
Variational Dropout Sparsifies Deep Neural Networks

Dmitry Molchanov^{1,2*} Arsenii Ashukha^{3,4*} Dmitry Vetrov^{3,1}

Abstract

We explore a recently proposed Variational Dropout technique that provided an elegant Bayesian interpretation to Gaussian Dropout. We extend Variational Dropout to the case when dropout rates are unbounded, propose a way to reduce the variance of the gradient estimator and report first experimental results with individual dropout rates per weight. Interestingly, it leads to extremely sparse solutions both in fully-connected and convolutional layers. This effect is similar to automatic relevance determination effect in empirical Bayes but has a number of advantages. We reduce the number of parameters up to 280 times on LeNet architectures and up to 68 times on VGG-like networks with a negligible decrease of accuracy.

1. Introduction

Deep neural networks (DNNs) are a widely popular family of models which is currently state-of-the-art in many important problems (Szegedy et al., 2016; Silver et al., 2016). However, DNNs often have many more parameters than the number of the training instances. This makes them prone to overfitting (Hinton et al., 2012; Zhang et al., 2016) and necessitates using regularization. A commonly used regularizer is Binary Dropout (Hinton et al., 2012) that prevents co-adaptation of neurons by randomly dropping them during training. An equally effective alternative is Gaussian Dropout (Srivastava et al., 2014) that multiplies the outputs of the neurons by Gaussian random noise.

Dropout requires specifying the dropout rates which are the

*Equal contribution ¹Yandex, Russia ²Skolkovo Institute of Science and Technology, Skolkovo Innovation Center, Moscow, Russia ³National Research University Higher School of Economics, Moscow, Russia ⁴Moscow Institute of Physics and Technology, Moscow, Russia. Correspondence to: Dmitry Molchanov <dmitry.molchanov@skolkovotech.ru>, Arsenii Ashukha <ars.ashuha@gmail.com>, Dmitry Vetrov <vetrovd@yandex.ru>.

probabilities of dropping a neuron. The dropout rates are typically optimized using grid search. To avoid the exponential complexity of optimizing multiple hyperparameters, the dropout rates are usually shared for all layers. Recently it was shown that dropout can be seen as a special case of Bayesian regularization (Gal & Ghahramani, 2015; Kingma et al., 2015). It is an important theoretical result that justifies dropout and at the same time allows us to tune individual dropout rates for each weight, neuron or layer in a Bayesian way.

Instead of injecting noise we can regularize a model by reducing the number of its parameters. This technique is especially attractive in the case of deep neural networks. Modern neural networks contain hundreds of millions of parameters (Szegedy et al., 2015; He et al., 2015) and require a lot of computational and memory resources. It restricts us from using deep neural networks when those resources are limited. Inducing sparsity during training of DNNs leads to regularization, compression, and acceleration of the resulting model (Han et al., 2015a; Scardapane et al., 2016).

Sparse Bayesian Learning (Tipping, 2001) provides a principled framework for training of sparse models without the manual tuning of hyperparameters. Unfortunately, this approach does not extend straightforwardly to DNNs. During past several years, a number of papers (Hoffman et al., 2013; Kingma & Welling, 2013; Rezende et al., 2014) on scalable variational inference have appeared. These techniques make it possible to train Bayesian Deep Neural Networks using stochastic optimization and provide us an opportunity to transfer Bayesian regularization techniques from simple models to DNNs.

In this paper, we study Variational Dropout (Kingma et al., 2015) in the case when each weight of a model has its individual dropout rate. We propose Sparse Variational Dropout that extends Variational Dropout to all possible values of dropout rates and leads to a sparse solution. To achieve this goal, we provide a new approximation of the KL-divergence term in Variational Dropout objective that is tight on the full domain. We also propose a way to greatly reduce the variance of the stochastic gradient estimator and show that it leads to a much faster convergence and a better value of the objective function. We show theoretically that

Sparse Variational Dropout applied to linear models can lead to a sparse solution. Like classical Sparse Bayesian models, our method provides the Automatic Relevance Determination effect, but overcomes certain disadvantages of empirical Bayes.

Our experiments show that Sparse Variational Dropout leads to a high level of sparsity in fully-connected and convolutional layers of Deep Neural Networks. Our method achieves a state-of-the-art sparsity level on LeNet architectures and scales on larger networks like VGG with negligible performance drop. Also we show that our method fails to overfit on randomly labeled data unlike Binary Dropout networks.

2. Related Work

Deep Neural Nets are prone to overfitting and regularization is used to address this problem. Several successful techniques have been proposed for DNN regularization, among them are Dropout (Srivastava et al., 2014), Drop-Connect (Wan et al., 2013), Max Norm Constraint (Srivastava et al., 2014), Batch Normalization (Ioffe & Szegedy, 2015), etc.

Another way to regularize deep model is to reduce the number of parameters. One possible approach is to use tensor decompositions (Novikov et al., 2015; Garipov et al., 2016). Another approach is to induce sparsity into weight matrices. Most recent works on sparse neural networks use pruning (Han et al., 2015b), elastic net regularization (Lebedev & Lempitsky, 2015; Liu et al., 2015; Scardapane et al., 2016; Wen et al., 2016) or composite techniques (Han et al., 2015a; Guo et al., 2016; Ullrich et al., 2017).

Sparsity can also be obtained by using the Sparse Bayesian Learning framework (Tipping, 2001). Automatic Relevance Determination was introduced in (Neal, 1996; MacKay et al., 1994), where small neural networks were trained with ARD regularization on the input layer. This approach was later studied on linear models like the Relevance Vector Machine (Tipping, 2001) and other kernel methods (Van Gestel et al., 2001). In the Relevance Tagging Machine model (Molchanov et al., 2015) Beta prior distribution is used to obtain the ARD effect in a similar setting.

Recent works on Bayesian DNNs (Kingma & Welling, 2013; Rezende et al., 2014; Scardapane et al., 2016) provide different ways to train deep models with a huge number of parameters in a Bayesian way. These techniques can be applied to improve latent variables models (Kingma & Welling, 2013), to prevent overfitting and to obtain model uncertainty (Gal & Ghahramani, 2015). Recently several works on efficient training of Sparse Bayesian Models have appeared (Challis & Barber, 2013; Titsias & Lázaro-

Gredilla, 2014). Soft Weights Sharing (Ullrich et al., 2017) uses an approach, similar to Sparse Bayesian Learning framework, to obtain a sparse and quantized Bayesian Deep Neural Network, but utilizes a more flexible family of prior distributions.

Variational Dropout (Kingma et al., 2015) is an elegant interpretation of Gaussian Dropout as a special case of Bayesian regularization. This technique allows us to tune dropout rate and can, in theory, be used to set individual dropout rates for each layer, neuron or even weight. However, that paper uses a limited family for posterior approximation that does not allow for ARD effect. Other Bayesian interpretations of dropout training have also appeared during past several years (Maeda, 2014; Gal & Ghahramani, 2015; Srinivas & Babu, 2016). Generalized Dropout (Srinivas & Babu, 2016) provides a way to tune individual dropout rates for neurons, but uses a biased gradient estimator. Also, the posterior distribution is modelled by a delta function, so the resulting neural network is effectively not Bayesian. Variational Spike-and-Slab Neural Networks (Louizos, 2015) is yet another Bayesian interpretation of Binary Dropout that allows for tuning of individual dropout rates and also leads to a sparse solution. Unfortunately, this procedure does not scale well with model width and depth.

3. Preliminaries

We begin by describing the Bayesian Inference and Stochastic Variational Inference frameworks. Then we describe Variational Dropout, a recently proposed Bayesian regularization technique (Kingma et al., 2015).

3.1. Bayesian Inference

Consider a dataset \mathcal{D} which is constructed from N pairs of objects $(x_n, y_n)_{n=1}^N$. Our goal is to tune the parameters w of a model $p(y|x, w)$ that predicts y given x and w . In Bayesian Learning we usually have some prior knowledge about weights w , which is expressed in terms of a prior distribution $p(w)$. After data \mathcal{D} arrives, this prior distribution is transformed into a posterior distribution $p(w|\mathcal{D}) = p(\mathcal{D}|w)p(w)/p(\mathcal{D})$. This process is called *Bayesian Inference*. Computing posterior distribution using the Bayes rule usually involves computation of intractable multidimensional integrals, so we need to use approximation techniques.

One of such techniques is *Variational Inference*. In this approach the posterior distribution $p(w|\mathcal{D})$ is approximated by a parametric distribution $q_\phi(w)$. The quality of this approximation is measured in terms of the Kullback-Leibler divergence $D_{KL}(q_\phi(w) \| p(w|\mathcal{D}))$. The optimal value of variational parameters ϕ can be found by maximization of

the *variational lower bound*:

$$\mathcal{L}(\phi) = L_{\mathcal{D}}(\phi) - D_{KL}(q_{\phi}(w) \parallel p(w)) \rightarrow \max_{\phi \in \Phi} \quad (1)$$

$$L_{\mathcal{D}}(\phi) = \sum_{n=1}^N \mathbb{E}_{q_{\phi}(w)} [\log p(y_n | x_n, w)] \quad (2)$$

It consists of two parts, the expected log-likelihood $L_{\mathcal{D}}(\phi)$ and the KL-divergence $D_{KL}(q_{\phi}(w) \parallel p(w))$, which acts as a regularization term.

3.2. Stochastic Variational Inference

In the case of complex models expectations in (1) and (2) are intractable. Therefore the variational lower bound (1) and its gradients can not be computed exactly. However, it is still possible to estimate them using sampling and optimize the variational lower bound using stochastic optimization.

We follow (Kingma & Welling, 2013) and use the Reparameterization Trick to obtain an unbiased differentiable minibatch-based Monte Carlo estimator of the expected log-likelihood (3). The main idea is to represent the parametric noise $q_{\phi}(w)$ as a deterministic differentiable function $w = f(\phi, \epsilon)$ of a non-parametric noise $\epsilon \sim p(\epsilon)$. This trick allows us to obtain an unbiased estimate of $\nabla_{\phi} L_{\mathcal{D}}(q_{\phi})$. Here we denote objects from a mini-batch as $(\tilde{x}_m, \tilde{y}_m)_{m=1}^M$.

$$\mathcal{L}(\phi) \simeq \mathcal{L}^{SGVB}(\phi) = L_{\mathcal{D}}^{SGVB}(\phi) - D_{KL}(q_{\phi}(w) \parallel p(w)) \quad (3)$$

$$L_{\mathcal{D}}(\phi) \simeq L_{\mathcal{D}}^{SGVB}(\phi) = \frac{N}{M} \sum_{m=1}^M \log p(\tilde{y}_m | \tilde{x}_m, f(\phi, \epsilon_m)) \quad (4)$$

$$\nabla_{\phi} L_{\mathcal{D}}(\phi) \simeq \frac{N}{M} \sum_{m=1}^M \nabla_{\phi} \log p(\tilde{y}_m | \tilde{x}_m, f(\phi, \epsilon_m)) \quad (5)$$

The Local Reparameterization Trick is another technique that reduces the variance of this gradient estimator even further (Kingma et al., 2015). The idea is to sample separate weight matrices for each data-point inside mini-batch. It is computationally hard to do it straight-forwardly, but it can be done efficiently by moving the noise from weights to activations (Wang & Manning, 2013; Kingma et al., 2015).

3.3. Variational Dropout

In this section we consider a single fully-connected layer with I input neurons and O output neurons before a non-linearity. We denote an output matrix as $B^{M \times O}$, input matrix as $A^{M \times I}$ and a weight matrix as $W^{I \times O}$. We index the elements of these matrices as b_{mj} , a_{mi} and w_{ij} respectively. Then $B = AW$.

Dropout is one of the most popular regularization methods for deep neural networks. It injects a multiplicative random

noise Ξ to the layer input A at each iteration of training procedure (Hinton et al., 2012).

$$B = (A \odot \Xi)W, \text{ with } \xi_{mi} \sim p(\xi) \quad (6)$$

The original version of dropout, so-called Bernoulli or Binary Dropout, was presented with $\xi_{mi} \sim \text{Bernoulli}(1-p)$ (Hinton et al., 2012). It means that each element of the input matrix is put to zero with probability p , also known as a *dropout rate*. Later the same authors reported that Gaussian Dropout with continuous noise $\xi_{mi} \sim \mathcal{N}(1, \alpha = \frac{p}{1-p})$ works as well and is similar to Binary Dropout with dropout rate p (Srivastava et al., 2014). It is beneficial to use continuous noise instead of discrete one because multiplying the inputs by a Gaussian noise is equivalent to putting Gaussian noise on the weights. This procedure can be used to obtain a posterior distribution over the model's weights (Wang & Manning, 2013; Kingma et al., 2015). That is, putting multiplicative Gaussian noise $\xi_{ij} \sim \mathcal{N}(1, \alpha)$ on a weight w_{ij} is equivalent to sampling of w_{ij} from $q(w_{ij} | \theta_{ij}, \alpha) = \mathcal{N}(w_{ij} | \theta_{ij}, \alpha \theta_{ij}^2)$. Now w_{ij} becomes a random variable parametrized by θ_{ij} .

$$w_{ij} = \theta_{ij} \xi_{ij} = \theta_{ij} (1 + \sqrt{\alpha} \epsilon_{ij}) \sim \mathcal{N}(w_{ij} | \theta_{ij}, \alpha \theta_{ij}^2) \quad (7)$$

$$\epsilon_{ij} \sim \mathcal{N}(0, 1)$$

Gaussian Dropout training is equivalent to stochastic optimization of the expected log likelihood (2) in the case when we use the reparameterization trick and draw a single sample $W \sim q(W | \theta, \alpha)$ per minibatch to estimate the expectation. Variational Dropout extends this technique and explicitly uses $q(W | \theta, \alpha)$ as an approximate posterior distribution for a model with a special prior on the weights. The parameters θ and α of the distribution $q(W | \theta, \alpha)$ are tuned via stochastic variational inference, i.e. $\phi = (\theta, \alpha)$ are the variational parameters, as denoted in Section 3.2. The prior distribution $p(W)$ is chosen to be improper log-scale uniform to make the Variational Dropout with fixed α equivalent to Gaussian Dropout (Kingma et al., 2015).

$$p(\log |w_{ij}|) = \text{const} \Leftrightarrow p(|w_{ij}|) \propto \frac{1}{|w_{ij}|} \quad (8)$$

In this model, it is the only prior distribution that makes variational inference consistent with Gaussian Dropout (Kingma et al., 2015). When parameter α is fixed, the $D_{KL}(q(W | \theta, \alpha) \parallel p(W))$ term in the variational lower bound (1) does not depend on θ (Kingma et al., 2015). Maximization of the variational lower bound (1) then becomes equivalent to maximization of the expected log-likelihood (2) with fixed parameter α . It means that Gaussian Dropout training is exactly equivalent to Variational Dropout with fixed α . However, Variational Dropout provides a way to train dropout rate α by optimizing the variational lower bound (1). Interestingly, dropout rate α now

becomes a variational parameter and not a hyperparameter. In theory, it allows us to train individual dropout rates α_{ij} for each layer, neuron or even weight (Kingma et al., 2015). However, no experimental results concerning the training of individual dropout rates were reported in the original paper. Also, the approximate posterior family was manually restricted to the case $\alpha \leq 1$.

4. Sparse Variational Dropout

In the original paper, authors reported difficulties in training the model with large values of dropout rates α (Kingma et al., 2015) and only considered the case of $\alpha \leq 1$, which corresponds to a binary dropout rate $p \leq 0.5$. However, the case of large α_{ij} is very exciting (here we mean separate α_{ij} per weight). High dropout rate $\alpha_{ij} \rightarrow +\infty$ corresponds to a binary dropout rate that approaches $p = 1$. It effectively means that the corresponding weight or neuron is always ignored and can be removed from the model. In this work, we consider the case of individual α_{ij} for each weight of the model.

4.1. Additive Noise Reparameterization

Training Neural Networks with Variational Dropout is difficult when dropout rates α_{ij} are large because of a huge variance of stochastic gradients (Kingma et al., 2015). The cause of large gradient variance arises from multiplicative noise. To see it clearly, we can rewrite the gradient of \mathcal{L}^{SGVB} w.r.t. θ_{ij} as follows.

$$\frac{\partial \mathcal{L}^{SGVB}}{\partial \theta_{ij}} = \frac{\partial \mathcal{L}^{SGVB}}{\partial w_{ij}} \cdot \frac{\partial w_{ij}}{\partial \theta_{ij}} \quad (9)$$

In the case of original parametrization (θ, α) the second multiplier in (9) is very noisy if α_{ij} is large.

$$\begin{aligned} w_{ij} &= \theta_{ij}(1 + \sqrt{\alpha_{ij}} \cdot \epsilon_{ij}), \\ \frac{\partial w_{ij}}{\partial \theta_{ij}} &= 1 + \sqrt{\alpha_{ij}} \cdot \epsilon_{ij}, \\ \epsilon_{ij} &\sim \mathcal{N}(0, 1) \end{aligned} \quad (10)$$

We propose a trick that allows us to drastically reduce the variance of this term in the case when α_{ij} is large. The idea is to replace the multiplicative noise term $1 + \sqrt{\alpha_{ij}} \cdot \epsilon_{ij}$ with an exactly equivalent additive noise term $\sigma_{ij} \cdot \epsilon_{ij}$, where $\sigma_{ij}^2 = \alpha_{ij} \theta_{ij}^2$ is treated as a new independent variable. After this trick we will optimize the variational lower bound w.r.t. (θ, σ) . However, we will still use α throughout the paper, as it has a nice interpretation as a dropout rate.

$$\begin{aligned} w_{ij} &= \theta_{ij}(1 + \sqrt{\alpha_{ij}} \cdot \epsilon_{ij}) = \theta_{ij} + \sigma_{ij} \cdot \epsilon_{ij} \\ \frac{\partial w_{ij}}{\partial \theta_{ij}} &= 1, \quad \epsilon_{ij} \sim \mathcal{N}(0, 1) \end{aligned} \quad (11)$$

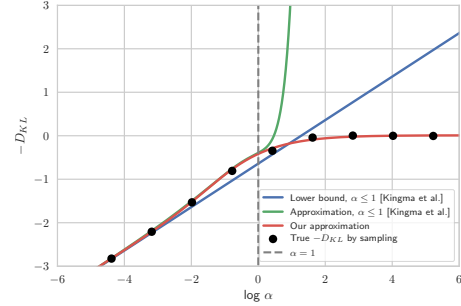


Figure 1: Different approximations of KL divergence: blue and green ones (Kingma et al., 2015) are tight only for $\alpha \leq 1$; black one is the true value, estimated by sampling; red one is our approximation.

From (11) we can see that $\frac{\partial w_{ij}}{\partial \theta_{ij}}$ now has no injected noise, but the distribution over $w_{ij} \sim q(w_{ij} | \theta_{ij}, \sigma_{ij}^2)$ remains exactly the same. The objective function and the posterior approximating family are unaltered. The only thing that changed is the parametrization of the approximate posterior. However, the variance of a stochastic gradient is greatly reduced. Using this trick, we avoid the problem of large gradient variance and can train the model within the full range of $\alpha_{ij} \in (0, +\infty)$.

It should be noted that the Local Reparametrization Trick does not depend on parametrization, so it can also be applied here to reduce the variance even further. In our experiments, we use both Additive Noise Reparameterization and the Local Reparameterization Trick. We provide the final expressions for the outputs of fully-connected and convolutional layers for our model in Section 4.4.

4.2. Approximation of the KL Divergence

As the prior and the approximate posterior are fully factorized, the full KL-divergence term in the lower bound (1) can be decomposed into a sum:

$$\begin{aligned} D_{KL}(q(W | \theta, \alpha) \| p(W)) &= \\ &= \sum_{ij} D_{KL}(q(w_{ij} | \theta_{ij}, \alpha_{ij}) \| p(w_{ij})) \end{aligned} \quad (12)$$

The log-scale uniform prior distribution is an improper prior, so the KL divergence can only be calculated up to an additive constant C (Kingma et al., 2015).

$$\begin{aligned} -D_{KL}(q(w_{ij} | \theta_{ij}, \alpha_{ij}) \| p(w_{ij})) &= \\ &= \frac{1}{2} \log \alpha_{ij} - \mathbb{E}_{\epsilon \sim \mathcal{N}(1, \alpha_{ij})} \log |\epsilon| + C \end{aligned} \quad (13)$$

In the Variational Dropout model this term is intractable, as the expectation $\mathbb{E}_{\epsilon \sim \mathcal{N}(1, \alpha_{ij})} \log |\epsilon|$ in (13) cannot be computed analytically (Kingma et al., 2015). However, this

term can be sampled and then approximated. Two different approximations were provided in the original paper, however they are accurate only for small values of the dropout rate α ($\alpha \leq 1$). We propose another approximation (14) that is tight for all values of alpha. Here $\sigma(\cdot)$ denotes the sigmoid function. Different approximations and the true value of $-D_{KL}$ are presented in Fig. 1. Original $-D_{KL}$ was obtained by averaging over 10^7 samples of ϵ with less than 2×10^{-3} variance of the estimation.

$$\begin{aligned} -D_{KL}(q(w_{ij} | \theta_{ij}, \alpha_{ij}) \| p(w_{ij})) &\approx \\ &\approx k_1 \sigma(k_2 + k_3 \log \alpha_{ij}) - 0.5 \log(1 + \alpha_{ij}^{-1}) + C \quad (14) \\ k_1 &= 0.63576 \quad k_2 = 1.87320 \quad k_3 = 1.48695 \end{aligned}$$

We used the following intuition to obtain this formula. The negative KL-divergence goes to a constant as $\log \alpha_{ij}$ goes to infinity, and tends to $0.5 \log \alpha_{ij}$ as $\log \alpha_{ij}$ goes to minus infinity. We model this behaviour with $-0.5 \log(1 + \alpha_{ij}^{-1})$. We found that the remainder $-D_{KL} + 0.5 \log(1 + \alpha_{ij}^{-1})$ looks very similar to a sigmoid function of $\log \alpha_{ij}$, so we fit its linear transformation $k_1 \sigma(k_2 + k_3 \log \alpha_{ij})$ to this curve. We observe that this approximation is extremely accurate (less than 0.009 maximum absolute deviation on the full range of $\log \alpha_{ij} \in (-\infty, +\infty)$); the original approximation (Kingma et al., 2015) has 0.04 maximum absolute deviation with $\log \alpha_{ij} \in (-\infty, 0]$.

One should notice that as α approaches infinity, the KL-divergence approaches a constant. As in this model the KL-divergence is defined up to an additive constant, it is convenient to choose $C = -k_1$ so that the KL-divergence goes to zero when α goes to infinity. It allows us to compare values of \mathcal{L}^{SGVB} for neural networks of different sizes.

4.3. Sparsity

From the Fig. 1 one can see that $-D_{KL}$ term increases with the growth of α . It means that this regularization term favors large values of α .

The case of $\alpha_{ij} \rightarrow \infty$ corresponds to a Binary Dropout rate $p_{ij} \rightarrow 1$ (recall $\alpha = \frac{p}{1-p}$). Intuitively it means that the corresponding weight is almost always dropped from the model. Therefore its value does not influence the model during the training phase and is put to zero during the testing phase.

We can also look at this situation from another angle. Infinitely large α_{ij} corresponds to infinitely large multiplicative noise in w_{ij} . It means that the value of this weight will be completely random and its magnitude will be unbounded. It will corrupt the model prediction and decrease the expected log likelihood. Therefore it is beneficial to put the corresponding weight θ_{ij} to zero in such a way that $\alpha_{ij} \theta_{ij}^2$ goes to zero as well. It means that $q(w_{ij} | \theta_{ij}, \alpha_{ij})$

is effectively a delta function, centered at zero $\delta(w_{ij})$.

$$\begin{aligned} \theta_{ij} \rightarrow 0, \quad \alpha_{ij} \theta_{ij}^2 \rightarrow 0 \\ \Downarrow \\ q(w_{ij} | \theta_{ij}, \alpha_{ij}) \rightarrow \mathcal{N}(w_{ij} | 0, 0) = \delta(w_{ij}) \end{aligned} \quad (15)$$

In the case of linear regression this fact can be shown analytically. We denote a data matrix as $X^{N \times D}$ and $\alpha, \theta \in \mathbb{R}^D$. If α is fixed, the optimal value of θ can also be obtained in a closed form.

$$\theta = (X^\top X + \text{diag}(X^\top X) \text{diag}(\alpha))^{-1} X^\top y \quad (16)$$

Assume that $(X^\top X)_{ii} \neq 0$, so that i -th feature is not a constant zero. Then from (16) it follows that $\theta_i = \Theta(\alpha_i^{-1})$ when $\alpha_i \rightarrow +\infty$, so both θ_i and $\alpha_i \theta_i^2$ tend to 0.

4.4. Sparse Variational Dropout for Fully-Connected and Convolutional Layers

Finally we optimize the *stochastic gradient variational lower bound* (3) with our approximation of KL-divergence (14). We apply Sparse Variational Dropout to both convolutional and fully-connected layers. To reduce the variance of \mathcal{L}^{SGVB} we use a combination of the Local Reparameterization Trick and Additive Noise Reparameterization. In order to improve convergence, optimization is performed w.r.t. $(\theta, \log \sigma^2)$.

For a fully connected layer we use the same notation as in Section 3.3. In this case, Sparse Variational Dropout with the Local Reparameterization Trick and Additive Noise Reparameterization can be computed as follows:

$$\begin{aligned} b_{mj} &\sim \mathcal{N}(\gamma_{mj}, \delta_{mj}) \\ \gamma_{mj} &= \sum_{i=1}^I a_{mi} \theta_{ij}, \quad \delta_{mj} = \sum_{i=1}^I a_{mi}^2 \sigma_{ij}^2 \end{aligned} \quad (17)$$

Now consider a convolutional layer. Take a single input tensor $A_m^{H \times W \times C}$, a single filter $w_k^{h \times w \times C}$ and corresponding output matrix $b_{mk}^{H' \times W'}$. This filter has corresponding variational parameters $\theta_k^{h \times w \times C}$ and $\sigma_k^{h \times w \times C}$. Note that in this case A_m , θ_k and σ_k are tensors. Because of linearity of convolutional layers, it is possible to apply the Local Reparameterization Trick. Sparse Variational Dropout for convolutional layers then can be expressed in a way, similar to (17). Here we use $(\cdot)^2$ as an element-wise operation, $*$ denotes the convolution operation, $\text{vec}(\cdot)$ denotes reshaping of a matrix/tensor into a vector.

$$\begin{aligned} \text{vec}(b_{mk}) &\sim \mathcal{N}(\gamma_{mk}, \delta_{mk}) \\ \gamma_{mk} &= \text{vec}(A_m * \theta_k), \quad \delta_{mk} = \text{diag}(\text{vec}(A_m^2 * \sigma_k^2)) \end{aligned} \quad (18)$$

These formulae can be used for the implementation of Sparse Variational Dropout layers. Lasagne and PyTorch

source code of Sparse Variational Dropout layers is available at <https://goo.gl/2D4tFW>. Both forward and backward passes through Sparse VD layers take twice as much time as passes through original layers.

4.5. Relation to RVM

The Relevance Vector Machine (RVM, (Tipping, 2001)) is a classical example of a Sparse Bayesian model. The RVM is essentially a Bayesian treatment of L_2 -regularized linear or logistic regression, where each weight has a separate regularization parameter α_i . These parameters are tuned by empirical Bayes. During training, a large portion of parameters α_i goes to infinity, and corresponding features are excluded from the model since those weights become zero. This effect is known as Automatic Relevance Determination (ARD) and is a popular way to construct sparse Bayesian models.

Empirical Bayes is a somewhat counter-intuitive procedure since we optimize prior distribution w.r.t. the observed data. Such trick has a risk of overfitting, and indeed it was reported in (Cawley, 2010). However, in our work the ARD-effect is achieved by straightforward variational inference rather than by empirical Bayes. Similarly to the RVM, in Sparse VD dropout rates α_i are responsible for the ARD-effect. However, in Sparse VD α_i are parameters of the approximate posterior distribution rather than parameters of the prior distribution. In our work, the prior distribution is fixed and does not have any parameters, and we tune α_i to obtain a more accurate approximation of the posterior distribution $p(w | \mathcal{D})$. Therefore there is no risk of additional overfitting from model selection unlike the case of empirical Bayes.

That said, despite this difference, the analytical solution for maximum a posteriori estimation is very similar for the RVM-regression

$$w^{MAP} = (X^T X + \text{diag}(\alpha))^{-1} X^T y \quad (19)$$

and Sparse Variational Dropout regression

$$\theta = (X^T X + \text{diag}(X^T X) \text{diag}(\alpha))^{-1} X^T y \quad (20)$$

Interestingly, the expression for Binary Dropout-regularized linear regression is exactly the same as (20) if we substitute α_i with $\frac{p_i}{1-p_i}$ (Srivastava et al., 2014).

5. Experiments

We perform experiments on classification tasks and use different neural network architectures including architectures with a combination of batch normalization and dropout layers. We explore the relevance determination performance of our algorithm as well as the classification accuracy of the

resulting sparse model. Our experiments show that Sparse Variational Dropout leads to extremely sparse models.

In order to make a Sparse Variational Dropout analog to an existing architecture, we only need to remove existing dropout layers and replace all dense and convolutional layers with their Sparse Variational Dropout counterparts as described in Section 4.4 and use \mathcal{L}^{SGVB} as the objective function. The value of the variational lower bound can be used to choose among several local optima.

5.1. General Empirical Observations

We provide a general intuition about training of Sparse Bayesian DNNs using Sparse Variational Dropout.

As it is impossible for the weights to converge exactly to zero in a stochastic setting, we explicitly put weights with high corresponding dropout rates to 0 during testing. In our experiments with neural networks, we use the value $\log \alpha = 3$ as a threshold. This value corresponds to a Binary Dropout rate $p > 0.95$. Unlike most other methods (Han et al., 2015b; Wen et al., 2016), this trick usually does not hurt the performance of our model. It means that Sparse VD does not require finetuning after thresholding.

Training our model from a random initialization is troublesome, as a lot of weights become pruned away early during training, before they could possibly learn a useful representation of the data. In this case we obtain a higher sparsity level, but also a high accuracy drop. The same problem is reported by (Sønderby et al., 2016) and is a common problem for Bayesian DNNs. One way to resolve this problem is to start from a pre-trained network. This trick provides a fair sparsity level with almost no drop of accuracy. Here by pre-training we mean training of the original architecture without Sparse Variational Dropout until full convergence.

Another way to approach this problem is to use warm-up, as described by (Sønderby et al., 2016). The idea is to rescale the KL-divergence term during training by a scalar term β_t , individual for each training epoch. During the first epochs we used $\beta_t = 0$, then increased β_t linearly from 0 to 1 and after that used $\beta_t = 1$. The final objective function remains the same, but the optimization trajectory becomes different. In some sense it is equivalent to choosing a better initial guess for the parameters.

We use the final value of the variational lower bound to choose the initialization strategy. We observe that the initialization does not matter much on simple models like LeNets, but in the case of more complex models like VGG, the difference is significant.

On most architectures we observe that the number of epochs required for convergence from a random initialization is roughly the same as for the original network. How-

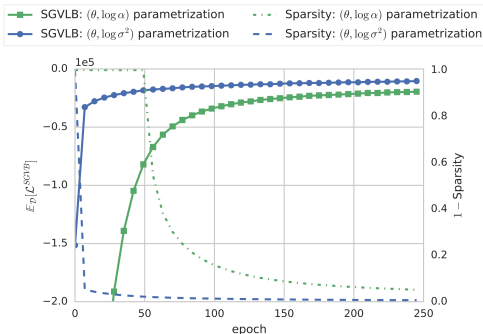


Figure 2: Original parameterization vs Additive Noise Reparameterization. Additive Noise Reparameterization leads to a much faster convergence, a better value of the variational lower bound and a higher sparsity level.

ever, we only need to make a several epochs (10-30) in order for our method to converge from a pre-trained network.

We train all networks using Adam (Kingma & Ba, 2014). When we start from a random initialization, we train for 200 epochs and linearly decay the learning rate from 10^{-4} to zero. When we start from a pre-trained model, we fine-tune for 10-30 epochs with learning rate 10^{-5} .

5.2. Variance Reduction

To see how Additive Noise Reparameterization reduces the variance, we compare it with the original parameterization. We used a fully-connected architecture with two layers with 1000 neurons each. Both models were trained with identical random initializations and with the same learning rate, equal to 10^{-4} . We did not rescale the KL term during training. It is interesting that the original version of Variational Dropout with our approximation of KL-divergence and with no restriction on alphas also provides a sparse solution. However, our method has much better convergence rate and provides higher sparsity and a better value of the variational lower bound.

5.3. LeNet-300-100 and LeNet5 on MNIST

We compare our method with other methods of training sparse neural networks on the MNIST dataset using a fully-connected architecture LeNet-300-100 and a convolutional architecture LeNet-5-Caffe¹. These networks were trained from a random initialization and without data augmentation. We consider pruning (Han et al., 2015b;a), Dynamic Network Surgery (Guo et al., 2016) and Soft Weight Sharing (Ullrich et al., 2017). In these architectures, our method achieves a state-of-the-art level of sparsity, while its accuracy is comparable to other methods. It should be noted

¹A modified version of LeNet5 from (LeCun et al., 1998). Caffe Model specification: <https://goo.gl/4yI3dL>

Network	Method	Error %	Sparsity per Layer %	$\frac{ W }{ W_{\neq 0} }$
LeNet-300-100	Original	1.64		1
	Pruning	1.59	92.0 – 91.0 – 74.0	12
	DNS	1.99	98.2 – 98.2 – 94.5	56
	SWS	1.94		23
	(ours) Sparse VD	1.92	98.9 – 97.2 – 62.0	68
LeNet-5-Caffe	Original	0.80		1
	Pruning	0.77	34 – 88 – 92.0 – 81	12
	DNS	0.91	86 – 97 – 99.3 – 96	111
	SWS	0.97		200
	(ours) Sparse VD	0.75	67 – 98 – 99.8 – 95	280

Table 1: Comparison of different sparsity-inducing techniques (Pruning (Han et al., 2015b;a), DNS (Guo et al., 2016), SWS (Ullrich et al., 2017)) on LeNet architectures. Our method provides the highest level of sparsity with a similar accuracy.

that we only consider the level of sparsity and not the final compression ratio.

5.4. VGG-like on CIFAR-10 and CIFAR-100

To demonstrate that our method scales to large modern architectures, we apply it to a VGG-like network (Zagoruyko, 2015) adapted for the CIFAR-10 (Krizhevsky & Hinton, 2009) dataset. The network consists of 13 convolutional and two fully-connected layers, each layer followed by pre-activation batch normalization and Binary Dropout. We experiment with different sizes of this architecture by scaling the number of units in each network by $k \in \{0.25, 0.5, 1.0, 1.5\}$. We use CIFAR-10 and CIFAR-100 for evaluation. The reported error of this architecture on the CIFAR-10 dataset with $k = 1$ is 7.55%. As no pre-trained weights are available, we train our own network and achieve 7.3% error. Sparse VD also achieves 7.3% error for $k = 1$, but retains $48\times$ less weights.

We observe underfitting while training our model from a random initialization, so we pre-train the network with Binary Dropout and L_2 regularization. It should be noted that most modern DNN compression techniques also can be applied only to pre-trained networks and work best with networks, trained with L_2 regularization (Han et al., 2015b).

Our method achieves over 65x sparsification on the CIFAR-10 dataset with no accuracy drop and up to 41x sparsification on CIFAR-100 with a moderate accuracy drop.

5.5. Random Labels

Recently it was shown that the CNNs are capable of memorizing the data even with random labeling (Zhang et al., 2016). The standard dropout as well as other regularization techniques did not prevent this behaviour. Following that work, we also experiment with the random labeling of data. We use a fully-connected network on the MNIST

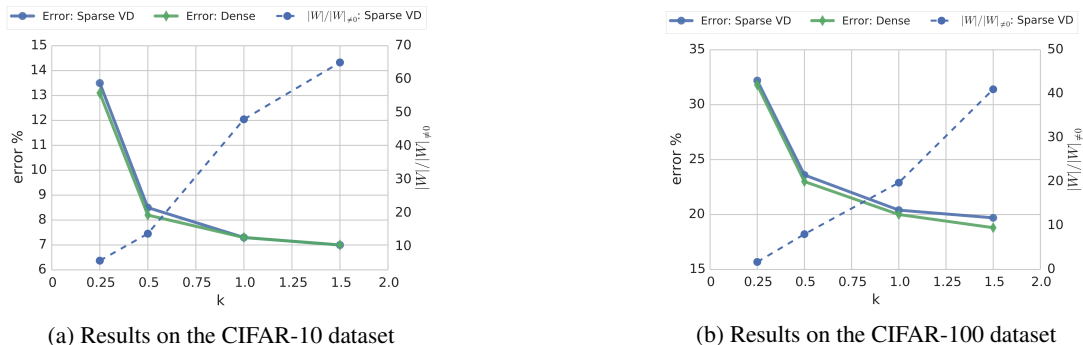


Figure 3: Accuracy and sparsity level for VGG-like architectures of different sizes. The number of neurons and filters scales as k . Dense networks were trained with Binary Dropout, and Sparse VD networks were trained with Sparse Variational Dropout on all layers. The overall sparsity level, achieved by our method, is reported as a dashed line. The accuracy drop is negligible in most cases, and the sparsity level is high, especially in larger networks.

dataset and VGG-like networks on CIFAR-10. We put Binary Dropout (BD) with dropout rate $p = 0.5$ on all fully-connected layers of these networks. We observe that these architectures can fit a random labeling even with Binary Dropout. However, our model decides to drop every single weight and provide a constant prediction. It is still possible to make our model learn random labeling by initializing it with a network, pre-trained on this random labeling, and then finetuning it. However, the variational lower bound $\mathcal{L}(\theta, \alpha)$ in this case is lower than in the case of 100% sparsity. These observations may mean that Sparse VD implicitly penalizes memorization and favors generalization. However, this still requires a more thorough investigation.

6. Discussion

The ‘‘Occam’s razor’’ principle states that unnecessarily complex should not be preferred to simpler ones (MacKay, 1992). Automatic Relevance Determination is effectively a Bayesian implementation of this principle that occurs in different cases. Previously, it was mostly studied in the case of factorized Gaussian prior in linear models, Gaussian Processes, etc. In the Relevance Tagging Machine model (Molchanov et al., 2015) the same effect was achieved using Beta distributions as a prior. Finally, in this work, the ARD-effect is reproduced in an entirely different setting. We consider a fixed prior and train the model using variational inference. In this case, the ARD effect is caused by the particular combination of the approximate posterior distribution family and prior distribution, and not by model selection. This way we can abandon the empirical Bayes approach that is known to overfit (Cawley, 2010).

We observed that if we allow Variational Dropout to drop irrelevant weights automatically, it ends up cutting most of the model weights. This result correlates with results of other works on training of sparse neural networks (Han et al., 2015a; Wen et al., 2016; Ullrich et al., 2017; So-

ravit Changpinyo, 2017). All these works can be viewed as a kind of regularization of neural networks, as they restrict the model complexity. Further investigation of such redundancy may lead to an understanding of generalization properties of DNNs and explain the phenomenon, observed by (Zhang et al., 2016). According to that paper, although modern DNNs generalize well in practice, they can also easily learn a random labeling of data. Interestingly, it is not the case for our model, as a network with zero weights has a higher value of objective than a trained network.

In this paper we study only the level of sparsity and do not report the actual network compression. However, our approach can be combined with other modern techniques of network compression, e.g. quantization and Huffman coding (Han et al., 2015a; Ullrich et al., 2017), as they use sparsification as an intermediate step. As our method provides a higher level of sparsity, we believe that it can improve these techniques even further. Another possible direction for future research is to find a way to obtain structured sparsity using our framework. As reported by (Wen et al., 2016), structured sparsity is crucial to acceleration of DNNs.

Acknowledgements

We would like to thank Michael Figurnov, Ekaterina Lobacheva and Max Welling for valuable feedback. Dmitry Molchanov was supported by the Ministry of Education and Science of the Russian Federation (grant 14.756.31.0001), Arsenii Ashukha was supported by HSE International lab of Deep Learning and Bayesian Methods which is funded by the Russian Academic Excellence Project ‘5-100’, Dmitry Vetrov was supported by the Russian Science Foundation grant 17-11-01027. We would also like to thank the Department of Algorithms and Theory of Programming, Faculty of Innovation and High Technology in Moscow Institute of Physics and Technology for the provided computational resources.

References

- Cawley, Nicola L. C. Talbot. On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11(Jul): 2079–2107, 2010.
- Challis, E and Barber, D. Gaussian kullback-leibler approximate inference. *Journal of Machine Learning Research*, 14:2239–2286, 2013.
- Gal, Yarin and Ghahramani, Zoubin. Dropout as a bayesian approximation: Insights and applications. In *Deep Learning Workshop, ICML*, 2015.
- Garipov, Timur, Podoprikin, Dmitry, Novikov, Alexander, and Vetrov, Dmitry. Ultimate tensorization: compressing convolutional and fc layers alike. *arXiv preprint arXiv:1611.03214*, 2016.
- Guo, Yiwen, Yao, Anbang, and Chen, Yurong. Dynamic network surgery for efficient dnns. In *Advances In Neural Information Processing Systems*, pp. 1379–1387, 2016.
- Han, Song, Mao, Huizi, and Dally, William J. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015a.
- Han, Song, Pool, Jeff, Tran, John, and Dally, William. Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems*, pp. 1135–1143, 2015b.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
- Hinton, Geoffrey E, Srivastava, Nitish, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan R. Improving neural networks by preventing co-adaptation of feature detectors. Technical report, 2012.
- Hoffman, Matthew D, Blei, David M, Wang, Chong, and Paisley, John William. Stochastic variational inference. *Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- Ioffe, Sergey and Szegedy, Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Kingma, Diederik and Ba, Jimmy. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, Diederik P and Welling, Max. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kingma, Diederik P, Salimans, Tim, and Welling, Max. Variational dropout and the local reparameterization trick. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 28*, pp. 2575–2583. Curran Associates, Inc., 2015.
- Krizhevsky, Alex and Hinton, Geoffrey. Learning multiple layers of features from tiny images. 2009.
- Lebedev, Vadim and Lempitsky, Victor. Fast convnets using group-wise brain damage. *arXiv preprint arXiv:1506.02515*, 2015.
- LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Liu, Baoyuan, Wang, Min, Foroosh, Hassan, Tappen, Marshall, and Pinsky, Marianna. Sparse convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 806–814, 2015.
- Louizos, Christos. Smart regularization of deep architectures. 2015.
- MacKay, David JC. Bayesian interpolation. *Neural computation*, 4(3):415–447, 1992.
- MacKay, David JC et al. Bayesian nonlinear modeling for the prediction competition. *ASHRAE transactions*, 100(2):1053–1062, 1994.
- Maeda, Shin-ichi. A bayesian encourages dropout. *arXiv preprint arXiv:1412.7003*, 2014.
- Molchanov, Dmitry, Kondrashkin, Dmitry, and Vetrov, Dmitry. Relevance tagging machine. *Machine Learning and Data Analysis*, 1(13):1877–1887, 2015.
- Neal, Radford M. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 1996.
- Novikov, Alexander, Podoprikin, Dmitrii, Osokin, Anton, and Vetrov, Dmitry P. Tensorizing neural networks. In *Advances in Neural Information Processing Systems*, pp. 442–450, 2015.
- Rezende, Danilo Jimenez, Mohamed, Shakir, and Wierstra, Daan. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Scardapane, Simone, Comminiello, Danilo, Hussain, Amir, and Uncini, Aurelio. Group sparse regularization for deep neural networks. *arXiv preprint arXiv:1607.00485*, 2016.

- Silver, David, Huang, Aja, Maddison, Chris J, Guez, Arthur, Sifre, Laurent, Van Den Driessche, George, Schrittwieser, Julian, Antonoglou, Ioannis, Panneershelvam, Veda, Lanctot, Marc, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- Sønderby, Casper Kaae, Raiko, Tapani, Maaløe, Lars, Sønderby, Søren Kaae, and Winther, Ole. *How to Train Deep Variational Autoencoders and Probabilistic Ladder Networks*. 2016.
- Soravit Changpinyo, Mark Sandler, Andrey Zhmoginov. The power of sparsity in convolutional neural networks. In *Under review on ICLR 2017*, 2017.
- Srinivas, Suraj and Babu, R Venkatesh. Generalized dropout. *arXiv preprint arXiv:1611.06791*, 2016.
- Srivastava, Nitish, Hinton, Geoffrey E, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1): 1929–1958, 2014.
- Szegedy, Christian, Liu, Wei, Jia, Yangqing, Sermanet, Pierre, Reed, Scott, Anguelov, Dragomir, Erhan, Dumitru, Vanhoucke, Vincent, and Rabinovich, Andrew. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–9, 2015.
- Szegedy, Christian, Ioffe, Sergey, Vanhoucke, Vincent, and Alemi, Alex. Inception-v4, inception-resnet and the impact of residual connections on learning. *arXiv preprint arXiv:1602.07261*, 2016.
- Tipping, Michael E. Sparse bayesian learning and the relevance vector machine. *Journal of machine learning research*, 1(Jun):211–244, 2001.
- Titsias, Michalis and Lázaro-Gredilla, Miguel. Doubly stochastic variational bayes for non-conjugate inference. *Proceedings of The 31st International Conference on Machine Learning*, 32:1971–1979, 2014.
- Ullrich, Karen, Meeds, Edward, and Welling, Max. Soft weight-sharing for neural network compression. *arXiv preprint arXiv:1702.04008*, 2017.
- Van Gestel, Tony, Suykens, JAK, De Moor, Bart, and Vandewalle, Joos. Automatic relevance determination for least squares support vector machine regression. In *Neural Networks, 2001. Proceedings. IJCNN'01. International Joint Conference on*, volume 4, pp. 2416–2421. IEEE, 2001.
- Wan, Li, Zeiler, Matthew, Zhang, Sixin, Cun, Yann L, and Fergus, Rob. Regularization of neural networks using dropconnect. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pp. 1058–1066, 2013.
- Wang, Sida I and Manning, Christopher D. Fast dropout training. In *ICML (2)*, pp. 118–126, 2013.
- Wen, Wei, Wu, Chunpeng, Wang, Yandan, Chen, Yiran, and Li, Hai. Learning structured sparsity in deep neural networks. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 2074–2082. Curran Associates, Inc., 2016.
- Zagoruyko, Sergey. 92.45 on cifar-10 in torch, 2015. URL <http://torch.ch/blog/2015/07/30/cifar.html>.
- Zhang, Chiyuan, Bengio, Samy, Hardt, Moritz, Recht, Benjamin, and Vinyals, Oriol. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.