CONTEXTUAL DROPOUT: AN EFFICIENT SAMPLE-DEPENDENT DROPOUT MODULE

Anonymous authors
Paper under double-blind review

ABSTRACT

Dropout has been demonstrated as a simple and effective module to not only regularize the training process of deep neural networks, but also provide the uncertainty estimation for prediction. However, the quality of uncertainty estimation is highly dependent on the dropout probabilities. Most current models use the same dropout distributions across all data samples due to its simplicity. Despite the potential gains in the flexibility of modeling uncertainty, sample-dependent dropout, on the other hand, is less explored as it often encounters scalability issues or involves non-trivial model changes. In this paper, we propose contextual dropout with an efficient structural design as a simple and scalable sample-dependent dropout module, which can be applied to a wide range of models at the expense of only slightly increased memory and computational cost. We learn the dropout probabilities with a variational objective, compatible with both Bernoulli dropout and Gaussian dropout. We apply the contextual dropout module to various models with applications to image classification and visual question answering and demonstrate the scalability of the method with large-scale datasets, such as ImageNet and VQA 2.0. Our experimental results show that the proposed method outperforms baseline methods in terms of both accuracy and quality of uncertainty estimation.

1 INTRODUCTION

Deep neural networks (NNs) have become ubiquitous and achieved state-of-the-art results in a wide variety of research problems (LeCun et al., 2015). To prevent over-parameterized NNs from overfitting, we often need to appropriately regularize their training. One way to do so is to use Bayesian NNs that treat the NN weights as random variables and regularize them with appropriate prior distributions (MacKay, 1992; Neal, 2012). More importantly, we can obtain the model's confidence on its predictions by evaluating the consistency between the predictions that are conditioned on different posterior samples of the NN weights. However, despite significant recent efforts in developing various types of approximate inference for Bayesian NNs (Graves, 2011; Welling & Teh, 2011; Li et al., 2016; Blundell et al., 2015; Louizos & Welling, 2017; Shi et al., 2018), the large number of NN weights makes it difficult to scale to real-world applications.

Dropout has been demonstrated as another effective regularization strategy, which can be viewed as imposing a distribution over the NN weights (Gal & Ghahramani, 2016). Relating dropout to Bayesian inference provides a much simpler and more efficient way than using vanilla Bayesian NNs to provide uncertainty estimation (Gal & Ghahramani, 2016), as there is no more need to explicitly instantiate multiple sets of NN weights. For example, Bernoulli dropout randomly shuts down neurons during training (Hinton et al., 2012; Srivastava et al., 2014). Gaussian dropout multiplies the neurons with independent, and identically distributed (iid) Gaussian random variables drawn from $\mathcal{N}(1, \alpha)$, where the variance α is a tuning parameter (Srivastava et al., 2014). Variational dropout generalizes Gaussian dropout by reformulating it under a Bayesian setting and allowing α to be learned under a variational objective (Kingma et al., 2015; Molchanov et al., 2017).

However, the quality of uncertainty estimation depends heavily on the dropout probabilities (Gal et al., 2017). To avoid grid-search over the dropout probabilities, Gal et al. (2017) and Boluki et al. (2020) propose to automatically learn the dropout probabilities, which not only leads to a faster experiment cycle but also enables the model to have different dropout probabilities for each layer, bringing greater flexibility into the uncertainty modeling. But, these methods still impose the restrictive assumption

that dropout probabilities are global parameters shared across all data samples. By contrast, we consider parameterizing dropout probabilities as a function of input covariates, treating them as data-dependent local variables. Applying covariate-dependent dropouts allows different data to have different distributions over the NN weights. This generalization has the potential to greatly enhance the expressiveness of a Bayesian NN. However, learning covariate-dependent dropout rates is challenging. Ba & Frey (2013) propose *standout*, where a binary belief network is laid over the original network, and develop a heuristic approximation to optimize free energy. But, as pointed out by Gal et al. (2017), it is not scalable due to its need to significantly increase the model size.

In this paper, we propose a simple and scalable contextual dropout module, whose dropout rates depend on the covariates $\boldsymbol{x}$, as a new approximate Bayesian inference method for NNs. With a novel design that reuses the main network to define how the covariate-dependent dropout rates are produced, it boosts the performance while only slightly increases the memory and computational cost. Our method greatly enhances the flexibility of modeling, maintains the inherent advantages of dropout over conventional Bayesian NNs, and is generally simple to implement and scalable to the large-scale applications. We plug the contextual dropout module into various types of NN layers, including fully connected, convolutional, and attention layers. On a variety of supervised learning tasks, contextual dropout achieves good performance in terms of accuracy and quality of uncertainty estimation.

2 CONTEXTUAL DROPOUT

We introduce an efficient solution for data-dependent dropout: (1) treat the dropout probabilities as sample-dependent local random variables, (2) propose an efficient parameterization of dropout probabilities by sharing parameters between the encoder and decoder, and (3) learn the dropout distribution with a variational objective.

2.1 BACKGROUND ON DROPOUT MODULES

Consider a supervised learning problem with training data $\mathcal{D} := \{\boldsymbol{x}_i, y_i\}_{i=1}^N$, where we model the conditional probability $p_{\boldsymbol{\theta}}(y_i \,|\, \boldsymbol{x}_i)$ using a NN parameterized by $\boldsymbol{\theta}$. Applying dropout to a NN often means element-wisely reweighing each layer with a data-specific Bernoulli/Gaussian distributed random mask $\boldsymbol{z}_i$, which are iid drawn from a prior $p_{\boldsymbol{\eta}}(\boldsymbol{z})$ parameterized by $\boldsymbol{\eta}$ (Hinton et al., 2012; Srivastava et al., 2014). This implies dropout training can be viewed as approximate Bayesian inference (Gal & Ghahramani, 2016). More specifically, one may view the learning objective of a supervised learning model with dropout as a log-marginal-likelihood: $\log \int \prod_{i=1}^N p(y_i \,|\, \boldsymbol{x}_i, \boldsymbol{z}) p(\boldsymbol{z}) d\boldsymbol{z}$. To maximize this often intractable log-marginal, it is common to resort to variational inference (Hoffman et al., 2013; Blei et al., 2017) that introduces a variational distribution $q(\boldsymbol{z})$ on the random mask $\boldsymbol{z}$ and optimizes an evidence lower bound (ELBO):

$$\mathcal{L}(\mathcal{D}) = \mathbb{E}_{q(\boldsymbol{z})} \left[\log \frac{\prod_{i=1}^N p_{\boldsymbol{\theta}}(y_i \,|\, \boldsymbol{x}_i, \boldsymbol{z}) p_{\boldsymbol{\eta}}(\boldsymbol{z})}{q(\boldsymbol{z})} \right] = \big(\textstyle\sum_{i=1}^N \mathbb{E}_{\boldsymbol{z}_i \sim q(\boldsymbol{z})} \left[\log p_{\boldsymbol{\theta}}(y_i \,|\, \boldsymbol{x}_i, \boldsymbol{z}_i) \right] \big) - \mathrm{KL}(q(\boldsymbol{z}) || p_{\eta}(\boldsymbol{z})), \quad (1)$$

where $\mathrm{KL}(q(\boldsymbol{z}) || p_{\boldsymbol{\eta}}(\boldsymbol{z})) = \mathbb{E}_{q(\boldsymbol{z})}[\log q(\boldsymbol{z}) - \log p(\boldsymbol{z})]$ is a Kullback-Leibler (KL) divergence based regularization term. Whether the KL term is explicitly imposed is a key distinction between regular dropout (Hinton et al., 2012; Srivastava et al., 2014) and their Bayesian generalizations (Gal & Ghahramani, 2016; Gal et al., 2017; Kingma et al., 2015; Molchanov et al., 2017; Boluki et al., 2020).

2.2 COVARIATE-DEPENDENT WEIGHT UNCERTAINTY

In regular dropout, as shown in (1), while we make the dropout masks data specific during optimization, we keep their distributions the same. This implies that while the NN weights can vary from data to data, their distribution is kept data invariant. In this paper, we propose *contextual dropout*, in which the distributions of dropout masks $\boldsymbol{z}_i$ depend on covariates $\boldsymbol{x}_i$ for each sample $(\boldsymbol{x}_i, y_i)$. Specifically, we define the variational distribution as $q_{\boldsymbol{\phi}}(\boldsymbol{z}_i \,|\, \boldsymbol{x}_i)$, where $\boldsymbol{\phi}$ denotes its NN parameters. In the framework of amortized variational Bayes (Kingma & Welling, 2013; Rezende et al., 2014), we can view $q_{\boldsymbol{\phi}}$ as an inference network (encoder) trying to approximate the posterior $p(\boldsymbol{z}_i \,|\, y_i, \boldsymbol{x}_i) \propto p(y_i \,|\, \boldsymbol{x}_i, \boldsymbol{z}_i) p(\boldsymbol{z}_i)$. Note as we have no access to y_i during testing, we parameterize our encoder in a way that it depends on $\boldsymbol{x}_i$ but not y_i. From the optimization point of view, what we propose corresponds to the ELBO of $\log \prod_{i=1}^N \int p(y_i \,|\, \boldsymbol{x}_i, \boldsymbol{z}_i) p(\boldsymbol{z}_i) d\boldsymbol{z}_i$ given $q_{\boldsymbol{\phi}}(\boldsymbol{z}_i \,|\, \boldsymbol{x}_i)$ as the

encoder, which can be expressed as

$$\mathcal{L}(\mathcal{D}) = \sum_{i=1}^{N} \mathcal{L}(\boldsymbol{x}_i, y_i),\ \mathcal{L}(\boldsymbol{x}_i, y_i) = \mathbb{E}_{\boldsymbol{z}_i \sim q_{\boldsymbol{\phi}}(\cdot \mid \boldsymbol{x}_i)}[\log p_{\boldsymbol{\theta}}(y_i \mid \boldsymbol{x}_i, \boldsymbol{z}_i)] - \mathrm{KL}(q_{\boldsymbol{\phi}}(\boldsymbol{z}_i \mid \boldsymbol{x}_i) || p_{\boldsymbol{\eta}}(\boldsymbol{z}_i)). \quad (2)$$

This ELBO differs from that of regular dropout in (1) in that the dropout distributions for $\boldsymbol{z}_i$ are now parameterized by $\boldsymbol{x}_i$ and a single KL regularization term is replaced with the aggregation of N data-dependent KL terms. Unlike conventional Bayesian NNs, as $\boldsymbol{z}_i$ is now a local random variable, the impact of the KL terms will not diminish as N increases, and from the viewpoint of uncertainty quantification, contextual dropout relies only on aleatoric uncertainty to model its uncertainty on y_i given $\boldsymbol{x}_i$. Like conventional BNNs, we may add epistemic uncertainty by imposing a prior distribution on $\boldsymbol{\theta}$ and/or $\boldsymbol{\phi}$, and infer their posterior given $\mathcal{D}$. As contextual dropout with a point estimate on both $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ is already achieving state-of-the-art performance, we leave that extension for future research. In what follows, we omit the data index i for simplification and formally define its model structure.

Cross-layer dependence: For a NN with L layers, we denote $\boldsymbol{z} = \{\boldsymbol{z}^1, \ldots, \boldsymbol{z}^L\}$, with $\boldsymbol{z}^l$ representing the dropout masks at layer l. As we expect $\boldsymbol{z}^l$ to be dependent on the dropout masks in previous layers $\{\boldsymbol{z}^j\}_{j<l}$, we introduce an autoregressive distribution as $q_{\boldsymbol{\phi}}(\boldsymbol{z} \mid \boldsymbol{x}) = \prod_{l=1}^{L} q_{\boldsymbol{\phi}}(\boldsymbol{z}^l \mid \boldsymbol{x}^{l-1})$, where $\boldsymbol{x}^{l-1}$, the output of layer $l-1$, is a function of $\{\boldsymbol{z}^1, \ldots, \boldsymbol{z}^{l-1}, \boldsymbol{x}\}$.

Parameter sharing between encoder and decoder: We aim to build an encoder by modeling $q_{\boldsymbol{\phi}}(\boldsymbol{z}^l \mid \boldsymbol{x}^{l-1})$, where $\boldsymbol{x}$ may come from complex and highly structured data such as images and natural languages. Thus, extracting useful features from $\boldsymbol{x}$ to learn the encoder distribution $q_{\boldsymbol{\phi}}$ itself becomes a problem as challenging as the original one, *i.e.*, extracting discriminative features from $\boldsymbol{x}$ to predict y. As intermediate layers in the decoder network $p_{\boldsymbol{\theta}}$ are already learning useful features from the input, we choose to reuse them in the encoder, instead of extracting the features from scratch. If we denote layer l of the decoder network by $g^l_{\boldsymbol{\theta}}$, then the output of layer l, given its input $\boldsymbol{x}^{l-1}$, would be $\mathbf{U}^l = g^l_{\boldsymbol{\theta}}(\boldsymbol{x}^{l-1})$. Considering this as a learned feature for $\boldsymbol{x}$, as illustrated in Figure 1, we build the encoder on this output as $\boldsymbol{\alpha}^l = h^l_{\boldsymbol{\varphi}}(\mathbf{U}^l)$, draw $\boldsymbol{z}^l$ conditioning on $\boldsymbol{\alpha}^l$, and element-wisely multiply $\boldsymbol{z}^l$ with $\mathbf{U}^l$ (with broadcast if needed) to produce the output of layer l as $\boldsymbol{x}^l$. In this way, we use $\{\boldsymbol{\theta}, \boldsymbol{\varphi}\}$ to parameterize the encoder, which reuses parameters $\boldsymbol{\theta}$ of the decoder. To produce the dropout rates of the encoder, we only need extra parameters $\boldsymbol{\varphi}$, the added memory and computational cost of which are often insignificant in comparison to these of the decoder.

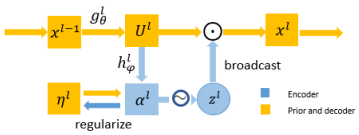

Figure 1: A contextual dropout module.

2.3 Efficient parameterization of contextual dropout module

Denote the output of layer l by a multidimensional array (tensor) $\mathbf{U}^l = g^l_{\boldsymbol{\theta}}(\boldsymbol{x}^{l-1}) \in \mathbb{R}^{C^l_1 \times \ldots \times C^l_{D^l}}$, where D^l denotes the summation of the dimensions of $\mathbf{U}^l$ and C^l_d denotes the number of elements along dimension $d \in \{1, \ldots, D^l\}$. For efficiency, the output shape of $h^l_{\boldsymbol{\varphi}}$ is not matched to the shape of $\mathbf{U}^l$. Instead, we make it smaller and broadcast the contextual dropout masks $\boldsymbol{z}^l$ across the dimensions of $\mathbf{U}^l$ (Tompson et al., 2015). Specifically, we parameterize dropout logits $\boldsymbol{\alpha}^l$ of the variational distribution to have C^l_d elements, where $d \in \{1, ..., D^l\}$ is a specified dimension of $\mathbf{U}^l$. We sample $\boldsymbol{z}^l$ from the encoder and broadcast them across all but dimension d of $\mathbf{U}^l$. We sample $\boldsymbol{z}^l \sim \mathrm{Ber}(\sigma(\boldsymbol{\alpha}^l))$ under contextual Bernoulli dropout, and follow Srivastava et al. (2014) to use $\boldsymbol{z}^l \sim N(1, \sigma(\boldsymbol{\alpha}^l)/(1-\sigma(\boldsymbol{\alpha}^l)))$ for contextual Gaussian dropout. To obtain $\boldsymbol{\alpha}^l \in \mathbb{R}^{C^l_d}$, we first take the average pooling of $\mathbf{U}^l$ across all but dimension d, with the output denoted as $F_{\text{avepool},d}(\mathbf{U}^l)$, and then apply two fully-connected layers Φ^l_1 and Φ^l_2 connected by F_{NL}, a (Leaky) ReLU based nonlinear activation function, as

$$\boldsymbol{\alpha}^l = h^l_{\boldsymbol{\varphi}}(\mathbf{U}^l) = \Phi^l_2(F_{\text{NL}}(\Phi^l_1(F_{\text{avepool},d}(\mathbf{U}^l)))), \quad (3)$$

where Φ^l_1 is a linear transformation mapping from $\mathbb{R}^{C^l_d}$ to $\mathbb{R}^{C^l_d/\gamma}$, while Φ^l_2 is from $\mathbb{R}^{C^l_d/\gamma}$ back to $\mathbb{R}^{C^l_d}$, with γ being a reduction ratio controlling the complexity of $h^l_{\boldsymbol{\varphi}}$. Below we describe how to apply contextual dropout to three representative types of NN layers.

Contextual dropout module for fully-connected layers: If layer l is a fully-connected layer and $\mathbf{U}^l \in \mathbb{R}^{C^l_1 \times \cdots \times C^l_{D^l}}$, we set $\boldsymbol{\alpha}^l \in \mathbb{R}^{C^l_{D^l}}$, where D^l is the dimension that the linear transformation is applied to. Note, if $\mathbf{U}^l \in \mathbb{R}^{C^l_1}$, then $\boldsymbol{\alpha}^l \in \mathbb{R}^{C^l_1}$, and $F_{\text{avepool},1}$ is an identity map, so $\boldsymbol{\alpha}^l = \Phi^l_2 F_{\text{NL}}(\Phi^l_1(\mathbf{U}^l))$.

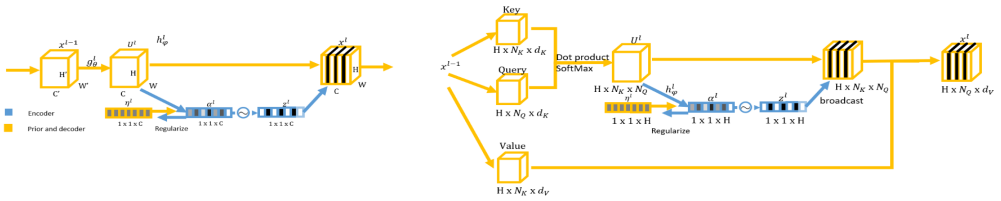

Figure 2: Left: Contextual dropout in convolution layers. Right: Contextual dropout in attention layers.

Contextual dropout module for convolutional layers: Assume layer l is a convolutional layer with C_3^l as convolutional channels and $\mathbf{U}^l \in \mathbb{R}^{C_1^l \times C_2^l \times C_3^l}$. Similar to Spatial Dropout (Tompson et al., 2015), we set $\boldsymbol{\alpha}^l \in \mathbb{R}^{C_3^l}$ and broadcast its corresponding $\boldsymbol{z}^l$ spatially as illustrated in Figure 2. Such parameterization is similar to the squeeze-and-excitation unit for convolutional layers, which has been shown to effective in image classification tasks (Hu et al., 2018). However, in squeeze-and-excitation, $\sigma(\boldsymbol{\alpha}^l)$ is used as channel-wise soft attention weights instead of dropout probabilities, therefore it serves as a deterministic mapping instead of a stochastic unit used in variational inference.

Contextual dropout module for attention layers: Dropout has been widely used in attention layers (Xu et al., 2015b; Vaswani et al., 2017; Yu et al., 2019). For example, it can be applied to multi-head attention weights after the softmax operation (see illustrations in Figure 2). The weights are of dimension $[H, N_K, N_Q]$, where H is the number of heads, N_K the number of keys, and N_Q the number of queries. In this case, we find that setting $\boldsymbol{\alpha}^l \in \mathbb{R}^H$ gives good performance. Intuitively, this coincides with the choice of channel dimension for convolutional layers, as heads in attention could be analogized as channels in convolution.

2.4 Variational inference for contextual dropout

In contextual dropout, we choose $\mathcal{L}(\mathcal{D}) = \sum_{(\boldsymbol{x},y)\in\mathcal{D}} \mathcal{L}(\boldsymbol{x}, y)$ shown in (2) as the optimization objective. Note in our design, the encoder $q_{\boldsymbol{\phi}}$ reuses the decoder parameters $\boldsymbol{\theta}$ to define its own parameters. Therefore, we copy the values of $\boldsymbol{\theta}$ into $\boldsymbol{\phi}$ and stop the gradient of $\boldsymbol{\theta}$ when optimizing $q_{\boldsymbol{\phi}}$. This is theoretically sound (Ba & Frey, 2013). Intuitively, the gradients to $\boldsymbol{\theta}$ from $p_{\boldsymbol{\theta}}$ are less noisy than that from $q_{\boldsymbol{\phi}}$ as the training of $p_{\boldsymbol{\theta}}(y \mid \boldsymbol{x}, \boldsymbol{z})$ is supervised while that of $q_{\boldsymbol{\phi}}(\boldsymbol{z})$ is unsupervised. As what we expected, allowing gradients from $q_{\boldsymbol{\phi}}$ to backpropagate to $\boldsymbol{\theta}$ is found to adversely affect the training of $p_{\boldsymbol{\theta}}$ in our experiments. We use a simple prior $p_{\boldsymbol{\eta}}$, making the prior distributions for dropout masks the same within each layer. The gradients with respect to $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ can be expressed as

$$\nabla_{\boldsymbol{\eta}}\mathcal{L}(\boldsymbol{x}, y) = \mathbb{E}_{\boldsymbol{z}\sim q_{\boldsymbol{\phi}}(\cdot \mid \boldsymbol{x})}[\nabla_{\boldsymbol{\eta}} \log p_{\boldsymbol{\eta}}(\boldsymbol{z})], \quad \nabla_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{x}, y) = \mathbb{E}_{\boldsymbol{z}\sim q_{\boldsymbol{\phi}}(\cdot \mid \boldsymbol{x})}[\nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(y \mid \boldsymbol{x}, \boldsymbol{z})], \quad (4)$$

which are both estimated via Monte Carlo integration, using a single $\boldsymbol{z} \sim q_{\boldsymbol{\phi}}(\boldsymbol{z} \mid \boldsymbol{x})$ for each $\boldsymbol{x}$.

Now, we consider the gradient of $\mathcal{L}$ with respect to $\boldsymbol{\varphi}$, the components of $\boldsymbol{\phi} = \{\boldsymbol{\theta}, \boldsymbol{\varphi}\}$ not copied from the decoder. For Gaussian contextual dropout, we estimate the gradients via the reparameterization trick (Kingma & Welling, 2013). For $\boldsymbol{z}^l \sim N(\mathbf{1}, \sigma(\boldsymbol{\alpha}^l)/(1-\sigma(\boldsymbol{\alpha}^l)))$, we rewrite it as $\boldsymbol{z}^l = 1 + \sqrt{\sigma(\boldsymbol{\alpha}^l)/(1-\sigma(\boldsymbol{\alpha}^l))}\boldsymbol{\epsilon}^l$, where $\boldsymbol{\epsilon}^l \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Similarly, sampling a sequence of $\boldsymbol{z} = \{\boldsymbol{z}^l\}_{l=1}^L$ from $q_{\boldsymbol{\phi}}(\boldsymbol{z} \mid \boldsymbol{x})$ can be rewritten as $f_{\boldsymbol{\phi}}(\boldsymbol{\epsilon}, \boldsymbol{x})$, where $f_{\boldsymbol{\phi}}$ is a deterministic differentiable mapping and $\boldsymbol{\epsilon}$ are iid standard Gaussian. The gradient $\nabla_{\boldsymbol{\varphi}}\mathcal{L}(\boldsymbol{x}, y)$ can now be expressed as (see pseudo code of Algorithm 3 in Appendix)

$$\nabla_{\boldsymbol{\varphi}}\mathcal{L}(\boldsymbol{x}, y) = \mathbb{E}_{\boldsymbol{\epsilon}\sim\mathcal{N}(\mathbf{0},\mathbf{1})}[\nabla_{\boldsymbol{\varphi}}(\log p_{\boldsymbol{\theta}}(y \mid \boldsymbol{x}, f_{\boldsymbol{\phi}}(\boldsymbol{\epsilon}, \boldsymbol{x})) - \tfrac{\log q_{\boldsymbol{\phi}}(f_{\boldsymbol{\phi}}(\boldsymbol{\epsilon},\boldsymbol{x}) \mid \boldsymbol{x})}{\log p_{\boldsymbol{\eta}}(f_{\boldsymbol{\phi}}(\boldsymbol{\epsilon},\boldsymbol{x}))})]. \quad (5)$$

For Bernoulli contextual dropout, backpropagating the gradient efficiently is not straightforward, as the Bernoulli distribution is not reparameterizable, restricting the use of the reparameterization trick. In this case, a commonly used gradient estimator is the REINFORCE estimator (Williams, 1992) (see details in Appendix A). This estimator, however, is known to have high Monte Carlo estimation variance. To this end, we estimate $\nabla_{\boldsymbol{\varphi}}\mathcal{L}$ with the augment-REINFORCE-merge (ARM) estimator (Yin & Zhou, 2018), which provides unbiased and low-variance gradients for the parameters of Bernoulli distributions. We defer the details of this estimator to Appendix A.

2.5 TESTING AND COMPLEXITY ANALYSIS

Testing stage: To obtain a point estimate, we follow the common practice in dropout (Srivastava et al., 2014) to multiply the neurons by the expected values of random dropout masks, which means that we predict y with $p_{\boldsymbol{\theta}}(y \mid \boldsymbol{x}, \bar{\boldsymbol{z}})$, where $\bar{\boldsymbol{z}} = \mathbb{E}_{q_{\boldsymbol{\phi}}(\boldsymbol{z} \mid \boldsymbol{x})}[\boldsymbol{z}]$ under the proposed contextual dropout. When uncertainty estimation is needed, we draw K random dropout masks to approximate the posterior predictive distribution of y given $\boldsymbol{x}$ using $\hat{p}(y \mid \boldsymbol{x}) = \frac{1}{K}\sum_{k=1}^{K} p_{\boldsymbol{\theta}}(y \mid \boldsymbol{x}, \boldsymbol{z}^{(k)})$, where $\boldsymbol{z}^{(1)}, \ldots, \boldsymbol{z}^{(K)} \overset{iid}{\sim} q_{\boldsymbol{\phi}}(\boldsymbol{z} \mid \boldsymbol{x})$.

Complexity analysis: The added computation and memory of contextual dropout are insignificant due to the parameter sharing between the encoder and decoder. Extra memory and computational cost mainly comes from the part of $h^l_{\boldsymbol{\varphi}}$, where both the parameter size and the number of operations are of order $O((C^l_d)^2/\gamma)$, where γ is from 8 to 16. This is insignificant, compared to the memory and computational cost of the main network, which are of order larger than $O((C^l_d)^2)$. We verify the point by providing memory and runtime comparisons between contextual dropout and other dropouts on ResNet in Table 3 (see more model size comparisons in Table 5 in Appendix).

2.6 RELATED WORK

Data-dependent variational distribution: Deng et al. (2018) model attentions as latent-alignment variables and optimize a tighter lower bound (compared to hard attention) using a learned inference network. To balance exploration and exploitation for contextual bandits problems, Wang & Zhou (2019) introduce local variable uncertainty under the Thompson sampling framework. However, their inference networks of are both independent of the decoder, which may considerably increase memory and computational cost for the considered applications. In addition, while the scope of Deng et al. (2018) is limited to attention units and that of Wang & Zhou (2019) limited to contextual bandits, we demonstrate the general applicability of contextual dropout to fully connected, convolutional, and attention layers in supervised learning models. *Conditional computation* (Bengio et al., 2015; 2013; Shazeer et al., 2017; Teja Mullapudi et al., 2018) tries to increase model capacity without a proportional increase in computation, where an independent gating network decides turning which part of a network active and which inactive for each example. In contextual dropout, the encoder works much like a gating network choosing the distribution of sub-networks for each sample. But the potential gain in model capacity is even larger, $e.g.$, there are potentially $\sim O((2^d)^L)$ combinations of nodes for L fully-connected layers, where d is the order of the number of nodes for one layer.

3 EXPERIMENTS

Our method can be straightforwardly deployed wherever regular dropout can be utilized. To test its general applicability and scalability, we apply the proposed method to three representative types of NN layers: fully connected, convolutional, and attention layers with applications on MNIST (LeCun et al., 2010), CIFAR (Krizhevsky et al., 2009), ImageNet (Deng et al., 2009), and VQA-v2 (Goyal et al., 2017). To investigate the model's robustness to noise, we also construct noisy versions of datasets by adding Gaussian noises to image inputs (Larochelle et al., 2007).

For evaluation, we consider both the accuracy and uncertainty on predicting y given $\boldsymbol{x}$. Many metrics have been proposed to evaluate the quality of uncertainty estimation. On one hand, researchers are generating calibrated probability estimates to measure model confidence (Guo et al., 2017; Naeini et al., 2015; Kuleshov et al., 2018). While expected calibration error and maximum calibration error have been proposed to quantitatively measure calibration, such metrics do not reflect how robust the probabilities are with noise injected into the network input, and cannot capture epistemic or model uncertainty (Gal & Ghahramani, 2016). On the other hand, the entropy of the predictive distribution as well as the mutual information, between the predictive distribution and posterior over network weights, are used as metrics to capture both epistemic and aleatoric uncertainty (Mukhoti & Gal, 2018). However, it is often unclear how large the entropy or mutual information is large enough to be classified as uncertain, so such metric only provides a relative uncertainty measure.

Hypothesis testing based uncertainty estimation: Unlike previous information theoretic metrics, we use a statistical test based method to estimate uncertainty, which works for both single-label and multi-label classification models. One advantage of using hypothesis testing over information

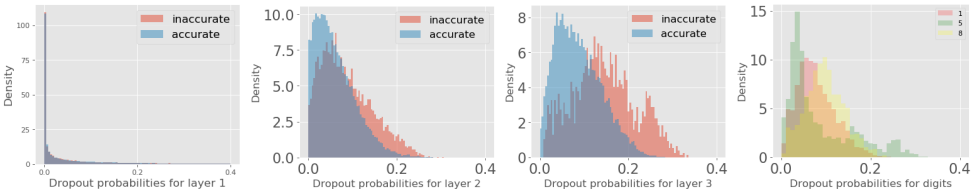

Figure 3: Visualization of dropout probabilities of Bernoulli contextual dropout on the MNIST dataset: the learned dropout probabilities seem to increase as we go to higher-level layers, as also observed in Gal et al. (2017). With contextual dropout, different samples own different dropout probabilities. Inaccurate ones often have higher dropout probabilities corresponding to higher uncertainties. On the further right figure, we compare the dropout distributions across 3 representative digits. The dropout probabilities are overall higher for digit 8 compared to digit 1, meaning 1 is easier to classify. The distribution for 5 has longer tails than others showing there are more variations in the uncertainty for digit 5.

theoretic metrics is that the p-value of the test can be more interpretable, making it easier to be deployed in practice to obtain a binary uncertainty decision. To quantify how confident our model is about this prediction, we evaluate whether the difference between the empirical distributions of the two most possible classes from multiple posterior samples is statistically significant. Please see Appendix D for a detailed explanation of the test procedure.

Uncertainty evaluation via PAvPU: With the p-value of the testing result and a given p-value threshold, we can determine whether the model is certain or uncertain about one prediction. To evaluate the uncertainty estimates, we uses Patch Accuracy vs Patch Uncertainty (PAvPU) (Mukhoti & Gal, 2018), which is defined as $\text{PAvPU} = (n_{ac} + n_{iu})/(n_{ac} + n_{au} + n_{ic} + n_{iu})$, where $n_{ac}, n_{au}, n_{ic}, n_{iu}$ are the numbers of accurate and certain, accurate and uncertain, inaccurate and certain, inaccurate and uncertain samples, respectively. This PAvPU evaluation metric would be higher if the model tends to generate the accurate prediction with high certainty and inaccurate prediction with high uncertainty.

3.1 Contextual dropout on fully connected layers

We consider an MLP with two hidden layers of size 300 and 100, respectively, with ReLU activations. Dropout is applied to the input layer and the outputs of first two full-connected layers. We use MNIST as the benchmark. We compare contextual dropout with MC dropout (Gal & Ghahramani, 2016), concrete dropout (Gal et al., 2017), Gaussian dropout (Srivastava et al., 2014), and Bayes by Backprop (Blundell et al., 2015). Please see the detailed experimental setting in Appendix C.1.

Results and analysis: In Table 1, we show accuracy, PAvPU (p-value threshold equal to 0.05) and test predictive loglikelihood with error bars (5 random runs) for models with different dropouts under the challenging noisy data[1] (added Gaussian noise with mean 0, variance 1). Note that the uncertainty results for p-value threshold 0.05 is in general consistent with the results for other p-value thresholds (see more in Table 6 in Appendix). We observe that contextual dropout outperforms other methods in all metrics. Moreover, compared to Bayes by Backprop, contextual dropout is more memory and computationally efficient. As shown in Table 5 in Appendix, contextual dropout only introduces 16% additional parameters. However, Bayes by Backprop doubles the memory and increases the computations significantly as we need multiple draws of NN weights for uncertainty. Due to this reason, we do not include it for the following large model evaluations. In Figure 3, we observe that Bernoulli contextual dropout learns different dropout probabilities for different samples adapting the sample-level uncertainty which further verifies our motivation and supports the empirical improvements.

Table 1: Results on noisy MNIST with MLP.

Methods	Accuracy	PAvPU(0.05)	log likelihood
MC - Bernoulli	86.36	85.63	-1.72
MC - Gaussian	86.31	85.64	-1.72
Concrete	86.52	86.77	-1.68
Bayes By Backprop	86.55	87.13	-2.30
Bernoulli Contextual	**87.43**±0.39	**87.81**±0.23	**-1.41** ±0.01
Gaussian Contextual	87.35±0.33	87.72±0.29	-1.43±0.01

Combine contextual dropout with Deep Ensemble: Deep ensemble proposed by Lakshminarayanan et al. (2017) is a simple way to obtain uncertainty by with ensembling models trained independently from different random initializations. In Figure 4, we show the performance of combining different dropouts with deep ensemble on noisy MNIST data. As the number of NNs increases, both accuracy and PAvPU increase for all dropouts. However, Bernoulli con-

[1]Results on original data is deferred to Table 6 in Appendix .

textual dropout outperforms other dropouts by a large margin in both metrics, showing contextual dropout is compatible with deep ensemble and their combination can lead to significant improvements. **Out of distribution (OOD) evaluation:** we evaluate different dropouts in an OOD setting, where we train our model with clean data but test it on noisy data. Contextual dropout achieves accuracy of 78.08 consistently higher than MC dropout (75.22) or concrete dropout (74.93). Meanwhile, the proposed method is also better at uncertainty estimation with PAvPU of 78.49 higher than MC (74.61) or Concrete (75.49).

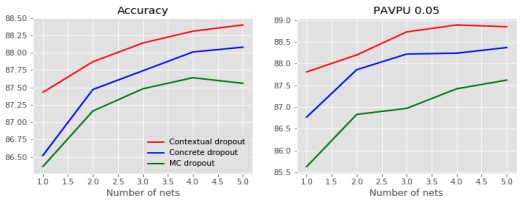

Figure 4: The performance of combining different dropouts with deep ensemble on noisy MNIST data.

3.2 Contextual dropout on convolutional layers

We apply dropout to the convolutional layers in WRN (Zagoruyko & Komodakis, 2016). In Figure 6 in Appendix, we show the architecture of WRN, where dropout is applied to the first convolutional layer in each network block; in total, dropout is applied to 12 convolutional layers. We evaluate on CIFAR-10 and CIFAR-100 . The detailed setting is provided in Appendix C.1.

Table 2: Results on CIFAR-100 with WRN.

Dropout	Original Data			Noisy Data		
	Accuracy	PAvPU (0.05)	log likelihood	Accuracy	PAvPU (0.05)	log likelihood
Bernoulli	79.03	61.54	-4.49	52.01	54.25	-4.55
Gaussian	76.63	78.05	-3.93	51.38	57.02	-4.23
Concrete	79.19	64.14	-4.50	51.58	56.61	-4.56
Bernoulli Contextual	80.85±0.05	81.56±0.31	-3.56±0.02	53.64±0.45	**58.63**±0.50	**-3.73**±0.04
Gaussian Contextual	**80.93**±0.18	**81.69**±0.16	**-3.43**±0.07	**53.72**±0.34	58.49±0.43	-3.81 ±0.03

Results and analysis: We show the results for CIFAR-100 in Table 2 (see CIFAR-10 results in Tables 8-9 in Appendix). Accuracies, PAvPUs and test predictive loglikelihoods are incorporated for both the original and noisy data. We consistently observe contextual dropout outperforms other models in accuracy, uncertainty estimation, and loglikelihood.

Uncertainty visualization: We conducted extensive qualitative analyses for uncertainty evaluation. In Figures 9-11 in Appendix F.2, we visualize 15 CIFAR images (with true label) and compare the corresponding probability outputs of different dropouts in boxplots. We observe (1) contextual dropout predicts the correct answer if it is certain, (2) contextual dropout is certain and predicts the correct answers on many images for which MC or concrete dropout is uncertain, (3) MC or concrete dropout is uncertain about some easy examples or certain on some wrong predictions (see details in Appendix F.2), (4) on an image that all three methods have high uncertainty, contextual dropout places a higher probability on the correct answer than the other two. These observations verify that contextual dropout provides better calibrated uncertainty.

Large-scale experiments with ImageNet: Contextual dropout is also applied to the convolutional layers in ResNet-18, where we plug contextual dropout into a pretrained model, and fine-tune the pretrained model on ImageNet. In Table 3, we show it is even possible to finetune a pretrained model with contextual dropout module and without much additional memory or run time cost it achieves better performance than both the original model and the one with regular Gaussian dropout. Training model with contextual dropout from scratch can further improve the performance. See detailed experimental setting in Appendix C.1.

Table 3: Results on ImageNet with ResNet-18.

Dropout	Top-1 Acc	PAvPU	Params	sec/step
Without	69.75	NA	11.70M	1.44
+Gaussian	69.46	72.86	11.70M	1.50
+Contextual	**70.03**±0.07	**74.68**±0.08	11.88M	1.64
+Contextual (scratch)	**70.29**±0.09	**76.47**±0.12	11.88M	1.64

3.3 Contextual dropout on attention layers

We further apply contextual dropout to the attention layers of VQA models, whose goal is to provide an answer to a question relevant to the content of a given image. We conduct experiments on the commonly used benchmark, VQA-v2 (Goyal et al., 2017), containing human-annotated question-answer (QA) pairs. There are three types of questions: Yes/No, Number, and Other. In Figure 5, we show one example for each question type. There are 10 answers provided by 10 different human annotators for each question (see explanation of evaluation metrics in Appendix C.2). As shown in the

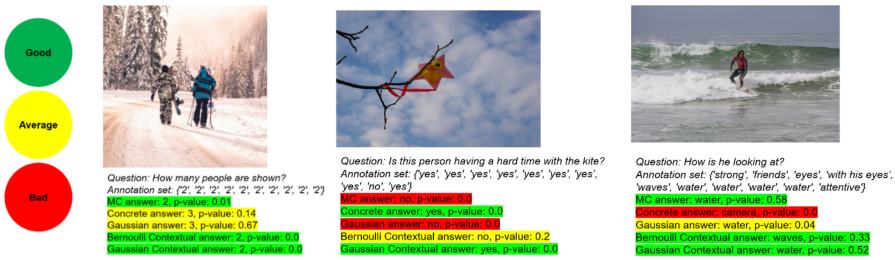

Figure 5: VQA visualization: for each question type, we present an image-question pair along with human annotations. We manually classify each prediction by different methods based on their answers and p-values. For questions that have a clear answer, we define the good as certain & accurate, the average as uncertain & accurate or uncertain & inaccurate, and the bad as certain & inaccurate. Otherwise, we define the good as uncertain & accurate, the average as certain & accurate or uncertain & inaccurate, and the bad as certain & inaccurate.

examples, VQA is generally so challenging that there are often several different human annotations for a given image. Therefore, good uncertainty estimation becomes even more necessary.

Model and training specifications: We use MCAN (Yu et al., 2019), an attention-based state-of-the-art model for VQA. Self-attention layers for question features and visual features, as well as the question-guided attention layers of visual features, are stacked one over another to build a deep model. Dropout is applied in every attention layer (after the softmax and before residual layer (Vaswani et al., 2017)) and fully-connected layer to prevent overfitting (Yu et al., 2019), resulting in 62 dropout layers in total. Detailed experiment setting is in Appendix C.2.

Results and analysis: We compare different dropouts on both the original VQA dataset and a noisy version, where Gaussian noise with standard deviation 5 is added to the visual features. In Tables 4, we show the overall accuracy and uncertainty estimation. The results show that on the original data, contextual dropout achieves better accuracy and uncertainty estimation than the others. Moreover, on noisy data, where the prediction becomes more challenging and requires more model flexibility and robustness, contextual dropouts outperform their regular dropout counterparts by a large margin in terms of accuracy with consistent improvement across all three question types.

Table 4: Accuracy and PAvPU on visual question answering.

Dropout	Accuracy		PAvPU	
	Original Data	Noisy Data	Original Data	Noisy Data
MC - Bernoulli	66.95	61.45	70.04	66.11
MC - Gaussian	66.96	62.75	70.77	67.42
Concrete	66.82	61.47	71.02	65.94
Bernoulli Contextual	**67.42**±0.06	63.73±0.08	**71.65**±0.06	68.57±0.11
Gaussian Contextual	67.35±0.03	**63.82**±0.05	71.62±0.02	**68.64**±0.04

Visualization: In Figures 12-15 in Appendix F.3, we visualize some image-question pairs, along with the human annotations and compare the predictions and uncertainty estimations of different dropouts. We show three of them in Figure 5. As shown in the plots, overall contextual dropout is more conservative on its wrong predictions and more certain on its correct predictions than other methods(see more detailed explanations in Appendix F.3).

4 Conclusion

We have proposed the contextual dropout as a simple and scalable data-dependent dropout module that achieves strong performance in both accuracy and uncertainty on a variety of tasks including large scale applications. With an efficient parameterization of the coviariate-dependent variational distribution, contextual dropout boosts the flexibility of Bayesian neural networks with only slightly increasing memory and computational cost. We demonstrate the general applicability of contextual dropout on fully connected, convolutional, and attention layers, and also show that contextual dropout masks are compatible with both Bernoulli and Gaussian distribution. On both image classification and visual question answering tasks, contextual dropout outperforms corresponding baselines consistently. Interestingly, on ImageNet, it is even possible to improve the performance of a pretrained model by adding the contextual dropout module during a finetuning stage. Based on these results, we believe contextual dropout can serve as an efficient alternative to data-independent dropouts in the versatile tool box of dropout modules.

REFERENCES

Jimmy Ba and Brendan Frey. Adaptive dropout for training deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 3084–3092, 2013.

Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, and Doina Precup. Conditional computation in neural networks for faster models. *arXiv preprint arXiv:1511.06297*, 2015.

Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.

Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424*, 2015.

Shahin Boluki, Randy Ardywibowo, Siamak Zamani Dadaneh, Mingyuan Zhou, and Xiaoning Qian. Learnable Bernoulli dropout for Bayesian deep learning. In *Artificial Intelligence and Statistics*, 2020.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.

Yuntian Deng, Yoon Kim, Justin Chiu, Demi Guo, and Alexander Rush. Latent alignment and variational attention. In *Advances in Neural Information Processing Systems*, pp. 9712–9724, 2018.

Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059, 2016.

Yarin Gal, Jiri Hron, and Alex Kendall. Concrete dropout. In *Advances in Neural Information Processing Systems*, pp. 3581–3590, 2017.

Asghar Ghasemi and Saleh Zahediasl. Normality tests for statistical analysis: a guide for non-statisticians. *International journal of endocrinology and metabolism*, 10(2):486, 2012.

Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6904–6913, 2017.

Alex Graves. Practical variational inference for neural networks. In *Advances in neural information processing systems*, pp. 2348–2356, 2011.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1321–1330. JMLR. org, 2017.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 2015.

Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347, 2013.

Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7132–7141, 2018.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Durk P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. In *Advances in Neural Information Processing Systems*, pp. 2575–2583, 2015.

Alex Krizhevsky et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.

Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. Accurate uncertainties for deep learning using calibrated regression. *arXiv preprint arXiv:1807.00263*, 2018.

Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, pp. 6402–6413, 2017.

Hugo Larochelle, Dumitru Erhan, Aaron Courville, James Bergstra, and Yoshua Bengio. An empirical evaluation of deep architectures on problems with many factors of variation. In *Proceedings of the 24th international conference on Machine learning*, pp. 473–480, 2007.

Yann LeCun, Corinna Cortes, and CJ Burges. MNIST handwritten digit database. *AT&T Labs [Online]. Available: http://yann. lecun. com/exdb/mnist*, 2:18, 2010.

Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

Chunyuan Li, Changyou Chen, David Carlson, and Lawrence Carin. Preconditioned stochastic gradient Langevin dynamics for deep neural networks. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.

Yang Li and Shihao Ji. L0-ARM: Network sparsification via stochastic binary optimization. In *The European Conference on Machine Learning (ECML)*, 2019.

Christos Louizos and Max Welling. Multiplicative normalizing flows for variational Bayesian neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2218–2227. JMLR. org, 2017.

David JC MacKay. A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3): 448–472, 1992.

Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Variational dropout sparsifies deep neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2498–2507. JMLR. org, 2017.

Jishnu Mukhoti and Yarin Gal. Evaluating bayesian deep learning methods for semantic segmentation. *arXiv preprint arXiv:1811.12709*, 2018.

Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using Bayesian binning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.

Yurii E Nesterov. A method for solving the convex programming problem with convergence rate o (1/kˆ 2). In *Dokl. akad. nauk Sssr*, volume 269, pp. 543–547, 1983.

Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.

Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pp. 91–99, 2015.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *ICML*, pp. 1278–1286, 2014.

Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.

Jiaxin Shi, Shengyang Sun, and Jun Zhu. Kernel implicit variational inference. In *International Conference on Learning Representations*, 2018. URL `https://openreview.net/forum?id=r1l4eQW0Z`.

Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1): 1929–1958, 2014.

Ravi Teja Mullapudi, William R Mark, Noam Shazeer, and Kayvon Fatahalian. Hydranets: Specialized dynamic architectures for efficient inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8080–8089, 2018.

Damien Teney, Peter Anderson, Xiaodong He, and Anton Van Den Hengel. Tips and tricks for visual question answering: Learnings from the 2017 challenge. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4223–4232, 2018.

Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christoph Bregler. Efficient object localization using convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 648–656, 2015.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.

Zhendong Wang and Mingyuan Zhou. Thompson sampling via local uncertainty. *arXiv preprint arXiv:1910.13673*, 2019.

Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 681–688, 2011.

Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pp. 5–32. Springer, 1992.

Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015a.

Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pp. 2048–2057, 2015b.

Mingzhang Yin and Mingyuan Zhou. ARM: Augment-REINFORCE-merge gradient for discrete latent variable models. *Preprint*, May 2018.

Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6281–6290, 2019.

Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.