

Fast Conditional Network Compression Using Bayesian HyperNetworks

Phuoc Nguyen, Truyen Tran, Ky Le, Sunil Gupta, Santu Rana, Dang Nguyen, Trong Nguyen, Shannon Ryan, and Svetha Venkatesh

A2I2, Deakin University, Geelong, Australia

{phuoc.nguyen, truyen.tran, k.le, sunil.gupta, santu.rana, d.nguyen, trong.nguyen, shannon.ryan, svetha.venkatesh}@deakin.edu.au

Abstract. We introduce a conditional compression problem and propose a fast framework for tackling it. The problem is how to quickly compress a pretrained large neural network into optimal smaller networks given target contexts, e.g. a context involving only a subset of classes or a context where only limited compute resource is available. To solve this, we propose an efficient Bayesian framework to compress a given large network into much smaller size tailored to meet each contextual requirement. We employ a hypernetwork to parameterize the posterior distribution of weights given conditional inputs and minimize a variational objective of this Bayesian neural network. To further reduce the network sizes, we propose a new input-output group sparsity factorization of weights to encourage more sparseness in the generated weights. Our methods can quickly generate compressed networks with significantly smaller sizes than baseline methods.

Keywords: Bayesian deep learning, meta-learning, network compression, Bayesian compression, hyper networks

1 Introduction

Modern deep neural networks used in computer vision and natural language processing are very large function approximators with millions of trainable parameters. Their flexibility has led to undesirable properties such as overparameterization [6] and memorization of random patterns [1]. The large size makes it difficult to run the networks on edge devices, and wastes computational resources when run on servers. While it is sometimes possible to handcraft or deliberately search for more compact networks, it is often easier for practitioners to take large proven networks and compress them to fit the hardware and time constraints.

Recent efforts in network compression [10, 19, 21, 20] have achieved significant progress in reducing large neural networks down to smaller sizes and cutting down run-time cost. The most popular approaches are to prune hidden units and filters of dense layers and convolutional filters or to quantize weights into lower bit precision. While the former helps slim down a network, thus reducing its size and total floating point operations (FLOPS), the latter can make these networks

run on limited compute hardware with integer arithmetic [14]. However, there is a compression-accuracy trade-off that prevent compressing large networks further and hinder their deployment in certain contexts with very low compute resources. It is largely unknown how to determine *a priori* if a certain degree of compression is achievable while satisfying context-specific constraints.

Addressing this limitation we introduce *the problem of contextual compression where a large network is conditionally compressed given operating constraints*. For instance, we need to deploy a large object recognition network trained on ImageNet with 1,000 classes in scenarios requiring only a subset of the classes. Using a single compressed network is clearly wasteful, and compressed networks are likely to bias minority classes [12]. Thus, there is an urgent need to optimally compress networks for different contexts.

Using only a partial network at inference time has been studied in a recent line of work known as conditional computation [3, 24, 2]. The strategy here is to use gating units to turn off unnecessary filters conditioned on inputs, therefore saving computation. While this reduces FLOPS at inference time, we still need to store and load the whole architecture onto run-time memory, making it unsuitable for edge devices.

To this end we propose a novel Bayesian framework for conditional compression given contextual run-time constraints. Key to the framework is a hypernetwork to parameterize the posterior distribution of weights of the target network given conditions. Given a pretrained large network, our hypernetwork generates compressed network weights for *on-the-fly* for each context. The conditional compression is implemented during training phase by minimizing a variational objective of the generated Bayesian neural network under different contexts. By using a sparsity inducing priors such as the scale mixture prior family [19] we can force the hypernetwork to generate very sparse hidden units networks, thus implementing context-based pruning. To encourage sparsifying the input and output neurons of layers jointly, we propose a novel input-output group sparsity factorization of the weight posterior distribution, in which the (input) group sparsity factorization of [19] is a special case. The framework is dubbed Bayesian Hypernetwork Compression (BHC).

We investigate the power of BHC on three general conditioning contexts: (a) given the subset of input features, (b) given the subset of output classes, and (c) constraints on expected model size. The model architectures include the popular LeNet-300-100 and LeNet-5-Caffe on MNIST dataset, and the VGG on CIFAR 10. Through a suite of experiments we demonstrate that BHC can quickly compress networks with sparse connectivity and low bit precision, resulting in significantly smaller memory footprints and computation than competing methods.

2 Related Work

Closest to our work are conditional computation methods [3, 5, 7, 2] which can predict and turn off unnecessary filters before running, thus saving computation

at inference time. However, these methods do not work for small compute devices. It also takes an extra step for the computation of which filters to include or remove. In [5] the masks for activation of each hidden layers are computed in parallel and for each sample independently, whereas in [27] the convolution kernels are parameterized using a condition from each input sample. In batch-shaping [2], large networks are slimmed down by using a residual network to predict masks. This is in contrast to our contextual compression proposal where a compressed network is generated once per context for run-time inference of any amount of data.

Also related to ours are meta-learning based compression methods. In [17], a hypernetwork is used to generate pruned convolution filters from sparse embedding vectors. In [18], a similar hypernetwork is used to generate various weights for a target network, then an evolutionary procedure is used to search for good-performing pruned networks. In contrast, we introduce a fast contextual compression problem where compressed networks is immediately generated given input conditions. We develop a principled Bayesian framework to model sparse posterior of weights which facilitates both pruning and bit precision reduction.

3 Preliminaries

3.1 Bayesian neural networks

A neural network is fully specified its computational graph whose nodes represent neural units equipped with activation functions and connectivity is specified by weight matrices. Let W be the model weights. Bayesian neural network [22] provides a principled way of modeling uncertainties over weights W . By choosing a suitable prior distribution $p(W)$, we can use Bayes formula to update the posterior distribution of weights given data \mathcal{D} as $p(W | \mathcal{D}) = \frac{p(\mathcal{D}|W)p(W)}{p(\mathcal{D})}$. This direct Bayesian inference is a long standing problem where the normalization term in the denominator requires marginalizing over the entire parameter space and is intractable to compute. Many recent advances in approximate inference for neural networks can be classified into sampling based or variational [4]. While the former can give better posterior approximation, recent literature focuses on the latter for neural networks due to its scalability [16, 23, 15]. Variational inference for the weight distribution often uses a simple posterior distribution such as normal distribution, $W \sim q(W) = \mathcal{N}(\mu_W, \Sigma_W)$ [22, 19]. Inference is then amounted to optimizing its parameters,

$$\max_{\mu_W, \Sigma_w} \mathbb{E}_{W \sim q(W|\mu_W, \Sigma_W)} p(\mathcal{D}|W) - D_{KL}(q(W|\mu_W, \Sigma_W) || p(W))$$

An emerging problem for training Bayesian deep networks with many layers is in mini-batch training to optimize its posterior distribution. While reparameterization trick [16, 23] can be used to sample one set of weights per update, it requires independent weights for independent inputs in each data mini-batch. Authors in [15] observed that by factorizing weights and noise and grouping noise

by input neurons, they can “apply” the noise into the input instead of weights. Then the (preactivation) output of the layer will be a normal distribution, linear transformation of the normal input vector [15],

$$y = x(Z \odot W) = (x \odot z)W$$

$$y_j \sim \mathcal{N}\left(\sum_i w_{ij} z_i x_i, \sum_i w_{ij}^2 z_i^2 x_i^2\right)$$

where the noise vector z is reparameterized from the noise matrix Z by grouping by input neurons, x is the input vector, and $y = (y_j)$ is the output vector of the layer. This interpretation is used in variational dropout [21] for sparsifying networks.

3.2 Bayesian compression

In [19], the authors explicitly describe the sparsity encouraging prior for neural network parameters by using a scale mixture prior from Bayesian statistics. Thereby, minimizing the KL divergence from this sparse prior to the weight posterior distribution will force weights to be sparse. This Bayesian compression framework for deep learning can be stated as follows.

Let us consider a particular weight matrix $[w_{ij}]_{i \in \overline{1, I}; j \in \overline{1, J}}$, where w_{ij} is the connection from the i -th input unit to the j -th output unit, for I input units and J output units. We define a scale mixtures of normals for the prior over parameter w ,

$$w_{ij} \sim \mathcal{N}(w_{ij}; 0, z_{ij}^2)$$

where $z \sim p(z)$ is some sparse prior. Marginalizing over z gives heavier tails as well as peaks at zero, thus resulting in a sparse network. Instead of using one scale for each weight, it is beneficial to use one scale z_i for all connections from neuron i . This has the effect of pruning (dropping) the entire neuron. The prior distribution of weight and scale variables are designed as:

$$p(z, W) = \prod_{i=1}^I p(z_i) \prod_{i,j}^{I,J} N(w_{ij} | 0, z_i^2) \quad (1)$$

A popular choice for the scale variable is the normal-Jeffreys prior $p(z) = \frac{1}{|z|}$. This improper prior (not normalizable) is non-informative and non parametric which make it suitable to model the scale for weights since it allows both small scales (weights are peak at zeros) and large scales (weights can spread out without being shrunk). Due to its irregularity, numerical approximation need to be made when computing KL divergences [15, 21]. We suggest the references in [19, 8] for more prior choices.

The posterior of weights and scales are approximated by normal distributions and they are factorized in the same way as the prior:

$$q_\phi(z, W) = \prod_{i=1}^I N(z_i | \mu_{z_i}, \sigma_z^2) \prod_{i,j}^{I,J} N(w_{ij} | z_i \mu_{ij}, z_i^2 \sigma_{ij}^2)$$

Note that σ_z^2 was reparameterized as $\mu_{z_i}^2 \alpha_i$ in [26, 15] for the Gaussian dropout interpretation, where α_i is the dropout rate of the neuron i . This dropout rate α_i can be later compared to a threshold to decide a dropout mask for the input neurons.

The optimization of the variational parameters $\phi = (\mu_{ij}, \sigma_{ij}^2, \mu_{z_i}, \sigma_{z_i}^2)$ is minimizing a variational lower bound of the data log-likelihood $\log p(\mathcal{D})$:

$$L(\phi) = \mathbb{E}_{q_\phi(z)q_\phi(W|z)} \log p(\mathcal{D} | W) - \mathbb{E}_{q_\phi(z)} D_{KL}(q_\phi(W | z) \| p(W | z)) - D_{KL}(q_\phi(z) \| p(z))$$

At test time, one can recover a deterministic network by replacing the distribution of W by its posterior mean times a dropout mask.

4 Conditional Compression

While the Bayesian compression framework of [19], explained in Section 3.2, provides a principled way for unconditional compression, it is not clear how to tune it to work in practical contexts with specific resource constraints. Thus we formulate a *new problem of conditional compression*, where we infer the network sparsity and bit precision *on-the-fly* conditioned on a specified context.

The compression context can be ideally arbitrary. In this work, we focus on three general types of conditioning context: (a) the subset of input features, (b) the subset of output classes, and (c) the expected model size. The input feature condition means that the feature or data distribution is limited to some subspace or input region, for example a subset of image channels or a subset of vocabulary. The output classes condition means that the classification output is limited to some sub-classes. The model size constraint reflects resource specification such as memory footprint, and this translates to reaching an equivalent compression rate for a given pre-trained network.

We then propose a new Bayesian framework to solve the problem, employing neural hypernetworks [9] to generate the distribution of compressed networks given the conditions. We term this new framework BHC, which stands for **B**ayesian **H**ypernetwork **C**ompression.

4.1 Conditional Bayesian hypernetworks

Let W_0 be the pretrained network weights trained on the dataset $\mathcal{D}_0 = \{(x_0^n, y_0^n)\}_{n=1}^N$. We aim to quickly generate a compressed network W_t from W_0 in some context t with a constraint on the data domain or the compute resource, denoted by the condition c_t , i.e., $W_t = f_\gamma(c_t, W_0)$ for some hypernetwork f_γ parameterized by γ . In theory we can train f by running external compression algorithms with varying conditions to collect training data, but this is extremely expensive to cover the large condition space.




Fig. 1. The graphical model of our conditional compression framework. W_0 is the pre-trained network weights trained on the dataset $\mathcal{D}_0 = \{(x_0^n, y_0^n)\}_{n=1}^N$. W_0 is used to initialize the hypernetwork γ to generate the initial means and variances of weights W_t . For each context t , the hypernetwork γ generates a compressed W_t conditioned on z_t (and c_t implicitly). Then it is used for the classification context on $\mathcal{D}_t = \{(x_t^m, y_t^m)\}_{m=1}^M$. When W_0 is absent, the framework can also be trained from scratch to generate W_t for each context t .

Inspired by the Bayesian compression framework of [19] (see Section 3.2), we introduce a new efficient Bayesian framework, in that W_t is probabilistically generated by the hypernetwork at training time and the posterior of W_t from the condition c_t is estimated at inference time. Fig. 1 shows the graphical model of the proposed method.

In this framework, we are given a set of T classification scenarios (or contexts) $\mathcal{D} = \{(\mathcal{D}_t, c_t)\}_{t=1}^T$, each containing classification data $\mathcal{D}_t = \{(x_t^m, y_t^m)\}_{m=1}^M$ and constraint c_t . Marginalizing over all condition c_t 's then gives the posterior distribution of W_t , $q(W) = \int q(W_t | c_t) p(c_t) dc_t$, which is a continuous mixture distribution with mixing density $p(c_t)$. The compression is done by minimizing the KL distance from a normal-Jeffreys group sparsity prior to the weight posterior [19]. At test time, the hypernetwork generates a mask and the posterior mean of W_t to produce a compressed deterministic network for classification on \mathcal{D}_t .

We solve this problem by modeling the joint distribution between the parameter W , the dropout mask z , and the condition c . First, we assume each mask is a latent variable associated with a given condition, and their joint distribution factors as $p(c, z) = p(c|z)p(z)$. We aim at learning an optimal variational posterior distribution $q_\phi(z|c_t)$ given this condition such that the weight W is sparse and still performs well on the (conditioned) dataset \mathcal{D} . That is we maximize the variational objective:

$$\begin{aligned} \max_{\phi} \mathbb{E}_{z \sim q_\phi(z|c)} [\log p(\mathcal{D} | W, z) + \log p(c | z)] \\ - \mathbb{E}_{q_\phi(z)} [D_{KL}(q_\phi(W|z) \| p(W|z))] - D_{KL}(q_\phi(z|c) \| p(z)) \end{aligned} \quad (2)$$

4.2 Compression contexts

For the sparsity condition that directly limits the network capacity, we observe that varying the cutoff threshold on the logarithm of dropout rates can produce networks with varying compression rate. The normal-Jeffreys prior give rise to a continuum of the dropout domain in the scale posterior which supports compression at various rate. Whereas the horseshoe prior tends to separate the compression rates into two isolated clusters. However, even for the normal-Jeffreys prior, we observe that the accuracy will degrade significantly if the required compression rate is significantly higher than its attained compression rate at convergence. Therefore we proposed an additional level of compression in addition to the group sparsity induced by the normal-Jeffreys prior and use the hypernetwork to select the level of compression given the required sparsity condition.

4.3 Group sparsity and input-output group sparsity reparameterization

In [19], the author proposed coupling the scales of weights belonging to the same input neuron by sharing the scale variable z in the joint prior as well as in the joint posterior,

$$p(W, z) \propto \prod_i^A \frac{1}{|z|} \prod_{i,j}^{A,B} \mathcal{N}(w_{ij}|0, z_i^2)$$

$$q_\phi(W, z) = \prod_i^A \mathcal{N}(z_i|\mu_{z_i}, \sigma_{z_i}^2) \prod_{i,j}^{A,B} \mathcal{N}(w_{ij}|z_i\mu_{ij}, z_i^2\sigma_{ij}^2)$$

This group sparsity distribution $p(z)$ is applied to the input of a layer with weight W . For deep networks with multiple layers $\{W_l\}_1^L$, we observe that the input sparsity z^l of a layer W_l should directly affect the input sparsity of the next layers. Therefore we introduce an additional scale variable s for the output and propose to use a joint input-output group sparsity as follow,

$$p(W, z, s) \propto \prod_i^A \frac{1}{|z_i|} \prod_j^B \frac{1}{|s_j|} \prod_{i,j}^{A,B} \mathcal{N}(w_{ij}|0, z_i^2 s_j^2)$$

$$q_\phi(W, z, s) = \prod_i^A \mathcal{N}(z_i|\mu_{z_i}, \sigma_{z_i}^2) \prod_j^B \mathcal{N}(s_j|\mu_{s_j}, \sigma_{s_j}^2) \prod_{i,j}^{A,B} \mathcal{N}(w_{ij}|z_i s_j \mu_{ij}, z_i^2 s_j^2 \sigma_{ij}^2)$$

Note that this method applies to activation functions which have a zero fixed point, such as tanh, RELU, and ELU families, in order for the output group sparsity to carry on to the next layer input¹. This creates a “pairwise” sparsity

¹ We assume the bias, if exist, can be included into the weight matrix and the input is augmented by a constant one.

coupling one layer and the next layer since the output feature map of one layer will be the input to the next layer, where its sparsity pattern will be taken into account. This coupling of the input-output sparsity directly help increase the sparsity level of the network, thus can be used to achieve higher compression level. In the forward step the weights W is sampled conditioned on the sparsity of both the input, via z , and the output, via s . In the backward step, the gradient of W is passed directly to update the input and output scales in pair.

The variational lower bound under this new prior and posterior is:

$$\begin{aligned} \mathcal{L}(\phi) = & \mathbb{E}_{q_\phi(z,s)q_\phi(W|z,s)}[\log p(\mathcal{D}|W)] + \\ & - \lambda_{KL} (\mathbb{E}_{q_\phi(z,s)}[D_{KL}(q_\phi(W|z,s)||p(W|z,s))] + D_{KL}(q_\phi(z,s)||p(z,s))) \end{aligned} \quad (3)$$

where λ_{KL} is the weighting factor to control the level of compression, which we will utilize in the experiments for different level of compression requirements. The negative KL divergence from the joint normal-Jeffreys scale prior to the joint Gaussian posterior in this new parameterization is as follow:

$$\begin{aligned} -D_{KL}(q_\phi(z,s)||p(z,s)) \approx & \sum_i^A (k_1\sigma(k_2 + k_3 \log \alpha_i) - 0.5m(-\log \alpha_i) - k_1) \\ & + \sum_j^B (k_1\sigma(k_2 + k_3 \log \beta_j) - 0.5m(-\log \beta_j) - k_1), \end{aligned}$$

where $\sigma(\cdot)$, $m(\cdot)$ are the sigmoid and softplus functions and $k_1 = 0.63576$, $k_2 = 1.87320$, $k_3 = 1.48695$, and $\log \alpha_i$ and $\log \beta_j$ are the dropout-rate vectors of the input and output respectively.

At test time, we use the posterior mean of W and apply the masks $Z = m_z m_s^T$ to output the deterministic weight matrix,

$$\hat{W} = Z \odot \mathbb{E}_{q_\phi(z,s)q_\phi(\tilde{W})} [\text{diag}(z)\tilde{W}\text{diag}(s)] = \text{diag}(m_z \odot \mu_z) M_w \text{diag}(m_s \odot \mu_s) \quad (4)$$

where the mask $Z = m_z m_s^T$ is the binary mask matrix determined using the input mask m_z and output mask m_s determined by thresholding z and s , and $M_w = (\mu_{ij})$ is the matrix of weight means. We use the following variational posterior marginal variance to determine the bit precision of the weights:

$$\mathbb{V}(w_{ij}) = \mathbb{V}(z_i \tilde{w}_{ij} s_j) = (\sigma_z^2 + \mu_z^2)(\sigma_{ij}^2 + \mu_{ij}^2)(\sigma_s^2 + \mu_s^2) - \mu_z^2 \mu_{ij}^2 \mu_s^2$$

4.4 Training methods

We train the Bayesian compression models then quickly adapt to each condition using our framework. Therefore, when there is no condition, our method generates the compressed network which matches the performance of existing Bayesian compression networks. However when there are extra context information, we can

input into our hypernetwork to generate a tailored, highly compressed network for that specific context.

Algorithm 1: BHC feedforward pass for fully connected layers.

Require: input H , hypernetwork f_ϕ , condition $c \sim p(c)$

- 1 $M_w, \Sigma_w, \mu_z, \sigma_z, \mu_s, \sigma_s = f_\phi(c)$
- 2 $E_z \sim \mathcal{N}(0, 1)$
- 3 $Z = \mu_z + \sigma_z \odot E_z$
- 4 $E_s \sim \mathcal{N}(0, 1)$
- 5 $S = \mu_s + \sigma_s \odot E_s$
- 6 $M_h = ((H \odot Z)M_w) \odot S$
- 7 $V_h = ((H \odot Z)^2 \Sigma_w) \odot S^2$
- 8 $E \sim \mathcal{N}(0, 1)$
- 9 return $M_h + \sqrt{V_h} \odot E$

Algorithm 2: BHC feedforward pass for convolution layers.

Require: input H , hypernetwork f_ϕ , condition $c \sim p(c)$

- 1 $M_w, \Sigma_w, \mu_z, \sigma_z, \mu_s, \sigma_s = f_\phi(c)$
- 2 $E_z \sim \mathcal{N}(0, 1)$
- 3 $Z = \mu_z + \sigma_z \odot E_z$
- 4 $E_s \sim \mathcal{N}(0, 1)$
- 5 $S = \text{reshape}(\mu_s + \sigma_s \odot E_s, [1, 1, F, N])$
- 6 $M_h = ((H \odot Z) * M_w) \odot S$
- 7 $V_h = ((H \odot Z)^2 * \Sigma_w) \odot S^2$
- 8 $E \sim \mathcal{N}(0, 1)$
- 9 return $M_h + \sqrt{V_h} \odot E$

Algorithms 1 and 2 show the forward pass of BHC networks for fully connected and convolution layers respectively. In each algorithm, $M_w = (\mu_{ij})$ and $\Sigma_w = (\sigma_{ij}^2)$ are the generated means and variances of each layers, μ_z, σ_z are the means and variances of the input group scale variables, μ_s, σ_s are the means and variances of the output group scale variables, H is the input to the current layer, N is the batch size, F is the number of convolutional filters, and $*$ is the convolutional operator. For fully connected layers, we use one input (output) scale variable for each input (output) neuron to group weights connected to that neuron. For convolutional layers, we use only one input scale variable for all filters and one output scale variable for each filter output. We use the [width, height, filters, batch] dimension ordering. Local reparameterization [15] is used to efficiently sample the input, output group scale variables, and the activations. Both algorithms work as follow. In step 1, the mean and variance parameters of weights, input and output masks posterior distributions are generated given the condition c . In step 2-3, an input noise matrix Z is sampled for the input matrix H . In step 4-5, an output noise matrix/tensor S is sampled for the output matrix/tensor. In step 6 and 7, the output mean M_h and variance V_h is calculated and masked accordingly. Finally, an output sample is return by step 8-9.

5 Experiments

5.1 Settings

We design the experiments with varying conditions and test whether our Bayesian Hypernetwork Compression (BHC) method can compress and speedup a target network tailored to each condition. We examine three classes of conditions: *conditions on the input domain*, *conditions on the output classes*, and *conditions on the network capacity*.

Network architectures: We investigate our method on the well-known neural network architectures of LeNet-300-100 and LeNet-5-Caffe on MNIST dataset, similarly with [21], and VGG on CIFAR 10, similarly with [19]. The input (output) groups of weight parameters are designed by coupling the scale variable for each input (pre-activation output) neuron for fully connected layers. For convolutional layers, we use a single scale variable for the input and coupling the scale variable for each filter. We simply use the same threshold for all layers for pruning and vary this threshold to determine the desired compression rate for each context². We simply ignore the reconstruction loss of the condition, $\log p(c | z)$, in the optimization objective, Eq. 2, and consider the conditions in the train and test phases come from similar distributions to simplify the optimization problem. We leave this investigation for future work³.

Data preparation: In each experiment, we divide the dataset $\mathcal{D}_{\text{train}}$ test set $\mathcal{D}_{\text{test}}$ into tasks (contexts), add in a condition in each task, and finally create conditional datasets $\mathcal{D}_{\text{train}} = \{(X_t, Y_t), c_t\}_{t=1}^T$ and $\mathcal{D}_{\text{test}} = \{(X_t, Y_t), c_t\}_{t=1}^{T'}$. Our proposed networks are initialized with the given target network weights, then trained in minibatches where each batch is a different task drawn randomly from $\mathcal{D}_{\text{train}}$. In the test phase, the hyper-network will generate the posterior mean over weights and masks given the condition. A threshold can then be varied and chosen to set the binary mask for weights such that a final compression rate is achieved for the final deterministic weights.

Baselines: We use the following baselines: sparse variational dropout (SparseVD) [21], Bayesian compression with group normal-Jeffreys (BC-GNJ) and with group horseshoe (BC-GHS) [20]. We also tested the input-output group sparsity method separately on the Lenet-300-100 architecture (denoted BC-IO-GNJ) and demonstrated that it has better compression rate than its BC-GNJ counterpart⁴. We will use the proposed input-output group sparsity method for our BHC networks.

² Better thresholds can be chosen by visually inspection on the log dropout rates.

³ This involves adding a generative network for the condition posterior given all layer-masks then maximizing its log-likelihood.

⁴ We also tested BC-IO-GNJ on Lenet-5-Caffe and got better performance.

5.2 Implementation

Our method was implemented in the Flux framework [13]. We used Adam optimizer with default parameters and train our models for 300 epochs, and found it converges better when clipping the variances of the first layers of the networks as in [19]. Hence the first layer variance of LeNet-300-100 was clipped to 0.2^2 , the first layer variance of LeNet-5-Caffe to 0.5^2 , and the 64 and 128 feature maps layer in VGG to 0.1^2 and the 256 feature maps layers to 0.2^2 .

For *subclass conditioning* we use a 2-layer MLP embedding network, each with hidden size 100 to embed the subclasses condition vector. For *input domain conditioning* we use a 2 layers convolutional neural network each with 20 filters, kernel width 3, and zero padding to embed the input domain condition vector.

The hypernetwork is a set of linear layers to map the embedding vectors into the posterior weights parameters: the posterior means, log variances, the input and output group scales. We initialize these hypernetwork weights using Kaiming uniform function [11] with gain 0.5. We initialize the hypernetworks such that it generates the pretrained weight W_0 with noise initially as follows. We set the bias of the mean-generating hypernetwork to W_0 , the bias of log-variance-generating hypernetwork to a small number drawn from $\mathcal{N}(-9, 1e-4)$. For the scales, we set the bias of the scale-mean-generating hypernetwork to $\mathcal{N}(1, 1e-8)$, and the bias of its log-variance-generating hypernetwork to a small number drawn from $\mathcal{N}(-9, 1e-4)$.

5.3 Conditioning on output subclasses and input domain

For the condition on output subclasses, we take the vector of Bernoulli probabilities of the subclasses as condition, thus each dimension contains a probability representing the appearance of a class. We assume that the number of practical conditions is only a subset of all possible subsets, the power set 2^C where C is the number of classes. In this experiments we limit our investigation to conditions that are the two consecutive random classes between 1 and C . In training phase, c_y is the Bernoulli vector of length C where each component represents the probability of appearance of that class. In test phase, the c_y is the binary vector representing the appearance of the classes.

For the condition on the input domain, we use minibatches as tasks and use features from the minibatches as condition. We take the sample mean of the input distribution as the condition c_x , i.e the arithmetic mean of the input tensor along the batch dimension. In practice, this can be the sample mean of the distribution of a representative dataset for the testing condition.

Results on compression rates We present the weight compression rates of our methods to the baseline methods in Table 1. We use (i) compression by pruning input and output neurons and (ii) compression by reducing bit precision in addition to pruning the input and output neurons. These scenarios are practical since pruned weights can be used to create slimmer network, where as reduced precision weights can be exploited in future hardware [25]. We can observe that

our method can achieve higher compression rates at lower error rates than the baselines. When condition on the subclasses, the compression rates is highest. When exploiting the weight variance to reduce bit precision in combination with pruning, a much higher compression rate is achieved, $216\times$ for the LeNet-300-100 architecture, $1004\times$ for the LeNet-5-Caffe architecture, and $73\times$ for VGG architecture in this conditional setting.

Table 1. Compression results of Bayesian-hypernet compression (BHC) compared to sparse variational dropout (SparseVD), group normal-Jeffreys (BC-GNJ) and group horseshoe (BC-GHS). Results marked with * are from [19]. BHC method conditionally compresses the target network based on the domain input condition c_x or the subclass condition c_y in each context. We report the mean compression rate across tasks and the overall error on the testset.

Model	Method	Pruning (Error %)	Pruning and bit reduction (Error %)
LeNet-300-100 1.6	SparseVD*	21 (1.8)	84 (1.8)
	BC-GNJ*	9 (1.8)	36 (1.8)
	BC-GHS*	9 (1.8)	23 (1.9)
	BC-IO-GNJ	12 (1.8)	40 (1.9)
	BHC- c_x	26 (1.9)	169 (1.9)
	BHC- c_y	54 (1.8)	216 (1.9)
LeNet-5-Caffe 0.9	SparseVD*	63 (1.0)	228 (1.0)
	BC-GNJ*	108 (1.0)	361 (1.0)
	BC-GHS*	156 (1.0)	419 (1.0)
	BHC- c_x	187 (0.2)	608 (0.4)
	BHC- c_y	286 (0.4)	1004 (0.5)
VGG 8.4	BC-GNJ*	14 (8.6)	56 (8.8)
	BC-GHS*	18 (9.0)	59 (9.0)
	BHC- c_x	20 (7.0)	66 (7.4)
	BHC- c_y	22 (2.0)	73 (2.8)

Compressed models We present the neuron group sparsity enforcing capabilities of our methods compared to baselines in Table. 2. We show the pruned architecture and bit-precision per layer. We observe that our methods can generate very sparse input and output neurons of the target architecture. Especially, the input layer and the first hidden layers have significantly sparser connection and only need as low as 5 bit precision compared to the baselines. For Lenet-5-Caffe and VGG architecture, the first and last layers require a much smaller number of neurons. For Lenet-5-Caffe architecture, as low as 10 bit precision can be achieved. For VGG, the number of convolutional filters is also reduced significantly compared to sparse variational dropout, group normal-Jeffreys, and group horseshoe.

5.4 Conditioning on the model compression rate

In this section, we investigate the Bayesian-hypernet compression method given different sparsity constraints. We compare BHC to Bayesian compression methods using normal-Jeffreys and horseshoe priors. We implement BHC using two broad levels of compression as condition: (1) high accuracy compression and (2) high compression rate with moderate accuracy drop. Finer compression rates can be chosen by varying the threshold in each level. Therefore, we use a binary condition input to the hypernetwork, $c_w = 0$ for high accuracy compression and $c_w = 1$ for high compression rate. In this case, the condition is discrete and BHC becomes a finite mixture model:

$$p(W) = \int p(W|c)p(c)dc = p(W|c)p(c=0) + p(W|c)p(c=1)$$

with mixing probabilities $p(c=0) = p(c=1) = \frac{1}{2}$. First, the user need to inputs $c_w \in \{0, 1\}$ into BHC hypernetwork and get as output the posterior mean of weights $M_w = (\mu_{ij})$ and a dropout-rate vector $\log \alpha$. Then the user varies a threshold τ to create an input mask $m_z = I(\log \alpha < \tau)$ and output mask $m_s = I(\log \beta < \tau)$ for z and s and calculates the deterministic weight as $\hat{W} = \text{diag}(m_z) \odot M_w \odot \text{diag}(m_s)$, as in Eq. 4. By checking the compression rate and the accuracy of \hat{W} on a validation set if it is available, the user can decide a final threshold τ to use.

We train BHC for $c_w = 0$ condition using a lower KL weighting, $\lambda_{KL} = 0.1$ in the training objective (Eq. 3), and we use a higher KL weighting, $\lambda_{KL} = 1$ for $c_w = 1$ condition. To speed up convergence, we choose a simple training strategy with this binary condition as follows⁵. We train the BHC hypernetwork with $c_w = 0$ condition for 100 epochs with $\lambda_{KL} = 0.1$ and freeze this generated weights in the the hypernetwork bias output, by setting the trainable flag of the bias to false. Next, we train the BHC hypernetwork with $c_w = 1$ condition for another 200 epochs with $\lambda_{KL} = 1$.

Results Fig. 2 compares our Bayesian-hypernet compression (BHC) with Bayesian compression with normal Jeffreys and horseshoe priors by varying the threshold to output deterministic weights. As can be observed, our methods can generate models with high accuracy at low compression rate and with only moderate accuracy drop at high compression rate.

6 Conclusion

We introduced a new problem concerning contextual compression of deep neural networks. Given a large uncompressed pre-trained network, we compress it *on-the-fly* to satisfy resource constraints at run time, e.g., compression rate, memory, input size or class subset. Our work provides the first principled Bayesian

⁵ In general, the condition can be sampled randomly from Bernoulli($\frac{1}{2}$) during training.

Table 2. Pruned architectures and bit precision of Bayesian-hypernet compression (BHC) compared to sparse variational dropout (SparseVD) [21], group normal-Jeffreys and group horseshoe [20]. Results marked with * are from [19].

Network & size	Method	Pruned architecture	Bit-precision
Lenet-300-100 784-300-100	SparseVD*	512-114-72	8-11-14
	BC-GNJ*	278-98-13	8-9-14
	BC-GHS*	311-86-14	13-11-10
	BC-IO- GNJ	255-90-15	7-9-15
	BHC- c_x	183-50-18	5-8-20
	BHC- c_y	141-30-18	6-11-20
Lenet-5-Caffe 20-50-800-500	SparseVD*	14-19-242-131	13-10-8-12
	BC-GNJ*	8-13-88-13	18-10-7-9
	BC-GHS*	5-10-76-16	10-10-14-13
	BHC- c_x	4-4-64-2	12-10-11-11
	BHC- c_y	4-4-64-3	12-9-11-9
VGG (2×64)-(2×128)- (3×256)-(8×512)	BC-GNJ*	63-64-128-128-245-155-63- -26-24-20-14-12-11-11-15	10-10-10-10-8-8-8- -5-5-5-5-5-6-7-11
	BC-GHS*	51-62-125-128-228-129-38- -13-9-6-5-6-6-6-20	11-12-9-14-10-8-5- -5-6-6-6-8-11-17-10
	BHC- c_x	52-63-126-128-161-71- -22-19-11-10-6-7-8-8-9-13	12-7-5-7-6-8-9-8- -11-12-12-12-14-14-26-20
	BHC- c_y	42-63-125-123-110-44-15- -12-8-6-5-5-6-6-6-8	12-11-8-13-6-9-10- -8-9-9-9-10-14-18-27-24

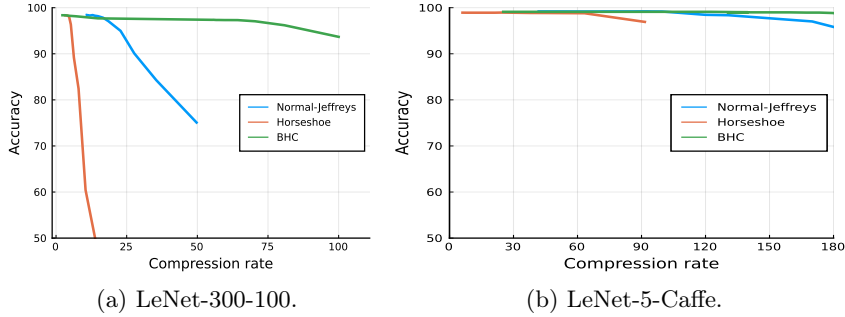


Fig. 2. Conditions on compression rate.

hypernetwork compression framework towards contextual compression for edge devices. Testing on several popular architectures on MNIST and CIFAR 10, we demonstrated that our methods can quickly generate compressed networks with significantly smaller memory footprints and computation than competing methods. Future challenges would be scaling up the Bayesian hypernetwork compression method to recent large architectures with hundred of millions of parameters and applying the method in real world contextual compression tasks and other data domains such as texts and sounds.

Acknowledgments

This research was a collaboration between the Commonwealth Australia (represented by Department of Defence) and Deakin University, through a Defence Science Partnerships agreement.

References

1. Devansh Arpit, Stanisław Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International Conference on Machine Learning*, pages 233–242. PMLR, 2017.
2. Babak Ehteshami Bejnordi, Tijmen Blankevoort, and Max Welling. Batch-shaping for learning conditional channel gated networks. In *International Conference on Learning Representations*, 2019.
3. Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
4. 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.
5. Zhouong Chen, Yang Li, Samy Bengio, and Si Si. You look twice: Gaternet for dynamic filter selection in cnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9172–9180, 2019.
6. Cong Fang, Hanze Dong, and Tong Zhang. Mathematical models of overparameterized neural networks. *Proceedings of the IEEE*, 2021.
7. Xitong Gao, Yiren Zhao, Łukasz Dudziak, Robert Mullins, and Cheng-zhong Xu. Dynamic channel pruning: Feature boosting and suppression. In *International Conference on Learning Representations*, 2018.
8. Soumya Ghosh, Jiayu Yao, and Finale Doshi-Velez. Structured variational learning of bayesian neural networks with horseshoe priors. In *International Conference on Machine Learning*, pages 1744–1753. PMLR, 2018.
9. David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *ICLR*, 2017.
10. Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *ICLR*, 2016.
11. 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*, pages 1026–1034, 2015.

12. Sara Hooker, Nyalleng Moorosi, Gregory Clark, Samy Bengio, and Emily Denton. Characterising bias in compressed models. *ICML Workshop on Human Interpretability in Machine Learning*, 2020.
13. Michael Innes, Elliot Saba, Keno Fischer, Dhairya Gandhi, Marco Concetto Rudilosso, Neethu Mariya Joy, Tejan Karmali, Avik Pal, and Viral Shah. Fashionable modelling with flux. *CoRR*, abs/1811.01457, 2018.
14. Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2704–2713, 2018.
15. Diederik P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. In *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 2*, pages 2575–2583, 2015.
16. Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *ICLR*, 2014.
17. Yawei Li, Shuhang Gu, Kai Zhang, Luc Van Gool, and Radu Timofte. Dhp: Differentiable meta pruning via hypernetworks. *arXiv preprint arXiv:2003.13683*, 2020.
18. Zechun Liu, Haoyuan Mu, Xiangyu Zhang, Zichao Guo, Xin Yang, Kwang-Ting Cheng, and Jian Sun. Metapruning: Meta learning for automatic neural network channel pruning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3296–3305, 2019.
19. Christos Louizos, Karen Ullrich, and Max Welling. Bayesian compression for deep learning. In *Advances in neural information processing systems*, pages 3288–3298, 2017.
20. Christos Louizos, Max Welling, and Diederik P Kingma. Learning sparse neural networks through l₀ regularization. In *International Conference on Learning Representations*, 2018.
21. Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Variational dropout sparsifies deep neural networks. In *International Conference on Machine Learning*, pages 2498–2507. PMLR, 2017.
22. Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
23. Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic back-propagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR, 2014.
24. 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.
25. Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated quantization with mixed precision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8612–8620, 2019.
26. Sida Wang and Christopher Manning. Fast dropout training. In *international conference on machine learning*, pages 118–126. PMLR, 2013.
27. Brandon Yang, Gabriel Bender, Quoc V Le, and Jiquan Ngiam. Condconv: Conditionally parameterized convolutions for efficient inference. In *Advances in Neural Information Processing Systems*, pages 1307–1318, 2019.