{ij} \mathbf{C}_j^{-1}$. Similarly, the squared norms of the full set of columns in the global dictionary \mathbf{B} are $\mathbf{N}_j^2 = \sum_{i=j}^l \mathbf{C}_i^2$. The full normalized dictionary can then be found as $\tilde{\mathbf{B}} = \mathbf{B} \mathbf{N}^{-1}$ where

the matrix \mathbf{N} is block diagonal with \mathbf{N}_j as its blocks. The blocks of the Gram matrix $\mathbf{G} = \tilde{\mathbf{B}}^\top \tilde{\mathbf{B}}$ are then given as:

$$\mathbf{G}_{jj'} = \sum_{i=j'}^l \mathbf{N}_j^{-1} \mathbf{B}_{ij}^\top \mathbf{B}_{ij'} \mathbf{N}_{j'}^{-1} \quad (13)$$

For chain networks, $\mathbf{G}_{jj'} \neq \mathbf{0}$ only when $j' = j + 1$, which represents the connections between adjacent layers. In this case, the blocks can be simplified as:

$$\mathbf{G}_{jj} = (\mathbf{C}_j^2 + \mathbf{I})^{-\frac{1}{2}} (\mathbf{B}_j^\top \mathbf{B}_j + \mathbf{I}) (\mathbf{C}_j^2 + \mathbf{I})^{-\frac{1}{2}} \quad (14)$$

$$\mathbf{G}_{j(j+1)} = -(\mathbf{C}_j^2 + \mathbf{I})^{-\frac{1}{2}} \mathbf{B}_{j+1} (\mathbf{C}_{j+1}^2 + \mathbf{I})^{-\frac{1}{2}} \quad (15)$$

$$\mathbf{G}_{ll} = \mathbf{B}_l^\top \mathbf{B}_l \quad (16)$$

Because the diagonal is removed in the deep frame potential computation, the contribution of \mathbf{G}_{jj} is simply a rescaled version of the local frame potential of layer j . The contribution of $\mathbf{G}_{j(j+1)}$, on the other hand, can essentially be interpreted as rescaled ℓ_2 weight decay where rows are weighted more heavily if the corresponding columns of the previous layer's parameters have higher magnitudes. Furthermore, since the global frame potential is averaged over the total number of nonzero elements in \mathbf{G} , if a layer has more parameters, then it will be given more weight in this computation. For more general networks with skip connections, however, the summation from Eq. 13 has additional terms that introduce more complicated interactions. In these cases, it cannot be evaluated from local properties of layers.

Essentially, the deep frame potential summarizes the structural properties of the global dictionary \mathbf{B} induced by a deep network architecture by balancing interactions within each individual layer through local coherence properties and between connecting layers.

4.1. Theoretical Lower Bound for Chain Networks

While the deep frame potential is a function of parameter values, its minimum value is determined only by the architecture-induced dictionary structure. Furthermore, we know that it must be bounded by a nonzero constant for overcomplete dictionaries. In this section, we derive this lower bound for the special case of fully-connected chain networks and provide intuition for why skip connections increase the capacity for low mutual coherence.

First, observe that a lower bound for the Frobenius norm of $\mathbf{G}_{j(j+1)}$ from Eq. 15 cannot be readily attained because the rows and columns are rescaled independently. This means that a lower bound for the norm of \mathbf{G} must be found by jointly considering the entire matrix structure, not simply through the summation of its components. To accomplish this, we instead consider the matrix $\mathbf{H} = \tilde{\mathbf{B}} \tilde{\mathbf{B}}^\top$, which is full rank and has the same norm as \mathbf{G} :

$$\|\mathbf{G}\|_F^2 = \|\mathbf{H}\|_F^2 = \sum_{j=1}^l \|\mathbf{H}_{jj}\|_F^2 + 2 \sum_{j=1}^{l-1} \|\mathbf{H}_{j(j+1)}\|_F^2 \quad (17)$$

We can then express the individual blocks of \mathbf{H} as:

$$\mathbf{H}_{11} = \mathbf{B}_1 (\mathbf{C}_1^2 + \mathbf{I}_{k_1})^{-1} \mathbf{B}_1^\top \quad (18)$$

$$\mathbf{H}_{jj} = \mathbf{B}_j (\mathbf{C}_j^2 + \mathbf{I})^{-1} \mathbf{B}_j^\top + (\mathbf{C}_{j-1}^2 + \mathbf{I})^{-1} \quad (19)$$

$$\mathbf{H}_{j(j+1)} = -\mathbf{B}_j (\mathbf{C}_j^2 + \mathbf{I})^{-1} \quad (20)$$

In contrast to $\mathbf{G}_{j(j+1)}$ in Eq. 15, only the columns of $\mathbf{H}_{j(j+1)}$ in Eq. 20 are rescaled. Since $\tilde{\mathbf{B}}_j$ has normalized columns, its norm can be exactly expressed as:

$$\|\mathbf{H}_{j(j+1)}\|_F^2 = \sum_{n=1}^{k_j} \left(\frac{c_{jn}}{c_{jn}^2 + 1} \right)^2 \quad (21)$$

For the other blocks, we find lower bounds for their norms through the same technique used in deriving the Welch bound, which expresses the minimum mutual coherence for unstructured overcomplete dictionaries [32]. Specifically, we apply the Cauchy-Schwarz inequality giving $\|\mathbf{A}\|_F^2 \geq r^{-1} (\text{Tr } \mathbf{A})^2$ for positive-semidefinite matrices \mathbf{A} with rank r . Since the rank of \mathbf{H}_{jj} is at most k_{j-1} , we can lower bound the norms of the individual blocks as:

$$\|\mathbf{H}_{11}\|_F^2 \geq \frac{1}{k_0} \left(\sum_{n=1}^{k_1} \frac{c_{1n}^2}{c_{1n}^2 + 1} \right)^2 \quad (22)$$

$$\|\mathbf{H}_{jj}\|_F^2 \geq \frac{1}{k_{j-1}} \left(\sum_{n=1}^{k_j} \frac{c_{jn}^2}{c_{jn}^2 + 1} + \sum_{p=1}^{k_{j-1}} \frac{1}{c_{(j-1)p}^2 + 1} \right)^2$$

In this case of dense shallow dictionaries, the Welch bound depends only on the data dimensionality and the number of dictionary atoms. However, due to the structure of the architecture-induced dictionaries, the lower bound of the deep frame potential depends on the data dimensionality, the number of layers, the number of units in each layer, the connectivity between layers, and the relative magnitudes between layers. Skip connections increase the number of nonzero elements in the Gram matrix over which to average and also enable off-diagonal blocks to have lower norms.

4.2. Model Selection

For more general architectures that lack a simple theoretical lower bound, we instead propose bounding the mutual coherence of the architecture-induced dictionary through empirical minimization of the deep frame potential $F^2(\mathbf{B})$ from Eq. 12. Frame potential minimization has been used effectively to construct finite normalized tight frames due to the lack of suboptimal local minima, which allows for effective optimization using gradient descent [4]. We propose using the minimum deep frame potential of an architecture—which is independent of data and individual parameter instantiations—as a means to compare different architectures. In practice, model selection is performed by choosing the candidate architecture with the lowest minimum frame potential subject to desired modeling constraints such as limiting the total number of parameters.

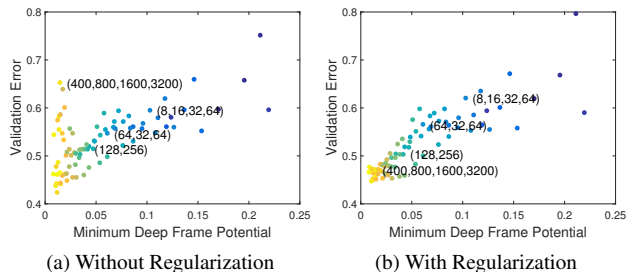


Figure 4: A comparison of fully connected deep network architectures with varying depths and widths. Warmer colors indicate models with more total parameters. (a) Some very large networks cannot be trained effectively resulting in unusually high validation errors. (b) This can be remedied through deep frame potential regularization, resulting in high correlation between minimum frame potential and validation error.

5. Experimental Results

In this section, we demonstrate correlation between the minimum deep frame potential and validation error on the CIFAR-10 dataset [19] across a wide variety of fully-connected, convolutional, chain, residual, and densely connected network architectures. Furthermore, we show that networks with skip connections can have lower deep frame potentials with fewer learnable parameters, which is predictive of the parameter efficiency of trained networks.

In Fig. 4, we visualize a scatter plot of trained fully-connected networks with between three and five layers and between 16 and 4096 units in each layer. The corresponding architectures are shown as a list of units per layer for a few representative examples. The minimum frame potential of each architecture is compared against its validation error after training, and the total parameter count is indicated by color. In Fig. 4a, some networks with many parameters have unusually high error due to the difficulty in training very large fully-connected networks. In Fig. 4b, the addition of a deep frame potential regularization term overcomes some of these optimization difficulties for improved parameter efficiency. This results in high correlation between minimum frame potential and validation error. Furthermore, it emphasizes the diminishing returns of increasing the size of fully-connected chain networks; after a certain point, adding more parameters does little to reduce both validation error and minimum frame potential.

To evaluate the effects of residual connections [13], we adapt the simplified CIFAR-10 ResNet architecture from [33], which consists of a single convolutional layer followed by three groups of residual blocks with activations given in Eq. 9. Before the second and third groups, the number of filters is increased by a factor of two and the spatial resolution is decreased by half through average pooling.

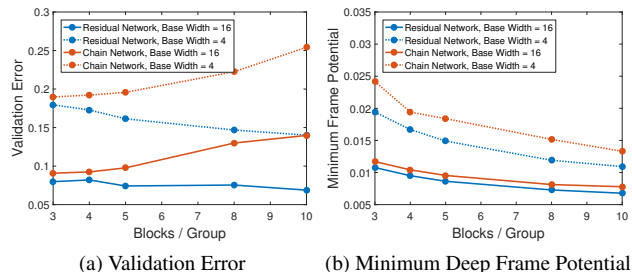


Figure 5: The effect of increasing depth in chain and residual networks. Validation error is compared against layer count for two different network widths. (a) In comparison to chain networks, even very deep residual networks can be trained effectively resulting in decreasing validation error. (b) Despite having the same number of total parameters, residual connections also induce dictionary structures with lower minimum deep frame potentials.

To compare networks with different sizes, we modify their depths by changing the number of residual blocks in each group from between 2 and 10 and their widths by changing the base number of filters from between 4 and 32. For our experiments with densely connected skip connections [15], we adapt the simplified CIFAR-10 DenseNet architecture from [21]. Like with residual networks, it consists of a convolutional layer followed by three groups of the activations from Eq. 11 with decreasing spatial resolutions and increasing numbers of filters. Within each group, a dense block is the concatenation of smaller convolutions that take all previous outputs as inputs with filter numbers equal to a fixed growth rate. Network depth and width are modified by respectively increasing both the number of layers per group and the base growth rate from between 2 and 12. Batch normalization [16] was also used in all convolutional networks.

In Fig. 5, we compare the validation errors and minimum frame potentials of residual networks and comparable chain networks with residual connections removed. In Fig. 5a, the validation error of chain networks increases for deeper networks while that of residual networks is lower and consistently decreases. This emphasizes the difficulty in training very deep chain networks. In Fig. 5b, we show that residual connections enable lower minimum frame potentials following a similar trend with respect to increasing model size, again demonstrating correlation between validation error and minimum frame potential.

In Fig. 6, we compare chain networks and residual networks with exactly the same number of parameters, where color indicates the number of residual blocks per group and connected data points have the same depths but different widths. The addition of skip connections reduces both validation error and minimum frame potential, as visualized by consistent placement below the diagonal line indicating lower values for residual networks than comparable chain

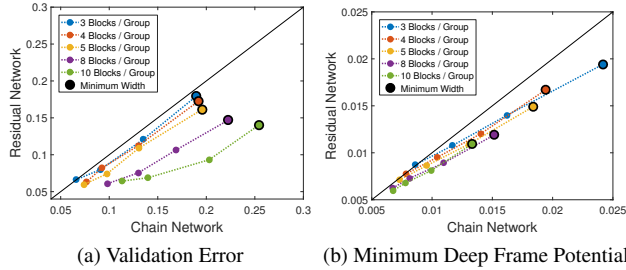


Figure 6: A comparison of (a) validation error and (b) minimum frame potential between residual networks and chain networks. Colors indicate different depths and datapoints are connected in order of increasing widths of 4, 8, 16, or 32 filters. Skip connections result in reduced error correlating with frame potential with dense networks showing superior efficiency with increasing depth.

networks. This effect becomes even more pronounced with increasing depths and widths.

In Fig. 7, we compare the parameter efficiency of chain networks, residual networks, and densely connected networks of different depths and widths. We visualize both validation error and minimum frame potential as functions of the number of parameters, demonstrating the improved scalability of networks with skip connections. While chain networks demonstrate increasingly poor parameter efficiency with respect to increasing depth in Fig. 7a, the skip connections of ResNets and DenseNets allow for further reducing error with larger network sizes in Figs. 7c,e. Considering all network families together as in Fig. 2d, we see that denser connections also allow for lower validation error with comparable numbers of parameters. This trend is mirrored in the minimum frame potentials of Figs. 7b,d,f which are shown together in Fig. 2e. Despite some fine variations in behavior across different families of architectures, minimum deep frame potential is correlated with validation error across network sizes and effectively predicts the increased generalization capacity provided by skip connections.

6. Conclusion

In this paper, we proposed a technique for comparing deep network architectures by approximately quantifying their implicit capacity for effective data representations. Based upon theoretical connections between sparse approximation and deep neural networks, we demonstrated how architectural hyper-parameters such as depth, width, and skip connections induce different structural properties of the dictionaries in corresponding sparse coding problems. We compared these dictionary structures through lower bounds on their mutual coherence, which is theoretically tied to their capacity for uniquely and robustly representing data via sparse approximation. A theoretical lower bound was derived for chain networks and the deep frame potential was

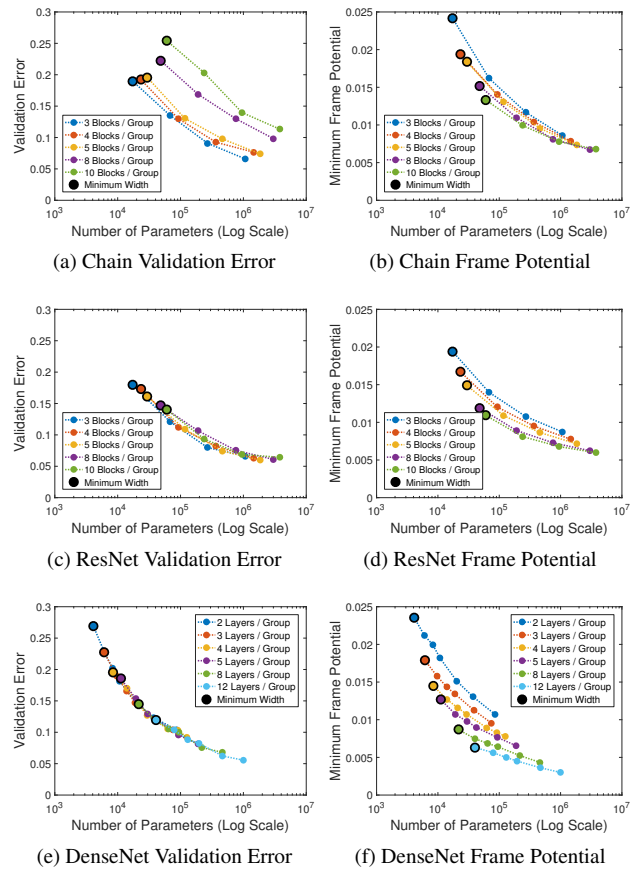


Figure 7: A demonstration of the improved scalability of networks with skip connections, where line colors indicate different depths and data points are connected showing increasing widths. (a) Chain networks with greater depths have increasingly worse parameter efficiency in comparison to (c) the corresponding networks with residual connections and (e) densely connected networks with similar size, of which performance scales efficiently with parameter count. This could potentially be attributed to correlated efficiency in reducing frame potential with fewer parameters, which saturates much faster with (b) chain networks than (d) residual networks or (f) densely connected networks.

proposed as an empirical optimization objective for constructing bounds for more complicated architectures.

Experimentally, we observed a correlation between minimum deep frame potential and validation error across different families of modern architectures with skip connections, including residual networks and densely connected convolutional networks. This suggests a promising direction for future research towards the theoretical analysis and practical construction of deep network architectures derived from connections between deep learning and sparse coding.

Acknowledgments: This work was supported by the CMU Argo AI Center for Autonomous Vehicle Research and by the National Science Foundation under Grant No.1925281.

References

- [1] Jose Alvarez and Mathieu Salzmann. Learning the number of neurons in deep networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016. 3
- [2] Devansh Arpit et al. A closer look at memorization in deep networks. In *International Conference on Machine Learning (ICML)*, 2017. 3
- [3] Pierre Baldi and Kurt Hornik. Neural networks and principal component analysis: Learning from examples without local minima. *Neural networks*, 2(1), 1989. 3
- [4] John Benedetto and Matthew Fickus. Finite normalized tight frames. *Advances in Computational Mathematics*, 18(2-4), 2003. 5, 6
- [5] Alfred M. Bruckstein et al. On the uniqueness of nonnegative sparse solutions to underdetermined systems of equations. *IEEE Transactions on Information Theory*, 54(11), 2008. 5
- [6] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3), 1995. 1
- [7] Misha Denil et al. Predicting parameters in deep learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2013. 3
- [8] David L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4), 2006. 3
- [9] David L. Donoho and Michael Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via l_1 minimization. *Proceedings of the National Academy of Sciences*, 100(5), 2003. 2, 4
- [10] David L. Donoho, Michael Elad, and Vladimir N. Temlyakov. Stable recovery of sparse overcomplete representations in the presence of noise. *IEEE Transactions on Information Theory*, 52(1), 2005. 4
- [11] Michael Elad. *Sparse and redundant representations: from theory to applications in signal and image processing*. Springer Science & Business Media, 2010. 2
- [12] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *Journal of Machine Learning Research (JMLR)*, 20(55), 2019. 3
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 2, 5, 7
- [14] Yihui He et al. Channel pruning for accelerating very deep neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [15] Gao Huang, Zhuang Liu, Kilian Weinberger, and Laurens van der Maaten. Densely connected convolutional networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 5, 7
- [16] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, 2015. 7
- [17] Jelena Kovačević, Amina Chebira, et al. An introduction to frames. *Foundations and Trends in Signal Processing*, 2(1), 2008. 5
- [18] Alex Krizhevsky et al. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2012. 1, 2
- [19] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 7
- [20] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 1998. 2
- [21] Yixuan Li. Tensorflow densenet. <https://github.com/YixuanLi/densenet-tensorflow>, 2018. 7
- [22] Cisse Moustapha, Bojanowski Piotr, Grave Edouard, Dauphin Yann, and Usunier Nicolas. Parseval networks: Improving robustness to adversarial examples. In *International Conference on Machine Learning (ICML)*, 2017. 3
- [23] Calvin Murdock, MingFang Chang, and Simon Lucey. Deep component analysis via alternating direction neural networks. In *European Conference on Computer Vision (ECCV)*, 2018. 1, 3, 4
- [24] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 1
- [25] Roman Novak et al. Sensitivity and generalization in neural networks: an empirical study. In *International Conference on Learning Representations (ICLR)*, 2018. 3, 4
- [26] Vardan Papyan, Yaniv Romano, and Michael Elad. Convolutional neural networks analyzed via convolutional sparse coding. *Journal of Machine Learning Research (JMLR)*, 18(83), 2017. 1, 3, 4
- [27] Yaniv Romano, Aviad Aberdam, Jeremias Sulam, and Michael Elad. Adversarial noise attacks of deep learning architectures-stability analysis via sparse modeled signals. *Journal of Mathematical Imaging and Vision*, 2018. 3
- [28] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015. 1
- [29] Jeremias Sulam, Aviad Aberdam, Amir Beck, and Michael Elad. On multi-layer basis pursuit, efficient algorithms and convolutional neural networks. *Pattern Analysis and Machine Intelligence (PAMI)*, 2019. 3
- [30] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning (ICML)*, 2019. 3
- [31] Vladimir N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and Its Applications*, XVI(2), 1971. 1
- [32] Lloyd Welch. Lower bounds on the maximum cross correlation of signals. *IEEE Transactions on Information theory*, 20(3), 1974. 6
- [33] Yuxin Wu et al. Tensorpack. <https://github.com/tensorpack/>, 2016. 7
- [34] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations (ICLR)*, 2017. 2, 3