
Compressing Neural Networks using the Variational Information Bottleneck

Bin Dai¹ Chen Zhu² Baining Guo³ David Wipf³

Abstract

Neural networks can be compressed to reduce memory and computational requirements, or to increase accuracy by facilitating the use of a larger base architecture. In this paper we focus on pruning individual neurons, which can simultaneously trim model size, FLOPs, and run-time memory. To improve upon the performance of existing compression algorithms we utilize the information bottleneck principle instantiated via a tractable variational bound. Minimization of this information theoretic bound reduces the redundancy between adjacent layers by aggregating useful information into a subset of neurons that can be preserved. In contrast, the activations of disposable neurons are shut off via an attractive form of sparse regularization that emerges naturally from this framework, providing tangible advantages over traditional sparsity penalties without contributing additional tuning parameters to the energy landscape. We demonstrate state-of-the-art compression rates across an array of datasets and network architectures.

1. Introduction

Although extremely effective across diverse application domains, it is nonetheless well-established that many popular deep neural network architectures are over-parameterized even with respect to datasets where predictive performance is high (Denil et al., 2013). Therefore, accuracy need not suffer per se, but unnecessarily large computational and memory footprints are required for practical deployment. Hence there is tremendous potential to compress trained networks while preserving the original accuracy. Alternatively, it has also been shown that even if a given network architecture is not necessarily over-parameterized, a larger

network pruned down to a smaller one of equivalent size often produces higher accuracy (Lin et al., 2017). Therefore network compression can potentially boost both efficiency and accuracy in some sense.

As reviewed in Section 2, a huge number of algorithmic pipelines have been recently proposed to instantiate some form of neural network compression; however, there is as of yet no perfect solution leaving room for new developments. In this paper, we borrow the idea of the information bottleneck (Tishby et al., 2000; Tishby & Zaslavsky, 2015), which provides a convenient mechanism for penalizing an information theoretic measure of redundancy between adjacent network layers which can be subsequently harnessed for compression.

More specifically, if we interpret these layers as forming a Markov chain, then typical training involves finding weights such that the maximal information pertaining to label y propagates from an input x to the network output. However, when a so-called information bottleneck is applied, superfluous information related to x but irrelevant for predicting y can be squeezed from the model layer-by-layer. There are multiple ways such a bottleneck could be introduced, but it is most helpful to consider strategies readily amenable to practical implementation.

To this end, we penalize the inter-layer mutual information using a variational approximation that both: (i) circumvents certain intractable integrations via a friendly bound, and (ii) reduces redundancy by aggregating useless information into certain expendable neurons using a latent sparsity-promotion mechanism. These neurons can naturally be identified and pruned to both reduce model size, FLOPs, and the run-time memory footprint. In accomplishing this, our main contributions are three-fold:

1. Beginning from the information bottleneck principle and a tractable variational approximation, we design a well-motivated neural network compression energy function that requires only a single, unavoidable tuning parameter for managing the compression/accuracy trade-off. No additional hyper-parameters for describing priors or other special constraints are required.
2. We carefully analyze an emergent tendency to accumulate useful information in a sparse set of neurons, while pushing the activations of others towards zero.

¹Institute for Advanced Study, Tsinghua University, Beijing, China ²Department of Computer Science, University of Maryland, USA ³Microsoft Research, Beijing, China. Correspondence to: Bin Dai <daib13@mails.tsinghua.edu.cn>, Chen Zhu <chen-zhu@umd.edu>, David Wipf <davidwipf@microsoft.com>.

Moreover, we quantify how this favoritism towards sparsity and implicit network pruning holds certain advantages over more traditional sparsity penalties that have previously been applied to network compression.

3. Finally, we present empirical results across the most common compression benchmarks, demonstrating improvement over numerous state-of-the-art approaches.

2. Related Network Compression Work

To obtain neural networks with low computational cost and/or memory footprint, prior work has involved designing more efficient architectures (Howard et al., 2017; Dong et al., 2017; Iandola et al., 2016), quantizing network weights (Courbariaux et al., 2016; 2015; Han et al., 2015a; Mellempudi et al., 2017; Rastegari et al., 2016), using efficient tensor or matrix decompositions to compress layers (Jaderberg et al., 2014; Zhang et al., 2016; Yu et al., 2017), or pruning existing network structures.

With respect to pruning, one option is to pre-train a network and then subsequently remove connections with small absolute values (Han et al., 2015b; Guo et al., 2016; LeCun et al., 1990). Other approaches instead employ more sophisticated Bayesian estimators (Blundell et al., 2015; Graves, 2011; Nalisnick et al., 2015; Ullrich et al., 2017; Molchanov et al., 2017). However, these cases cannot significantly reduce computation times and memory without some special coding and processing because the dimensionality of the neurons/activations has not been changed.

To address the latter, it is necessary to target activations for pruning, which is our primary focus herein. For this purpose, multiple interesting Bayesian approaches have been proposed (Louizos et al., 2017a; Neklyudov et al., 2017) that rely on sparse priors (e.g., Jeffreys, horseshoe) on either groups of weights along dimensions useful for compression, or directly on the activations themselves. A variational free energy approximation/bound on the log-likelihood is then optimized, potentially using the warm-start procedure from (Sønderby et al., 2016) to gradually increase the influence of regularization effects (at the expense of altering the original bound). Although the final objectives minimized by these approaches significantly overlap with each other and our method, they are derived from a different perspective and involve several key, differentiating assumptions.

As an alternative deterministic strategy, (Liu et al., 2017; Pan et al., 2016; Wen et al., 2016) address similar pruning effects by applying convex group Lasso or ℓ_1 norm-based regularizers in various different ways. In contrast, (Louizos et al., 2017b) first adopts an ideal ℓ_0 norm penalty on rows or columns of weight matrices, and then deals with the resulting discontinuous, non-convex energy surface via a probabilistic smoothing approximation. Finally, (Lin et al., 2017) introduces a radically different approach based on

reinforcement learning while (Li et al., 2016) specifically prunes convolutional neural network filters that have small norms.

3. Model Development

We first introduce some brief notational details. Denote the input variables to a neural network with L layers as $\mathbf{x} \in \mathbb{R}^d$ and the associated label (or target output) as $\mathbf{y} \in \mathcal{Y}$. We represent the network hidden layer activations as $\{\mathbf{h}_i\}_{i=1}^L$, where $\mathbf{h}_i \in \mathbb{R}^{r_i}$ and r_i denotes the number of neurons.

Now if we view \mathbf{x} as a stochastic input, feedforward network layers are sometimes interpreted as a Markov chain of successive representations (Tishby & Zaslavsky, 2015), *i.e.*,

$$\mathbf{y} \rightarrow \mathbf{x} \rightarrow \mathbf{h}_1 \rightarrow \dots \rightarrow \mathbf{h}_L \rightarrow \hat{\mathbf{y}}. \quad (1)$$

Every hidden layer in the network defines the conditional probability $p(\mathbf{h}_i|\mathbf{h}_{i-1})$, where we use $\mathbf{x} = \mathbf{h}_0$ for convenience. For a deterministic network model, $p(\mathbf{h}_i|\mathbf{h}_{i-1})$ can be regarded as a Dirac-delta function; however, there are also many situations where the hidden layers are stochastic even when conditioned on their inputs. For example, when using dropout (Srivastava et al., 2014), each hidden neuron has some probability of being set to zero. Likewise, in Bayesian neural networks (Blundell et al., 2015) and variational autoencoder models (Rezende et al., 2014; Kingma & Welling, 2014) some or all layer activations are assigned a non-degenerate distribution. In most such cases, the neurons within each stochastic hidden layer are assumed to be independent when conditioned on the activations from the previous layer, *e.g.*, a diagonal Gaussian distribution.

The role of the hidden layers is to extract information from the previous ones, while the output layer attempts to approximate the true distribution $p(\mathbf{y}|\mathbf{h}_L)$ via some tractable alternative $q(\mathbf{y}|\mathbf{h}_L)$. However, even if sufficient information for accurately predicting \mathbf{y} percolates through the network to the output, many of the internal representations may still contain superfluous content. Removal of this redundant content through some form of pruning or network ablation therefore represents a viable route to model compression.

Our starting point for accomplishing this goal is to explicitly penalize an information theoretic measure of redundancy between each adjacent layer, a concept originally introduced as the information bottleneck (Tishby et al., 2000). More specifically, for every hidden layer \mathbf{h}_i , we would like to *minimize* the mutual information $\mathbf{I}(\mathbf{h}_i; \mathbf{h}_{i-1})$ between \mathbf{h}_i and \mathbf{h}_{i-1} to remove inter-layer redundancy, while simultaneously *maximizing* the mutual information $\mathbf{I}(\mathbf{h}_i; \mathbf{y})$ between \mathbf{h}_i and the output \mathbf{y} to encourage accurate predictions of \mathbf{y} . Consequently, the layer-wise energy \mathcal{L}_i becomes

$$\mathcal{L}_i = \gamma_i \mathbf{I}(\mathbf{h}_i; \mathbf{h}_{i-1}) - \mathbf{I}(\mathbf{h}_i; \mathbf{y}), \quad (2)$$

where $\gamma_i \geq 0$ is a coefficient that determines the strength of the bottleneck, or the degree to which we value com-

pression over prediction accuracy. Summing over layers, the goal then is to minimize $\sum_i \mathcal{L}_i$ with respect to both network weights and any additional parameters describing the distributions $q(\mathbf{y}|\mathbf{h}_L)$ and $p(\mathbf{h}_i|\mathbf{h}_{i-1})$ for all i . However, reasonable model choices reflecting popular network architectures do not facilitate tractable computation of (2). Fortunately though, certain variational bounds can serve as a convenient surrogate. In this work, we invoke the upper bound

$$\begin{aligned} \tilde{\mathcal{L}}_i &= \gamma_i \mathbb{E}_{\{\mathbf{x}, \mathbf{y}\} \sim \mathcal{D}, \mathbf{h}_{1:i-1} \sim p(\mathbf{h}_{1:i-1}|\mathbf{x})} [\mathbb{KL}[p(\mathbf{h}_i|\mathbf{h}_{i-1})||q(\mathbf{h}_i)]] \\ &\quad - \mathbb{E}_{\{\mathbf{x}, \mathbf{y}\} \sim \mathcal{D}, \mathbf{h} \sim p(\mathbf{h}|\mathbf{x})} [\log q(\mathbf{y}|\mathbf{h}_L)] \geq \mathcal{L}_i, \end{aligned} \quad (3)$$

where $\mathbf{h}_{1:i} \triangleq \{\mathbf{h}_j\}_{j=1}^i$, $\mathbf{h} \triangleq \mathbf{h}_{1:L}$, \mathcal{D} denotes the true data distribution, and $q(\mathbf{h}_i)$ and $q(\mathbf{y}|\mathbf{h}_L)$ represent two variational distributions designed to approximate $p(\mathbf{h}_i)$ and $p(\mathbf{y}|\mathbf{h}_L)$ respectively. Details of the derivations can be found in (Dai et al., 2018) for space considerations.

$\tilde{\mathcal{L}}_i$ from (3) is composed of two terms. The first is the KL divergence between $p(\mathbf{h}_i|\mathbf{h}_{i-1})$ and $q(\mathbf{h}_i)$, which approximates how much information \mathbf{h}_i extracts from \mathbf{h}_{i-1} . The second term reflects fidelity with respect to the data distribution. The final variational information bottleneck loss function then becomes $\tilde{\mathcal{L}} = \sum_i \tilde{\mathcal{L}}_i$ to assimilate information management across all layers.

Of course to actually optimize $\tilde{\mathcal{L}}$ we need to specify a parametric form for the distributions $p(\mathbf{h}_i|\mathbf{h}_{i-1})$, $q(\mathbf{h}_i)$, and $q(\mathbf{y}|\mathbf{h}_L)$. For the latter, the final network layer with weights \mathbf{W}_y provides the necessary structure, often a multinomial distribution for classification problems and a Gaussian for standard regression tasks. And with respect to the conditional layer-wise distributions, we assume that each $p(\mathbf{h}_i|\mathbf{h}_{i-1})$ is defined via the relation

$$\mathbf{h}_i = (\boldsymbol{\mu}_i + \boldsymbol{\epsilon}_i \odot \boldsymbol{\sigma}_i) \odot f_i(\mathbf{h}_{i-1}), \quad (4)$$

where $\boldsymbol{\sigma}_i$ and $\boldsymbol{\mu}_i$ are learnable parameters and $\boldsymbol{\epsilon}_i$ is a random vector sampled from $\mathcal{N}(\mathbf{0}, \mathbf{I})$. In contrast, the function f_i represents a typical, deterministic network layer, meaning the concatenation of a linear transformation (or convolution operation), batch normalization, and some nonlinear activation function. In fact, if we were to fix $\boldsymbol{\mu}_i = \mathbf{1}$ and $\boldsymbol{\sigma}_i = \mathbf{0}$, then the model would default to a regular feed-forward neural network. Additionally, we use \mathbf{W}_i to indicate the weights embedded in f_i , while $\mathbf{W}_{i,j}$ represents the j -th row of \mathbf{W}_i and $\mathbf{W}_{i,j}$ denotes the j -th column. Consequently, $\mathbf{W}_{i,j}$ corresponds with the j -th neuron in the i -th hidden layer, *i.e.*, $h_{i,j}$; likewise for $\mathbf{W}_{i,j}$ and the j -th neuron in the $(i-1)$ -th hidden layer, *i.e.* $h_{i-1,j}$. To avoid unnecessary clutter, we omit the bias in all layers.

With the above definitions in mind, it follows that

$$p(\mathbf{h}_i|\mathbf{h}_{i-1}) = \mathcal{N}(\mathbf{h}_i; f_i(\mathbf{h}_{i-1}) \odot \boldsymbol{\mu}_i, \text{diag}[f_i(\mathbf{h}_{i-1})^2 \odot \boldsymbol{\sigma}_i]). \quad (5)$$

Note that Gaussian noise has previously been multiplied with layer-wise activations as a path towards network regularization (Kingma et al., 2015; Achille & Soatto, 2018); however, these works are not concerned with network compression and other modeling details are significantly different than ours. It has also been used in (Neklyudov et al., 2017) for neuron pruning in conjunction with a truncated, approximate Jeffreys prior. But this requires additional hyper-parameters for balancing the approximation, without which their alternative energy function is ill-defined. Moreover, the exact relationship between these extra parameters and compression performance is not always transparent.

Proceeding further, with our model we simply specify that $q(\mathbf{h}_i)$ is also Gaussian via

$$q(\mathbf{h}_i) = \mathcal{N}(\mathbf{h}_i; \mathbf{0}, \text{diag}[\boldsymbol{\xi}_i]), \quad (6)$$

where $\boldsymbol{\xi}_i$ is an unknown vector of variances that can be learned from the data. Note that if any of these variances are pushed to zero during training, this action will in turn pressure the corresponding coordinates of $p(\mathbf{h}_i|\mathbf{h}_{i-1})$ towards a degenerate Dirac-delta, effectively pruning the associated neuron from the model. As we will later argue, both theoretically and empirically, this form of regularization can serve as a powerful basis for network compression and redundancy reduction.¹ And with regard to practical implementations, it is far easier to exploit network compression instantiated via neural pruning than by, for example, reducing the information content uniformly across all the neurons in a layer.

Our Gaussian assumptions are also advantageous in that they lead to an interpretable, closed-form approximation for the KL term from (3), allowing us to directly optimize $\boldsymbol{\xi}_i$ out of the model. Specifically, following several algebraic manipulations (Dai et al., 2018), we have that

$$\begin{aligned} &\inf_{\boldsymbol{\xi}_i > \mathbf{0}} 2\mathbb{E}_{\mathbf{h}_{i-1} \sim p(\mathbf{h}_{i-1})} [\mathbb{KL}[p(\mathbf{h}_i|\mathbf{h}_{i-1})||q(\mathbf{h}_i)]] \\ &\equiv \sum_j \left[\log \left(1 + \frac{\mu_{i,j}^2}{\sigma_{i,j}^2} \right) + \psi_{i,j} \right], \end{aligned} \quad (7)$$

where

$$\begin{aligned} \psi_{i,j} &\triangleq \log \mathbb{E}_{\mathbf{h}_{i-1} \sim p(\mathbf{h}_{i-1})} [f_{i,j}(\mathbf{h}_{i-1})^2] \\ &\quad - \mathbb{E}_{\mathbf{h}_{i-1} \sim p(\mathbf{h}_{i-1})} [\log f_{i,j}(\mathbf{h}_{i-1})^2] \end{aligned} \quad (8)$$

and $\mu_{i,j}$, $\sigma_{i,j}$, and $f_{i,j}(\mathbf{h}_{i-1})$ denote the j -th element of $\boldsymbol{\mu}_i$, $\boldsymbol{\sigma}_i$, and $f_i(\mathbf{h}_{i-1})$ respectively. The quantity $\psi_{i,j}$ is always positive by Jensen's inequality, but likely to be smaller when the variance of $p(\mathbf{h}_{i-1})$ is not too large and the gap between the log of an expectation and the expectation of the log is reduced. For computational convenience, and because empirically we found the contribution of $\psi_{i,j}$ to be unnecessary for excellent compression performance, we remove

¹In the past a similar prior has been used in conjunction with learning sparse kernel machines (Tipping, 2001).

this factor from the model. By plugging this simplified KL approximation into (3) across all layers i , we obtain the revised final loss function²

$$\tilde{\mathcal{L}} = \sum_{i=1}^L \gamma_i \sum_{j=1}^{r_i} \log \left(1 + \frac{\mu_{i,j}^2}{\sigma_{i,j}^2} \right) - L \mathbb{E}_{\{\mathbf{x}, \mathbf{y}\} \sim \mathcal{D}, \mathbf{h} \sim p(\mathbf{h}|\mathbf{x})} [\log q(\mathbf{y}|\mathbf{h}_L)]. \quad (9)$$

Several items are worth pointing out with respect to this objective. First, the parameters γ_i grant us the flexibility to individually tailor the degree of compression across each layer like some prior methods. While in many situations the simple choice $\gamma_i = \gamma > 0$ for all i is sufficient, in cases where there are significant complexity differences across layers a simple modification can be warranted. Regardless, γ_i serves a useful, transparent purpose, and our energy function requires no additional hyper-parameters as in (Louizos et al., 2017a;b; Neklyudov et al., 2017) to describe priors, approximations, or any other constraints.

Secondly, the weighting factor L on the data term naturally provides balance for deeper networks, preventing the KL factors from accumulating such that the prediction accuracy is completely ignored by the globally optimal solution. In contrast, with many probabilistic network models, a related KL term must be heuristically down-weighted during training (Louizos et al., 2017a; Sønderby et al., 2016), but this then interferes with the associated free-energy bound on the log-likelihood. This is unlike our approach, where L naturally emerges from the formulation itself, and represents an integral part of the variational information bottleneck bounding process.

And finally, although the remaining integrals from (9) have no closed form, unbiased stochastic approximations of the required expectations provide a natural remedy for training purposes (Kingma & Welling, 2014; Rezende et al., 2014). First, a pair $\{\mathbf{x}, \mathbf{y}\}$ is randomly sampled from the training data and fed into the network. For the forward pass, at each layer we sample $\epsilon_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and then compute \mathbf{h}_i via (4). For the backward pass, the gradients can naturally flow via back-propagation to $\{\mu_i, \sigma_i, \mathbf{W}_i\}_{i=1}^L$ and \mathbf{W}_y .

We refer to our model as the *Variational Information Bottleneck Network* (VIBNet).³ The layer-wise sampling strategy

²With slight abuse of notation, we reuse $\tilde{\mathcal{L}}$ to describe this updated objective even though in reality it is no longer a strict upper bound, satisfying only the looser requirement $\tilde{\mathcal{L}} \geq \sum_{i=1}^L \left(\mathcal{L}_i - \gamma_i \sum_{j=1}^{r_i} \psi_{i,j} \right)$. Note though that if the activation function is such that $f_{i,j}(\mathbf{h}_{i-1})^2 \approx 0$ across a region with nonzero probability measure, then the associated $\psi_{i,j}$ can potentially be arbitrarily large, trivializing the bound. Regardless, in later sections we will provide theoretical justification for adopting $\tilde{\mathcal{L}}$ that is independent of the tightness of this approximation anyway. Additionally, for simplicity and without loss of generality we have absorbed the factor of 2 from (7) into each γ_i .

³Code available at <https://github.com/zhuchen03/>




Figure 1. VIBNet Structure. The conditional distribution $p(\mathbf{h}_i|\mathbf{h}_{i-1})$ is given by (5). \mathbf{h}_i is sampled by multiplying $f_i(\mathbf{h}_{i-1})$ with a random variable $\mathbf{z}_i \triangleq \mu_i + \epsilon_i \odot \sigma_i$.

is shown in Figure 1. In subsequent sections we will provide supporting analyses that help to justify our choice of objective function.

4. Reduced Redundancy via Intrinsic Sparsity

Previously we motivated the VIBNet compression model using the concept of the information bottleneck. However, given that multiple bounds/approximations were required to obtain a tractable energy function, it is quite reasonable to question whether or not our original design principles were somehow compromised. To address this concern, both this section and the sequel will attempt to provide independent justification for the VIBNet cost. In this way, we can naturally sidestep issues related to the tightness or legitimacy of the various underlying bounds involved in arriving at (9). From this high-level perspective, we can then view the information bottleneck as having merely provided a form of loose inspiration for a candidate energy function, but one that must still be further subject to careful examination before confidence is warranted.

To begin, recall that (9) is constructed from two factors: a regularizer based on the KL divergence, and a data fit term involving an expectation over latent hidden states. With respect to the former, it is easily shown that $\log(1 + u)$ is a concave, non-decreasing function on the domain $[0, \infty)$, canonical characteristics of a sparsity-promotion regularizer (Chen et al., 2017). Therefore, rather than favoring a solution with many smaller, partially shrunken versions of the ratios

$$\alpha_{i,j} \triangleq \mu_{i,j}^2 \sigma_{i,j}^{-2}, \quad \forall i, j, \quad (10)$$

this type of sparsity archetype instead prefers pushing some percentage to exactly zero while leaving others mostly unchanged (Rao et al., 2003). But how exactly does this favoritism relate to our original information bottleneck criterion? We characterize this relationship with the following result; all proofs are deferred to (Dai et al., 2018):

Proposition 1. *At any minimum of (9), $\alpha_{i,j} = 0$ is a nec-*

VIBNet. Note also that the variational information bottleneck has been referenced in the past as a means of improving generalization performance and robustness to adversarial attacks (Alemi et al., 2016), but this is unrelated to our present purposes here.

essary condition for $\mathbf{I}(h_{i,j}; \mathbf{h}_{i-1}) = 0$ and a sufficient condition for $\mathbf{I}(h_{i,j}; \mathbf{h}_{i-1}) \leq \psi_{i,j}$. Additionally, further assume that there exists a ball of radius $\rho > 0$ around a given minimum such that within this ball, the data term $-\mathbb{E}_{\{\mathbf{x}, \mathbf{y}\} \sim \mathcal{D}, \mathbf{h} \sim p(\mathbf{h}|\mathbf{x})} [\log q(\mathbf{y}|\mathbf{h}_L)]$ is an increasing function of $\sigma_{i,j} \geq 0$. Then $\alpha_{i,j} = 0$ is a sufficient condition for $\mathbf{I}(h_{i,j}; \mathbf{h}_{i-1}) = 0$ at this minimum.

Based on the data processing inequality (Cover & Thomas, 2012) and the Markovian structure from (1), $\mathbf{I}(h_{i,j}; \mathbf{y}) \leq \mathbf{I}(h_{i,j}; \mathbf{h}_{i-1})$. It directly follows that any neuron with $\mathbf{I}(h_{i,j}; \mathbf{h}_{i-1}) = 0$ actually contains no information about \mathbf{y} and is therefore redundant. Hence such a neuron can be safely removed without hurting the predictive performance of the network. And if $\mathbf{I}(h_{i,j}; \mathbf{h}_{i-1}) = 0$, then $\alpha_{i,j} = 0$ per the necessary condition from Proposition 1. Consequently, if we remove all neurons with $\alpha_{i,j} = 0$, we can be sure that any survivors are informative, with $\mathbf{I}(h_{i,j}; \mathbf{h}_{i-1}) > 0$.

But these observations do not address the potential risk of over-pruning useful neurons. Fortunately though, if we manage to minimize the VIBNet cost and find that some $\alpha_{i,j} = 0$, then the sufficiency qualification of Proposition 1 implies that $\mathbf{I}(h_{i,j}; \mathbf{h}_{i-1}) = 0$, provided at least that, within an arbitrarily small neighborhood around this minimum, the data term $-\mathbb{E}_{\{\mathbf{x}, \mathbf{y}\} \sim \mathcal{D}, \mathbf{h} \sim p(\mathbf{h}|\mathbf{x})} [\log q(\mathbf{y}|\mathbf{h}_L)]$ is an increasing function of $\sigma_{i,j}$. This latter condition will provably hold in many special cases, e.g., if $\log q(\mathbf{y}|\mathbf{h}_L)$ is a quadratic function of the activations, but it is likely to be true in broader practical scenarios given that increasing the variance will generally be disruptive to the data fit, and hence increase the negative log-likelihood.

Of course if the weights from the next layer to which $h_{i,j}$ feeds are equal to zero, meaning $\mathbf{W}_{i+1,j} = \mathbf{0}$, then the stated data term can no longer be a strictly increasing function of $\sigma_{i,j}$ within any ball no matter how small (just flat or non-decreasing). However, in this situation, even though technically $\mathbf{I}(h_{i,j}; \mathbf{h}_{i-1})$ may not be provably equal to zero at this minimum, it is irrelevant since no information $h_{i,j}$ retains about \mathbf{h}_{i-1} can be passed on to the next layer. It is therefore still a useless neuron and should be pruned.

This raises a larger point in terms of superfluous information unrelated to the relationship between $\mathbf{I}(h_{i,j}; \mathbf{h}_{i-1})$ and $\alpha_{i,j}$. In general, if a neuron contains some information pertaining to \mathbf{y} , but this information is never inherited by the next layer because $\mathbf{W}_{i+1,j} = \mathbf{0}$, then it is effectively useless and represents a prime candidate for compression. In this case, the corresponding $\alpha_{i,j}$ should necessarily also be zero if it is to serve as an ideal Bellwether for expendable neurons. Fortunately, this is indeed the case with our model:

Proposition 2. *At any minimum of (9), $\alpha_{i,j} = 0$ is a necessary condition for $\mathbf{W}_{i+1,j} = \mathbf{0}$.*

Stated differently, this result implies that if $\mathbf{W}_{i+1,j} = \mathbf{0}$

and $\alpha_{i,j} \neq 0$, then we cannot be at a minimum of (9). And so again, the state $\alpha_{i,j} = 0$ is naturally aligned with our goal of unmasking ineffectual neurons for pruning.

At a high level then, given that our chosen penalty encourages $\alpha_{i,j} \rightarrow 0$, Propositions 1 and 2 then loosely suggest that this process may naturally aggregate useless information into certain expendable neurons, as opposed to distributing the bottleneck equally across all neurons in a layer. And this aggregation strategy provides a simple rule for exposing these redundant neurons such that they can be readily pruned from the model for practical compression purposes: namely, those neurons for which $\alpha_{i,j} \approx 0$ are ideal candidates for removal, and both of the potentially redundant neuron types described above will be discarded based on this criteria (providing further justification for drop-out (Molchanov et al., 2017)). We also conjecture that these relatively straightforward channels for reducing redundancy are indicative of broader mechanisms for compression.

However, there nonetheless remains a lingering issue with this overall line of reasoning. Although the KL-penalty should favor neuron pruning in a generic context as we have argued, the component factors μ and σ are also nonlinearly combined within the neural-network-dependent data term. It therefore remains ambiguous exactly how the suggested sparsity mechanism will operate within this particular, practically-relevant setting. Moreover, it is still unclear how the stochastic, sparsity-promoting objective of VIBNet may exhibit any advantage over standard, deterministic alternatives such as the use of ℓ_1 norm or related penalties. We consider these issues next.

5. Analysis of Tractable Upper Bounds

In general, it is extremely difficult to analyze the complex energy surface of a deep network, and the problem is only compounded when we include the high-dimensional integrals from (9). Fortunately though, certain convenient bounds and analyses inspired by sparse Bayesian methods (Wipf et al., 2011) allow us to nonetheless gain insights into operational behaviors of VIBNet. At an intuitive level, the basic idea here is that if tractable upper bounds can reasonably describe a local neighborhood while displaying useful properties with respect to compression and neural pruning, then we may expect that the underlying energy itself may inherit these desirable attributes, at least to some extent in local regions well-matched to the bounds.

To begin, let $\theta \triangleq \{\mu_i, \sigma_i\}_{i=1}^L$ and $\mathbf{W} \triangleq \{\{\mathbf{W}_i\}_{i=1}^L, \mathbf{W}_y\}$. We then define

$$g(\epsilon; \theta, \mathbf{W}) \triangleq -L \int p(\mathbf{x}, \mathbf{y}) \log q[\mathbf{y}|\mathbf{h}_L(\epsilon, \mathbf{x}; \theta, \mathbf{W})] d\mathbf{x}d\mathbf{y}, \quad (11)$$

where the last hidden layer activation $\mathbf{h}_L(\epsilon, \mathbf{x}; \theta, \mathbf{W})$ is described recursively via (4) and we have explicitly included its dependence on the random variables $\{\epsilon, \mathbf{x}\}$, with $\epsilon \triangleq$

$\{\epsilon_i\}_{i=1}^L$, and the parameters $\{\theta, \mathbf{W}\}$. It then follows that the VIBNet objective from (9) can be re-expressed as

$$\begin{aligned} \tilde{\mathcal{L}}(\theta, \mathbf{W}) &= \int p(\epsilon)g(\epsilon; \theta, \mathbf{W}) d\epsilon \\ &+ \sum_{i=1}^L \gamma_i \sum_{j=1}^{r_i} \log(1 + \alpha_{i,j}). \end{aligned} \quad (12)$$

Note that we have included the parametric dependence of $\tilde{\mathcal{L}}$ on $\{\theta, \mathbf{W}\}$ which serves to clarify certain usages later. Now suppose for any fixed $\mathbf{W} = \mathbf{W}'$ we construct a positive semi-definite quadratic upper bound on g with respect to $\mathbf{z}(\epsilon; \theta) \triangleq \{\mathbf{z}_i(\epsilon_i; \theta_i)\}_{i=1}^L$ (stacked in vectorized form), with $\mathbf{z}_i(\epsilon_i; \theta_i) \triangleq \boldsymbol{\mu}_i + \boldsymbol{\sigma}_i \odot \epsilon_i$. Specifically, this leads to the generic bound $\bar{g}(\epsilon; \theta) \geq g(\epsilon; \theta, \mathbf{W}')$, where

$$\bar{g}(\epsilon; \theta) \triangleq \mathbf{z}(\epsilon; \theta)^\top \mathbf{A}^\top \mathbf{A} \mathbf{z}(\epsilon; \theta) + \mathbf{b}^\top \mathbf{z}(\epsilon; \theta) + c, \quad (13)$$

and for simplicity we have ignored the implicit dependency of \mathbf{A} , \mathbf{b} , and c on the value of \mathbf{W}' . Such a bound is always possible, with $\bar{g}(\epsilon; \theta) = g(\epsilon; \theta, \mathbf{W}')$ for at least some value(s) of $\mathbf{z}(\epsilon; \theta)$ provided that, for example, $g(\epsilon; \theta, \mathbf{W}')$ has Lipschitz continuous gradients. Moreover, (13) can likewise be used to bound the overall cost via

$$\begin{aligned} \tilde{\mathcal{L}}(\theta) &\triangleq \int p(\epsilon) \bar{g}(\epsilon; \theta) d\epsilon + \sum_{i=1}^L \gamma_i \sum_{j=1}^{r_i} \log(1 + \alpha_{i,j}) \\ &\geq \tilde{\mathcal{L}}(\theta, \mathbf{W}') \end{aligned} \quad (14)$$

This leads to the following:

Proposition 3. *If $\theta^* = \{\boldsymbol{\mu}^*, \boldsymbol{\sigma}^*\}$ is a local minimum of the bound $\tilde{\mathcal{L}}(\theta)$ from (14), then*

$$\|\boldsymbol{\mu}^*\|_0 = \|\boldsymbol{\alpha}^*\|_0 \leq \text{rank}[\mathbf{A}] + 1. \quad (15)$$

Here $\|\cdot\|_0$ denotes the ℓ_0 (quasi)-norm, or a count of the number of nonzero elements in a vector. It is typically viewed as the canonical or ideal measure of sparse solutions (Donoho & Elad, 2003). In general, Proposition 3 only provides a non-trivial upper bound on the estimated sparsity of $\boldsymbol{\mu}$ and the ratios $\boldsymbol{\alpha} = \boldsymbol{\mu}^2 \odot \boldsymbol{\sigma}^{-2}$ when $\mathbf{A}^\top \mathbf{A}$ is rank deficient, or equivalently, $\text{rank}[\mathbf{A}] < \dim[\boldsymbol{\mu}]$. However, given an overparameterized neural network where significant compression is possible, we expect that many regions of the energy landscape will be heavily skewed with long valleys of constant cost such that a low rank \mathbf{A} contributes to a reasonable local approximation. So in this situation the upper bound indicates that some neuron pruning will occur even if only the worst local minimizer is obtained. Of course in practice far more significant sparsity is likely.

But there is another more subtle benefit of the VIBNet pruning mechanism: loosely speaking, *the shape/concavity of the implicit VIBNet sparsification effect is automatically calibrated with the local curvature of $g(\epsilon; \theta, \mathbf{W})$ in such a way that the ideal ℓ_0 norm can be adaptively approximated while reducing the risk of bad local minimum.* To better appreciate this claim, it is helpful to first consider a more traditional, deterministic sparsity-based regularization analogue.

Suppose we were to remove the stochastic elements from the bound defined in (14), meaning we set $\boldsymbol{\sigma} = \mathbf{0}$, and we then replaced the KL-based penalty with some generic function π promoting sparse values of $\boldsymbol{\mu}$. The result would be the deterministic energy

$$\Psi_\pi(\boldsymbol{\mu}) \triangleq \boldsymbol{\mu}^\top \mathbf{A}^\top \mathbf{A} \boldsymbol{\mu} + \mathbf{b}^\top \boldsymbol{\mu} + \sum_{i=1}^L \gamma_i \sum_{j=1}^{r_i} \pi(\mu_{i,j}). \quad (16)$$

If $\pi(\mu_{i,j})$ is the non-convex indicator function $\mathcal{I}[\mu_{i,j} \neq 0]$, then a weighted ℓ_0 norm emerges; however, minimizing $\Psi_\pi(\boldsymbol{\mu})$ to obtain a maximally sparse solution is NP-hard because of a combinatorial number of local minima (like-wise for smooth yet non-convex approximations (Chen et al., 2017)). In contrast, if $\pi(\mu_{i,j}) = |\mu_{i,j}|$, then we obtain a weighted ℓ_1 norm regularizer, which represents the tightest convex relaxation of the ℓ_0 norm. While the overall energy is now convex, minimization will often fail to produce maximally sparse (or maximally compressible) estimates except in highly idealized scenarios (Donoho & Elad, 2003). This is because the ℓ_1 norm tends to over-shrink large coefficients at the expense of sparsity (Fan & Li, 2001).

Against this backdrop, we can directly contrast the adaptive, data-dependent VIBNet regularization effect. Let $a_{i,j}$ denote the diagonal element of $\mathbf{A}^\top \mathbf{A}$ corresponding with $\mu_{i,j}$ and define $\omega_{i,j} \triangleq \gamma_i a_{i,j}^{-1}$. Then based on the proof of Proposition 3, it can be shown that bound $\tilde{\mathcal{L}}(\theta)$ satisfies

$$\begin{aligned} \inf_{\boldsymbol{\sigma} > \mathbf{0}} \tilde{\mathcal{L}}(\theta) &= \boldsymbol{\mu}^\top \mathbf{A}^\top \mathbf{A} \boldsymbol{\mu} + \mathbf{b}^\top \boldsymbol{\mu} \\ &+ \sum_{i=1}^L \gamma_i \sum_{j=1}^{r_i} \rho(\mu_{i,j}; \omega_{i,j}), \end{aligned} \quad (17)$$

where $\rho(\mu; \omega)$

$$\triangleq \frac{2|\mu|}{|\mu| + \sqrt{\mu^2 + 4\omega}} + \log\left(2\omega + \mu^2 + |\mu|\sqrt{\mu^2 + 4\omega}\right). \quad (18)$$

Functions like ρ have previously been used for blind image deblurring (Wipf & Zhang, 2014), in part because they provide an attractive means of interpolating between scaled versions of the ℓ_0 norm as $\omega \rightarrow 0$, and the ℓ_1 norm as $\omega \rightarrow \infty$. This interpolation can, for example, be useful in adapting to different blur kernel estimates. Of paramount importance though, in the present context here this interpolation is directly modulated by the parameters $a_{i,j}$, which collectively represent a measure of the local curvature of $\bar{g}(\epsilon; \theta)$ when $\boldsymbol{\sigma} = \mathbf{0}$, *i.e.*, a proxy for the deterministic, data-dependent neural network loss after the stochastic hidden layer latent variables have been removed.

The cumulative effect is that the penalty function shape is roughly matched to this data term. When the latter is relatively smooth and unconstrained, many $a_{i,j}$ values become small within the most representative bound. This pushes the corresponding $\omega_{i,j}$ to be large, and the regularizer is likewise comparably smooth and flat. This helps to avoid aggressive or premature dominance of a highly non-convex sparsity penalty in regions where the deep network energy is relatively flat.

Conversely, if the network’s local region is highly curved and constrained, the bound reflecting local curvature will have many $a_{i,j}$ values that are large, the associated $\omega_{i,j}$ then becomes small, and a more ℓ_0 -norm-like regularizer emerges. But here a stronger penalty can be employed with relatively limited risk of a quick, dominant descent to far away spurious optima. For these reasons, we believe that such a data-dependent, adaptive regularizer is particularly appropriate for compression purposes.

6. Experiments and Discussion

In the majority of recent neural network compression work, models are evaluated with respect to some subset of the following architecture/dataset combinations: LeNet-300-100 (LeCun et al., 1998) and LeNet-5-Caffe⁴ networks on MNIST (LeCun, 1998), and VGG-16 (Simonyan & Zisserman, 2014)⁵ networks on CIFAR10 and CIFAR100 (Krizhevsky & Hinton, 2009). Using all of these benchmarks, we compare our VIBNet with published results from a variety of contemporary state-of-the-art methods including *Generalized Dropout* (GD) (Srinivas & Babu, 2016), *Group Lasso* (GL) (Wen et al., 2016), *Sparse Variational Dropout* (VD) (Molchanov et al., 2017), *Bayesian Compression with Group Normal Jeffreys Prior* (BC-GNJ) and *Group Horseshoe Prior* (BC-GHS) (Louizos et al., 2017a), *Sparse ℓ_0 Regularization* (L0) and L0 with *separate* λ for each layer (L0-sep) (Louizos et al., 2017b), *Drop Neuron* (DN) (Pan et al., 2016), *Runtime Neural Pruning* (RNP) (Lin et al., 2017), *Pruning Filter* (PF) (Li et al., 2016), *Network Slimming* (NS) (Liu et al., 2017), and *Structured Bayesian Pruning* (SBP) and SBP with *KL scaling* (SBPa) (Neklyudov et al., 2017).

Note that GL, DN, and NS all rely on an ℓ_1 -norm-like group-sparsity penalty in some form, and therefore may be at least partially exposed to some of the weaknesses described in Section 5. Likewise, L0-sep is based on a smoothed version of the ℓ_0 -norm obtained via an expectation operator over an additional set of latent variables, and therefore represents another interesting comparison. We also emphasize that because many of the above methods were introduced concurrently, none of them are actually compared against all of the others on standard benchmarks as we do here.

Evaluation Metrics: Beyond classification error on test sets, we also evaluate with respect to three metrics that relate to the compression ratio and model complexity: (i) *Model size* (r_W) - The ratio of number of nonzero weights in the compressed network versus the original model. (ii) *Floating point operations* (FLOPs) - The number of floating point operations required to predict \mathbf{y} from \mathbf{x} during test-

Method	r_W (%)	r_N (%)	error(%)	Pruned Model
VD	25.28	58.95	1.8	512-114-72
BC-GNJ	10.76	32.85	1.8	278-98-13
BC-GHS	10.55	34.71	1.8	311-86-14
L0	26.02	45.02	1.4	219-214-100
L0-sep	10.01	32.69	1.8	266-88-33
DN	23.05	57.94	1.8	542-83-61
VIBNet	3.59	16.98	1.6	97-71-33

Table 1. Compression results on MNIST using LeNet-300-100. VIBNet achieves much better compression than all previous methods while the error rate is nearly the best. Note that for all experiments reported herein, the performance of baseline networks (prior to any compression) can be found in (Dai et al., 2018).

time. (iii) *Run-time memory footprint* (r_N) - The ratio of the space for storing hidden feature maps during run-time in the pruned network versus the original model. This involves calculating the feature map sizes (product of the channel, height, and width) across each layer.

Training and Testing: During training, we set the trade-off parameter γ using a simple heuristic such that VIBNet can roughly match the best previously reported accuracy values. In doing so we obtain a meaningful calibration of the corresponding compression results. See (Dai et al., 2018) for full details regarding this and other aspects of our model training set-up. As for the test phase, we only use the mean values of $p(\mathbf{h}_i|\mathbf{h}_{i-1})$ rather than sampling, which is computationally expensive. Hence we are ultimately only using a probabilistic modeling paradigm to obtain a useful energy function for compressing what amounts to a deterministic network. But the option remains for sampling to improve accuracy as suggested in (Louizos et al., 2017a).

MNIST Results with LeNet-300-100 and LeNet-5: Perhaps the most commonly used pipeline for evaluating existing compression algorithms is MNIST hand-written image data paired with either the LeNet-300-100 or LeNet-5-Caffe network architecture. In evaluations, we follow the conventional training and testing protocols, initializing the weights from scratch like most methods applied to this data. Also since LeNet-300-100 is a fully connected network treating the input as an abstract vector rather than a 2D image, it makes sense to add an additional information bottleneck to the input layer. Results for available models are shown in Table 1, where VIBNet achieves a much smaller r_W and r_N while the error rates for all methods are nearly the same. And the marginal 0.2 accuracy advantage of L0 is offset by the worse compression.

Next we evaluate VIBNet on the LeNet-5-Caffe network, which includes two convolutional layers and two fully connected layers. Results are shown in Table 2. Although VIBNet does not have the lowest r_W (it is the second best in the table), it achieves the lowest FLOP and r_N . The accuracy is almost the same for all methods.

⁴<https://github.com/BVLC/caffe/tree/master/examples/mnist>

⁵The original VGG-16 is applied on 224×224 images. Modified VGG16 versions are used in our experiments.

Compressing Neural Networks using the Variational Information Bottleneck

Method	$r_W(\%)$	FLOP(Mil)	$r_N(\%)$	error(%)
GD	1.38	0.250	32.00	1.1
GL	23.69	0.201	19.35	1.0
VD	9.29	0.660	60.78	1.0
SBP	19.66	0.213	21.15	0.9
BC-GNJ	0.95	0.283	35.03	1.0
BC-GHS	0.64	0.153	22.80	1.0
L0	8.92	1.113	85.82	0.9
L0-sep	1.08	0.389	40.36	1.0
VIBNet	0.83	0.094	15.55	1.0

Table 2. Compression results on MNIST using Lenet-5-Caffe. VIBNet has the smallest FLOPs and r_N , while its r_W is the second best. All methods achieve similar accuracy.

CIFAR10 and CIFAR100 Results Using VGG-16: Evaluations with larger VGG-16 networks on real-world CIFAR10 and CIFAR100 data are complicated by several factors. The primary issue is that, unlike MNIST, it becomes necessary to disentangle various sources of variation unrelated to compression algorithms. For example, different compression pipelines alter network structures and the form of the training data (e.g., data augmentation). Given then that there is no accepted standard for comparison, we evaluate against each competing pipeline individually, training and testing following each different published protocol. Additionally, in all cases models are initialized from a pre-trained network. Hence we obtain three separate VIBNet results for CIFAR10, and two separate results for CIFAR100.

We stress however, that VIBNet can work well under diverse conditions, including training from scratch instead of from a pre-trained network, and our optimal performance (which can be application-dependent) may not ultimately be represented by any of these training protocols originally developed for other models. Nonetheless, VIBNet still obtains improved results in all cases, which speaks to its versatility.

Table 3 displays separate VIBNet results against BC-GNJ/BC-GHS, PF/SBP/SBP_a, and NS-Single/NS-Best following each separate protocol in turn. For the BC models, the dimension of the fully connected layers is simply changed from 4096 to 512, while leaving convolutional layers unaltered (Louizos et al., 2017a). PF and SBP/SBP_a further remove one additional fully-connected layer (Li et al., 2016) and apply standard CIFAR10 data augmentation (e.g., cropping and flipping). Finally, the NS model replaces two fully connected layers with three convolutional layers, which can improve accuracy at the expense of FLOPs and model size (Liu et al., 2017); data augmentation is also used.

From the Table 3, VIBNet achieves the best performance across all three compression metrics, and comparable or better accuracy as well. SBP_a achieves the second best compression, but this comes at the significant cost of a nearly 50% decrease in accuracy. Note that the NS algorithm involves multiple iterations of training and pruning; however,

Method	$r_W(\%)$	FLOP(Mil)	$r_N(\%)$	error(%)
BC-GNJ	6.57	141.5	81.68	8.6
BC-GHS	5.40	121.9	74.82	9.0
VIBNet	5.30	70.63	49.57	8.8 (8.5)
PF	35.99	206.3	83.97	6.6
SBP	7.01	136.0	80.72	7.5
SBP _a	5.78	99.20	66.46	9.0
VIBNet	5.45	86.82	57.86	6.5 (6.1)
NS-Single	11.50	195.5	-	6.2
NS-Best	8.60	147.0	-	5.9
VIBNet	5.79	116.0	59.60	6.2 (5.8)

Table 3. Compression results on CIFAR10 using VGG-16. We compare VIBNet to three different methods, in each case adopting the the training protocols of the original work. Although accuracy measures are similar, VIBNet produces the best compression by a significant margin. Error rates in parentheses were obtained by fine-tuning the pruned architecture. Note also that NS-Best involves multiple iterations of training and fine-tuning.

Method	$r_W(\%)$	FLOP(Mil)	$r_N(\%)$	error(%)
RNP	-	160	-	38.0
VIBNet	22.75	133.6	59.80	37.6 (37.4)
NS-Single	24.90	250.5	-	26.5
NS-Best	20.80	214.8	-	26.0
VIBNet	15.08	203.1	73.80	25.9 (25.7)

Table 4. Compression results on CIFAR100 using VGG-16. VIBNet is compared with two different protocols adopted from the corresponding prior work. Again, while accuracy measures are similar, VIBNet produces the best compression by a significant margin. The error rate in parentheses was obtained by fine-tuning the pruned architecture.

if this procedure is carried out too far, the accuracy drops significantly as the compression increases. Hence NS-Best refers to the iteration result with the best accuracy, while NS-Single refers to the first iteration, a comparable training process to VIBNet. Additionally, multiple iterations of training can be tedious in practice, especially since we do not know in advance how many iterations will be necessary. In contrast, lower accuracy with higher compression can be achieved via VIBNet by simply increasing γ .

Lastly, we compare performance on CIFAR100 against RNP and NS, where RNP (Lin et al., 2017) applies a similar network adaptation as PF on CIFAR10 without data augmentation, and NS is as described above. Result are shown in Table 4, with VIBNet showing consistent improvements.

Discussion: We have compared our approach against arguably the largest number of competing methods primarily designed to prune neurons for computational and memory efficiency, achieving comparable accuracy with improved compression. For additional experiments elucidating how the intrinsic sparsity mechanism of VIBNet contributes to a layer-wise decrease in mutual information (consistent with the information bottleneck), please see (Dai et al., 2018).

References

- Achille, A. and Soatto, S. Information dropout: Learning optimal representations through noisy computation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, To appear, 2018.
- Alemi, A. A., Fischer, I., Dillon, J. V., and Murphy, K. Deep variational information bottleneck. *arXiv:1612.00410*, 2016.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural networks. *arXiv:1505.05424*, 2015.
- Chen, Y., Ge, D., Wang, M., Wang, Z., Ye, Y., and Yin, H. Strong np-hardness for sparse optimization with concave penalty functions. *International Conference on Machine Learning (ICML)*, 2017.
- Courbariaux, M., Bengio, Y., and David, J.-P. Binaryconnect: Training deep neural networks with binary weights during propagations. *Advances in neural information processing systems (NIPS)*, 2015.
- Courbariaux, M., Hubara, I., Soudry, D., El-Yaniv, R., and Bengio, Y. Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1. *arXiv:1602.02830*, 2016.
- Cover, T. M. and Thomas, J. A. *Elements of information theory*, pp. 32–33. John Wiley & Sons, 2012.
- Dai, B., Zhu, C., and Wipf, D. Compressing neural networks using the variational information bottleneck. *arXiv:1802.10399*, 2018.
- Denil, M., Shakibi, B., Dinh, L., Ranzato, M., and De Freitas, N. Predicting parameters in deep learning. *Advances in neural information processing systems (NIPS)*, 2013.
- Dong, X., Huang, J., Yang, Y., and Yan, S. More is less: A more complicated network with less inference complexity. *arXiv:1703.08651*, 2017.
- Donoho, D. L. and Elad, M. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization. *Proc. National Academy of Sciences*, 100(5), 2003.
- Fan, J. and Li, R. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.
- Graves, A. Practical variational inference for neural networks. *Advances in neural information processing systems (NIPS)*, 2011.
- Guo, Y., Yao, A., and Chen, Y. Dynamic network surgery for efficient dnns. *Advances in neural information processing systems (NIPS)*, 2016.
- Han, S., Mao, H., and Dally, W. J. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *Advances in Neural Information Processing Systems (NIPS)*, 2015a.
- Han, S., Pool, J., Tran, J., and Dally, W. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems (NIPS)*, 2015b.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv:1704.04861*, 2017.
- Iandola, F. N., Han, S., Moskewicz, M. W., Ashraf, K., Dally, W. J., and Keutzer, K. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size. *arXiv:1602.07360*, 2016.
- Jaderberg, M., Vedaldi, A., and Zisserman, A. Speeding up convolutional neural networks with low rank expansions. *arXiv:1405.3866*, 2014.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *International Conference on Learning Representations (ICLR)*, 2014.
- Kingma, D. P., Salimans, T., and Welling, M. Variational dropout and the local reparameterization trick. *Advances in neural information processing systems (NIPS)*, 2015.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. *Master's thesis, Department of Computer Science, University of Toronto*, 2009.
- LeCun, Y. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- LeCun, Y., Denker, J. S., and Solla, S. A. Optimal brain damage. *Advances in Neural Information Processing Systems (NIPS)*, 1990.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Li, H., Kadav, A., Durdanovic, I., Samet, H., and Graf, H. P. Pruning filters for efficient convnets. *arXiv:1608.08710*, 2016.
- Lin, J., Rao, Y., Lu, J., and Zhou, J. Runtime neural pruning. *Advances in neural information processing systems (NIPS)*, 2017.

- Liu, Z., Li, J., Shen, Z., Huang, G., Yan, S., and Zhang, C. Learning efficient convolutional networks through network slimming. *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- Louizos, C., Ullrich, K., and Welling, M. Bayesian compression for deep learning. *Advances in Neural Information Processing Systems (NIPS)*, 2017a.
- Louizos, C., Welling, M., and Kingma, D. P. Learning sparse neural networks through ℓ_0 regularization. *arXiv:1712.01312*, 2017b.
- Mellempudi, N., Kundu, A., Mudigere, D., Das, D., Kaul, B., and Dubey, P. Ternary neural networks with fine-grained quantization. *arXiv:1705.01462*, 2017.
- Molchanov, D., Ashukha, A., and Vetrov, D. Variational dropout sparsifies deep neural networks. *arXiv:1701.05369*, 2017.
- Nalisnick, E., Anandkumar, A., and Smyth, P. A scale mixture perspective of multiplicative noise in neural networks. *arXiv:1506.03208*, 2015.
- Neklyudov, K., Molchanov, D., Ashukha, A., and Vetrov, D. P. Structured bayesian pruning via log-normal multiplicative noise. *Advances in neural information processing systems (NIPS)*, 2017.
- Pan, W., Dong, H., and Guo, Y. Dropneuron: Simplifying the structure of deep neural networks. *arXiv:1606.07326*, 2016.
- Rao, B. D., Engan, K., Cotter, S. F., Palmer, J., and Kretz-Delgado, K. Subset selection in noise based on diversity measure minimization. *IEEE Trans. Signal Processing*, 51(3):760–770, 2003.
- Rastegari, M., Ordonez, V., Redmon, J., and Farhadi, A. Xnor-net: Imagenet classification using binary convolutional neural networks. *European Conference on Computer Vision (ECCV)*, 2016.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. *International Conference on Machine Learning (ICML)*, 2014.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014.
- Sønderby, C. K., Raiko, T., Maaløe, L., Sønderby, S. K., and Winther, O. How to train deep variational autoencoders and probabilistic ladder networks. *arXiv preprint arXiv:1602.02282*, 2016.
- Srinivas, S. and Babu, R. V. Generalized dropout. *arXiv:1611.06791*, 2016.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Tipping, M. E. Sparse Bayesian learning and the relevance vector machine. *Journal of Machine Learning Research*, 1(6):211–244, 2001.
- Tishby, N. and Zaslavsky, N. Deep learning and the information bottleneck principle. *IEEE Information Theory Workshop (ITW)*, 2015.
- Tishby, N., Pereira, F. C., and Bialek, W. The information bottleneck method. *arXiv:physics/0004057*, 2000.
- Ullrich, K., Meeds, E., and Welling, M. Soft weight-sharing for neural network compression. *arXiv:1702.04008*, 2017.
- Wen, W., Wu, C., Wang, Y., Chen, Y., and Li, H. Learning structured sparsity in deep neural networks. *Advances in neural information processing systems (NIPS)*, 2016.
- Wipf, D. and Zhang, H. Revisiting bayesian blind deconvolution. *Journal of Machine Learning Research*, 15(1):3595–3634, 2014.
- Wipf, D. P., Rao, B. D., and Nagarajan, S. Latent variable Bayesian models for promoting sparsity. *IEEE Trans. Information Theory*, 57(9):6236–6255, 2011.
- Yu, X., Liu, T., Wang, X., and Tao, D. On compressing deep models by low rank and sparse decomposition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- Zhang, X., Zou, J., He, K., and Sun, J. Accelerating very deep convolutional networks for classification and detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(10):1943–1955, 2016.