

REDUCING THE COMPUTATIONAL COST OF DEEP GENERATIVE MODELS WITH BINARY NEURAL NETWORKS

Thomas Bird*, Friso H. Kingma[†] & David Barber*

* Department of Computer Science, University College London

<firstname>.<surname>@cs.ucl.ac.uk

[†] fhkingma@gmail.com

ABSTRACT

Deep generative models provide a powerful set of tools to understand real-world data. But as these models improve, they increase in size and complexity, so their computational cost in memory and execution time grows. Using binary weights in neural networks is one method which has shown promise in reducing this cost. However, whether binary neural networks can be used in generative models is an open problem. In this work we show, for the first time, that we can successfully train generative models which utilize binary neural networks. This reduces the computational cost of the models massively. We develop a new class of binary weight normalization, and provide insights for architecture designs of these binarized generative models. We demonstrate that two state-of-the-art deep generative models, the ResNet VAE and Flow++ models, can be binarized effectively using these techniques. We train binary models that achieve loss values close to those of the regular models but are 90%-94% smaller in size, and also allow significant speed-ups in execution time.

1 INTRODUCTION

As machine learning models continue to grow in number of parameters, there is a corresponding effort to try and reduce the ever-increasing memory and computational requirements that these models incur. One method to make models more efficient is to use neural networks with weights and possibly activations restricted to be binary-valued (Courbariaux et al., 2015; Courbariaux et al., 2016; Rastegari et al., 2016; McDonnell, 2018; Gu et al., 2018). Binary weights and activations require significantly less memory, and also admit faster low-level implementations of key operations such as linear transformations than when using the usual floating-point precision.

Although the application of binary neural networks for classification is relatively well-studied, there has been no research that we are aware of that has examined whether binary neural networks can be used effectively in *unsupervised* learning problems. Indeed, many of the deep generative models that are popular for unsupervised learning do have high parameter counts and are computationally expensive (Vaswani et al., 2017; Maaløe et al., 2019; Ho et al., 2019a). These models would stand to benefit significantly from converting the weights and activations to binary values, which we call *binarization* for brevity.

In this work we focus on non-autoregressive models with explicit densities. One such class of density model is the variational autoencoder (VAE) (Kingma & Welling, 2014; Rezende et al., 2014), a latent variable model which has been used to model many high-dimensional data domains accurately. The state-of-the-art VAE models tend to have deep hierarchies of latent layers, and have demonstrated good performance relative to comparable modelling approaches (Ranganath et al., 2016; Kingma et al., 2016; Maaløe et al., 2019). Whilst this deep hierarchy makes the model powerful, the model size and compute requirements increases with the number of latent layers, making very deep models resource intensive.

Another class of density model which has shown promising results are flow-based generative models (Dinh et al., 2014; Rezende & Mohamed, 2015; Dinh et al., 2017). These models perform a series

of invertible transformations to a simple density, with the transformed density approximating the data-generating distribution. Flow models which achieve state-of-the-art performance compose many transformations to give flexibility to the learned density (Kingma & Dhariwal, 2018; Ho et al., 2019a). Again the model computational cost increases as the number of transformations increases.

To examine how to binarize hierarchical VAEs and flow models successfully, we take two models which have demonstrated excellent modelling performance - the ResNet VAE (Kingma et al., 2016) and the Flow++ model (Ho et al., 2019a) - and implement the majority of each model with binary neural networks. Using binary weights and activations reduces the computational cost, but also decreases the representational capability of the model. Therefore our aim is to strike a balance between reducing the computational cost and maintaining good modelling performance. We show that it is possible to decrease the model size drastically, and allow for significant speed ups in run time, with only a minor impact on the achieved loss value. We make the following key contributions:

- We propose an efficient binary adaptation of weight normalization, a reparameterization technique often used in deep generative models to accelerate convergence. *Binary weight normalization* is the generative-modelling alternative to the usual batch normalization used in binary neural networks.
- We show that we can binarize the majority of weights and activations in deep hierarchical VAE and flow models, without significantly hurting performance. We demonstrate the corresponding binary architecture designs for both the ResNet VAE and the Flow++ model.
- We perform experiments on different levels of binarization, clearly demonstrating the trade-off between binarization and performance.

2 BACKGROUND

In this section we give background on the implementation and training of binary neural networks. We also describe the generative models that we implement with binary neural networks in detail.

2.1 BINARY NEURAL NETWORKS

In order to reduce the memory and computational requirements of neural networks, there has been recent research into how to effectively utilise networks which use binary-valued weights $w_{\mathbb{B}}$ and possibly also activations $\alpha_{\mathbb{B}}$ rather than the usual real-valued¹ weights and activations (Courbariaux et al., 2015; Courbariaux et al., 2016; Rastegari et al., 2016; McDonnell, 2018; Gu et al., 2018). In this work, we use the convention of binary values being in $\mathbb{B} := \{-1, 1\}$.

Motivation. The primary motivation for using binary neural networks is to decrease the memory and computational requirements of the model. Clearly binary weights require less memory to be stored: $32\times$ less than the usual 32-bit floating-point weights.

Binary neural networks also admit significant speed-ups. A reported $2\times$ speed-up can be achieved by a layer with binary weights and real-valued inputs (Rastegari et al., 2016). This can be made an additional $29\times$ faster if the inputs to the layer are also constrained to be binary (Rastegari et al., 2016). With both binary weights and inputs, linear operators such as convolutions can be implemented using the inexpensive XNOR and bit-count binary operations. A simple way to ensure binary inputs to a layer is to have a binary activation function before the layer (Courbariaux et al., 2016; Rastegari et al., 2016).

Optimization. Taking a trained model with real-valued weights and binarizing the weights has been shown to lead to significant worsening of performance (Alizadeh et al., 2019). So instead the binary weights are optimized. It is common to not optimize the binary weights directly, but instead optimize a set of underlying real-valued weights $w_{\mathbb{R}}$ which can then be binarized in some fashion for inference. In this paper we will adopt the convention of binarizing the underlying weights using the sign function (see Equation 2). We also use the sign function as the activation function when we use binary activations (see Equation 5, where $\alpha_{\mathbb{R}}$ are the real-valued pre-activations). We define the sign

¹We use real-valued throughout the paper to be synonymous with "implemented with floating-point precision".

function as:

$$\text{sign}(x) := \begin{cases} -1, & \text{if } x < 0 \\ 1, & \text{if } x \geq 0 \end{cases} \quad (1)$$

Since the derivative of the sign function is zero almost everywhere², the gradients of the underlying weights $w_{\mathbb{R}}$ and through binary activations are zero almost everywhere. This makes gradient-based optimization challenging. To overcome this issue, the straight-through estimator (STE) (Bengio et al., 2013) can be used. When computing the gradient of the loss \mathcal{L} , the STE replaces the gradient of the sign function (or other discrete output functions) with an approximate surrogate. A straightforward and widely used surrogate gradient is the identity function, which we use to calculate the gradients of the real-valued weights $w_{\mathbb{R}}$ (see Equation 3). It has been shown useful to cancel the gradients when their magnitude becomes too large (Courbariaux et al., 2015; Alizadeh et al., 2019). Therefore we use a clipped identity function for the gradients of the pre-activations (see Equation 6). This avoids saturating a binary activation. Lastly, the loss value only depends on the sign of the real-valued weights. Therefore, the values of the weights are generally clipped to be in $[-1, 1]$ after each gradient update (see Equation 4). This restricts the magnitude of the weights and thus makes it easier to flip the sign.

	Weights		Activations	
Forward pass:	$w_{\mathbb{B}} = \text{sign}(w_{\mathbb{R}})$	(2)	$\alpha_{\mathbb{B}} = \text{sign}(\alpha_{\mathbb{R}})$	(5)
Backward pass:	$\frac{\partial \mathcal{L}}{\partial w_{\mathbb{R}}} := \frac{\partial \mathcal{L}}{\partial w_{\mathbb{B}}}$	(3)	$\frac{\partial \mathcal{L}}{\partial \alpha_{\mathbb{R}}} := \frac{\partial \mathcal{L}}{\partial \alpha_{\mathbb{B}}} * 1_{ \alpha_{\mathbb{R}} \leq 1}$	(6)
After update:	$w_{\mathbb{R}} \leftarrow \max(-1, \min(1, w_{\mathbb{R}}))$	(4)	-	

2.2 DEEP GENERATIVE MODELS

Hierarchical VAEs. The variational autoencoder (Kingma & Welling, 2014; Rezende et al., 2014) is a latent variable model for observed data \mathbf{x} conditioned on unobserved latent variables \mathbf{z} . It consists of a generative model $p_{\theta}(\mathbf{x}, \mathbf{z})$ and an inference model $q_{\phi}(\mathbf{z}|\mathbf{x})$. The generative model can be decomposed into the prior on the latent variables $p_{\theta}(\mathbf{z})$ and the likelihood of our data given the latent variables $p_{\theta}(\mathbf{x}|\mathbf{z})$. The inference model is a variational approximation to the true posterior, since the true posterior is usually intractable in models of interest. Training is generally performed by maximization of the evidence lower bound (ELBO), a lower bound on the log-likelihood of the data:

$$\log p_{\theta}(\mathbf{x}) \geq \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\theta}(\mathbf{x}, \mathbf{z}) - \log q_{\phi}(\mathbf{z}|\mathbf{x})] \quad (7)$$

To give a more expressive model, the latent space can be structured into a hierarchy of latent variables $\mathbf{z}_{1:L}$. In the generative model each latent layer is conditioned on deeper latents $p_{\theta}(\mathbf{z}_i|\mathbf{z}_{i+1:L})$. A common problem with hierarchical VAEs is that the deeper latents can struggle to learn, often "collapsing" such that the layer posterior matches the prior: $q_{\phi}(\mathbf{z}_i|\mathbf{z}_{i+1:L}, \mathbf{x}) \approx p_{\theta}(\mathbf{z}_i|\mathbf{z}_{i+1:L})$ ³. One method to help prevent posterior collapse is to use skip connections between latent layers (Kingma et al., 2016; Maaløe et al., 2019), turning the layers into residual layers (He et al., 2016).

We focus on the ResNet VAE (RVAE) model (Kingma et al., 2016). In this model, both the generative and inference model structure their layers as residual layers. The ResNet VAE uses a bi-directional inference structure with both a bottom-up and top-down residual channel. This is a similar structure to the BIVA model (Maaløe et al., 2019), which has demonstrated state-of-the-art results for a latent variable model. We give a more detailed description of the model in Appendix C.

Flow models. Flow models consist of a parameterized invertible transformation, $\mathbf{z} = \mathbf{f}_{\theta}(\mathbf{x})$, and a known density $p_{\mathbf{z}}(\mathbf{z})$ usually taken to be a unit normal distribution. Given observed data \mathbf{x} we obtain the objective for θ by applying a change-of-variables to the log-likelihood:

$$\log p_{\theta}(\mathbf{x}) = \log p_{\mathbf{z}}(\mathbf{f}_{\theta}(\mathbf{x})) + \log \left| \det \frac{d\mathbf{f}_{\theta}}{d\mathbf{x}} \right| \quad (8)$$

For training to be possible, it is required that computation of the Jacobian determinant $\det(d\mathbf{f}_{\theta}/d\mathbf{x})$ is tractable. We therefore aim to specify of flow model \mathbf{f}_{θ} which is sufficiently flexible to model the

²Apart from at 0, where it is non-differentiable.

³We have assumed here that the inference model is factored "top-down".

data distribution well, whilst also being invertible and having a tractable Jacobian determinant. One common approach is to construct \mathbf{f}_θ as a composition of many simpler functions: $\mathbf{f}_\theta = \mathbf{f}_1 \circ \mathbf{f}_2 \circ \dots \circ \mathbf{f}_L$, with each \mathbf{f}_i invertible and with tractable Jacobian. So the objective becomes:

$$\log p_\theta(\mathbf{x}) = \log p_z(\mathbf{f}_\theta(\mathbf{x})) + \sum_{i=1}^L \log \left| \det \frac{d\mathbf{f}_i}{d\mathbf{f}_{i-1}} \right| \quad (9)$$

There are many approaches to construct the \mathbf{f}_i layers (Dinh et al., 2014; Rezende & Mohamed, 2015; Dinh et al., 2017; Kingma & Dhariwal, 2018; Ho et al., 2019a). In this work we will focus on the Flow++ model (Ho et al., 2019a), which has state-of-the-art results for flow models. In the Flow++ model, the \mathbf{f}_i are coupling layers which partition the input into \mathbf{x}_1 and \mathbf{x}_2 , then transform only \mathbf{x}_2 :

$$\mathbf{f}_i(\mathbf{x}_1) = \mathbf{x}_1, \quad \mathbf{f}_i(\mathbf{x}_2) = \sigma^{-1}(\text{MixLogCDF}(\mathbf{x}_2; \mathbf{t}(\mathbf{x}_1))) \cdot \exp(\mathbf{a}(\mathbf{x}_1)) + \mathbf{b}(\mathbf{x}_1) \quad (10)$$

Where MixLogCDF is the CDF for a mixture of logistic distributions. This is an iteration on the affine coupling layer (Dinh et al., 2014; 2017). Note that keeping part of the input fixed ensures that the layer is invertible. To ensure that all dimensions are transformed in the composition, adjacent coupling layers will keep different parts of the input fixed, often using an alternating checkerboard or stripe pattern to choose the fixed dimensions (Dinh et al., 2017). The majority of parameters in this flow model come from the functions \mathbf{t} , \mathbf{a} and \mathbf{b} in the coupling layer, and in Flow++ these are parameterized as stacks of convolutional residual layers. In this work we will focus on how to binarize these functions whilst maintaining good modelling performance. We give a more detailed description of the full flow model in Appendix D.

3 BINARIZING DEEP GENERATIVE MODELS

In this section we first introduce a technique to effectively binarize weight normalized layers, which are used extensively in deep generative model architectures. Afterwards, we elaborate on which components of the models we can binarize without significantly hurting performance.

3.1 NORMALIZATION

It is important to apply some kind of normalization after a binary layer. Binary weights are often large in magnitude relative to the usual real-valued weights, and can result in large outputs which can destabilize training. Previous binary neural network implementations have largely used batch normalization, which can be executed efficiently using a shift-based implementation (Courbariaux et al., 2016).

However, it is common in generative modelling to use weight normalization (Salimans & Kingma, 2016) instead of batch normalization. For example, it is used in the Flow++ (Ho et al., 2019a) and state-of-the-art hierarchical VAE models (Kingma et al., 2016; Maaløe et al., 2019). Weight normalization factors a vector of weights $\mathbf{w}_\mathbb{R}$ into a vector of the same dimension $\mathbf{v}_\mathbb{R}$ and a magnitude g , both of which are learned. The weight vector is then expressed as:

$$\mathbf{w}_\mathbb{R} = \mathbf{v}_\mathbb{R} \cdot \frac{g}{\|\mathbf{v}_\mathbb{R}\|} \quad (11)$$

Where $\|\cdot\|$ denotes the Euclidean norm. This implies that the norm of $\mathbf{w}_\mathbb{R}$ is g .

Now suppose we wish to binarize the parameters of a weight normalized layer. We are only able to binarize $\mathbf{v}_\mathbb{R}$, since binarizing the magnitude g and bias b could result in large outputs of the layer. However, g and b do not add significant compute or memory requirements, as they are applied elementwise and are much smaller than the binary weight vector.

Let $\mathbf{v}_\mathbb{B} = \text{sign}(\mathbf{v}_\mathbb{R})$ be a binarized weight vector of dimension n . Since every element of $\mathbf{v}_\mathbb{B}$ is one of ± 1 , we know that $\|\mathbf{v}_\mathbb{B}\| = \sqrt{n}$ ⁴. We then have:

$$\mathbf{w}_\mathbb{R} = \mathbf{v}_\mathbb{B} \cdot \frac{g}{\sqrt{n}} \quad (12)$$

⁴ $\|\mathbf{v}_\mathbb{B}\| = \sqrt{\sum_i (v_{\mathbb{B},i})^2} = \sqrt{\sum_i 1} = \sqrt{n}$

We refer to this as *binary weight normalization*, or BWN. Importantly, this is faster to compute than the usual weight normalization (Equation 11), since we do not have to calculate the norm of $\mathbf{v}_{\mathbb{B}}$. The binary weight normalization requires only $O(1)$ FLOPs to calculate the scaling for $\mathbf{v}_{\mathbb{B}}$, whereas the regular weight normalization requires $O(n)$ FLOPs to calculate the scaling for $\mathbf{v}_{\mathbb{R}}$. For a model of millions of parameters, this can be a significant speed-up. Binary weight normalization also has a more straightforward backward pass, since we do not need to take gradients of the $1/||\mathbf{v}||$ term.

Furthermore, convolutions \mathcal{F} and other linear transformations can be implemented using cheap binary operations when using binary weights, $\mathbf{w}_{\mathbb{B}}$, as discussed in Section 2.1⁵. However, after applying binary weight normalization, the weight vector is real-valued, $\mathbf{w}_{\mathbb{R}}$. Fortunately, since a convolution is a linear transformation, we can apply the normalization factor $\alpha = g/\sqrt{n}$ either before or after applying the convolution to input \mathbf{x} .

$$\mathcal{F}(\mathbf{x}, \mathbf{v}_{\mathbb{B}} \cdot \alpha) = \mathcal{F}(\mathbf{x}, \mathbf{v}_{\mathbb{B}}) \cdot \alpha \quad (13)$$

So if we wish to utilize fast binary operations for the binary convolution layer, we need to apply binary weight normalization *after* the convolution. This means that the weights are binary for the convolution operation itself. This couples the convolution operation and the weight normalization, and we refer to the overall layer as a binary weight normalized convolution, or BWN convolution. Note that the above process applies equally well to other linear transformations. We initialize BWN layers in the same manner as regular weight normalization, but give a more thorough description of alternatives in Appendix E.

3.2 BINARIZING RESIDUAL LAYERS

We aim to binarize deep generative models, in which it is common to utilize residual layers extensively. Residual layers are functions with skip connections:

$$\mathbf{g}_{\text{res}}(\mathbf{x}) = \mathcal{T}(\mathbf{x}) + \mathbf{x} \quad (14)$$

Indeed, the models we target in this work, the ResNet VAE and Flow++ models, have the majority of their parameters within residual layers. Therefore they are natural candidates for binarization, since binarizing them would result in a large decrease in the computational cost of the model. To binarize them we implement $\mathcal{T}(\mathbf{x})$ in Equation 14 using binary weights and possibly activations.

The motivation for using residual layers is that they can be used to add more representative capability to a model without suffering from the *degradation problem* (He et al., 2016). That is, residual layers can easily learn the identity function by driving the weights to zero. So, if sufficient care is taken with initialization and optimization, adding residual layers to the model should not degrade performance, helping to precondition the problem.

Degradation of performance is of particular concern when using binary layers. Binary weights and activations are both less expressive than their real-valued counterparts, and more difficult to optimize. These disadvantages of binary layers are more pronounced for generative modelling than for classification. Generative models need to be very expressive, since we wish to model complex data such as images. Optimization can also be difficult, since the likelihood of a data point is highly sensitive to the distribution statistics output by the model, and can easily diverge. This provides an additional justification for binarizing only the residual layers of a generative model. By restricting binarization to the residual layers, it decreases the chance that using binary layers harms performance.

Crucially, if we were to use a residual binary layer *without* weight normalization, then the layer would not be able to learn the identity function, as the binary weights cannot be set to zero. This would remove the primary motivation to use binary residual layers. In contrast, using a binary weight normalized layer in the residual layer, the gain g and bias b can be set to zero to achieve the identity function. As such, we binarize the ResNet VAE and Flow++ models by implementing the residual layers using BWN layers.

⁵This applies when the inputs are real-valued or binary, but the speed-ups are far greater for binary inputs

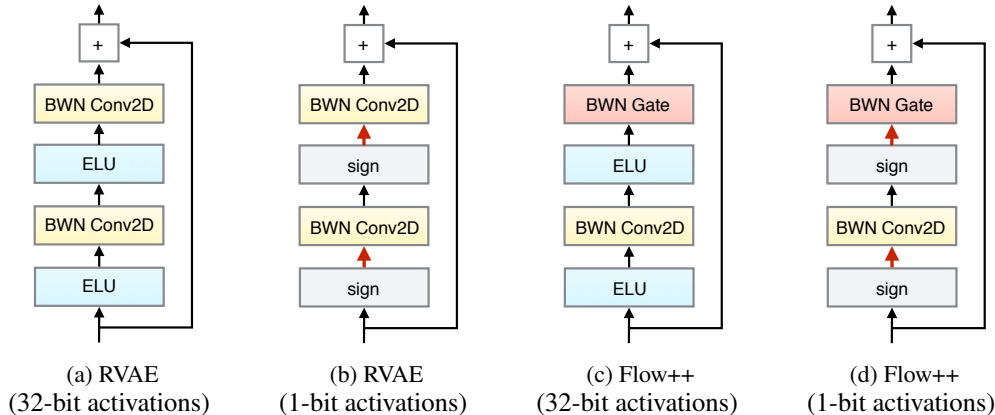


Figure 1: The residual blocks used in the binarized ResNet VAE and Flow++ models, using both binary and floating-point activations. The BWN Gate layer is a binary weight normalized 1×1 convolution followed by a gated linear unit. We display the binary valued tensors with thick red arrows.

4 DEEP GENERATIVE MODELS WITH BINARY WEIGHTS

We now describe the binarized versions of the ResNet VAE and Flow++ model, using the techniques and considerations from Section 3. Note that, for both the ResNet VAE and Flow++ models, we still retain a valid probabilistic model after binarizing the weights. In both cases, the neural networks are simply used to output distribution parameters, which define a normalized density for any set of parameter values.

ResNet VAE. As per Section 3.2, we wish to binarize the residual layers of the ResNet VAE. The residual layers are constructed as convolutional residual blocks, consisting of two 3×3 convolutions and non-linearities, with a skip connection. This is shown in Figure 1(a)-(b). To binarize the block, we change the convolutions to BWN convolutions, as described in Section 3.1. We can either use real-valued activations or binary activations. Binary activations allow the network to be executed much faster, but are less expressive. We use the ELU function as the real-valued activation, and the sign function as the binary activation.

Flow++. As with the ResNet VAE, in the Flow++ model the residual layers are structured as stacks of convolutional residual blocks. To binarize the residual blocks, we change both the 3×3 convolution and the gated 1×1 convolution in the residual block to be BWN convolutions. The residual block design is shown in Figure 1(c)-(d). We have the option of using real-valued or binary activations.

5 EXPERIMENTS

We run experiments with the ResNet VAE and the Flow++ model, to demonstrate the effect of binarizing the models. We train and evaluate on the CIFAR and ImageNet (32×32) datasets. For both models we use the Adam optimizer (Kingma & Ba, 2015), which has been demonstrated to be effective in training binary neural networks (Alizadeh et al., 2019).

For the ResNet VAE, we decrease the number of latent variables per latent layer and increase the width of the residual channels, as compared to the original implementation. We found that increasing the ResNet blocks in the first latent layer slightly increased modelling performance. Furthermore, we chose not to model the posterior using IAF layers Kingma et al. (2016), since we want to keep the model class as general as possible.

For the Flow++ model, we decrease the number of components in the mixture of logistics for each coupling layer and increase the width of the residual channels, as compared to the original implementation. For simplicity, we also remove the attention mechanism from the model, since the ablations the authors performed showed that this had only a small effect on the model performance.

Note that we do not use any techniques to try and boost the test performance of our models, such as importance sampling or using weighted averages of the model parameters. These are often used in

generative modelling, but since we are trying to establish the relative performance of models with various degrees of binarization, we believe that these techniques are irrelevant.

5.1 DENSITY MODELLING

We display results in Table 1. We can see that the models with binary weights and real-valued activations perform only slightly worse than those with real-valued weights, for both the ResNet VAE and the Flow++ models. For the models with binary weights, we observe better performance when using real-valued activations than with the binary activations. These results are as expected given that binary values are by definition less expressive than real values. All models with binary weights perform better than a baseline model with the residual layers set to the identity, indicating that the binary layers do learn. We display samples from the binarized models in Appendix A.

Importantly, we see that the model size is significantly smaller when using binary weights - 94% smaller for the ResNet VAE and 90% smaller for the Flow++ model.

These results demonstrate the fundamental trade-off that using binary layers in generative models allows. By using binary weights the size of the model can be drastically decreased, but there is a slight degradation in modelling performance. The model can then be made much faster by using binary activations as well as weights, but this decreases performance further.

5.2 INCREASING THE RESIDUAL CHANNELS

Binary models are less expensive in terms of memory and compute. This raises the question of whether binary models could be made *larger* in parameter count than the model with real-valued weights, with the aim of trying to improve performance for a fixed computational budget. We examine this by increasing the number of channels in the residual layers (from 256 to 336) of the ResNet VAE. This increases the number of binary weights by approximately 40 million, but leaves the number of real-valued weights roughly constant⁶. The results are shown in Table 1 and Figure 3(c). We can see the increase the binary parameter count does have a noticeable improvement in performance. The model size increases from 13 MB to 20 MB, which is still an order of magnitude smaller than the model with real-valued weights (255 MB). It is an open question as to how much performance could be improved by increasing the size of the binary layers even further. The barrier to this approach currently is training, since we need to maintain and optimize a set of real-valued weights during training. These get prohibitively large as we increase the model size significantly.

5.3 ABLATIONS

We perform ablations to verify our hypothesis from Section 3.2 that we should only binarize the residual layers of the generative models. We attempt to binarize all layers in the ResNet VAE using BWN layers, using both binary and real-valued activations. The results are shown in Figure 3(d). As expected, the loss values attained are significantly worse than when binarizing only the residual layers. We also perform an ablation comparing the performance of BWN against batch normalization in Appendix B, demonstrating the advantages of using BWN.

6 DISCUSSION AND FUTURE WORK

We have demonstrated that is possible to drastically reduce model size and compute requirements for the ResNet VAE and Flow++ models, whilst maintaining good modelling performance. We chose these models because of their demonstrated modelling power, but the methods we used to binarize them are readily applicable to other hierarchical VAEs or flow models. We believe this is a useful result for any real-world application of these generative models. such as learned lossless compression (Townsend et al., 2019; Kingma et al., 2019; Townsend et al., 2020; Ho et al., 2019b; Hoogeboom et al., 2019), which could be made practical with these reduced memory and compute benefits.

A key technical challenge that needs to be overcome for binary neural networks as a whole is the availability of implementations of the fast binary linear operations such as binary convolutions.

⁶There will be a slight increase, since we use real-valued weights to map to and from the residual channels.

Table 1: Results for binarized ResNet VAE and Flow++ model on CIFAR and ImageNet (32×32) test sets. Loss values are reported in bits per dimension. We give the percentage of the model parameters that are binary and the overall size of the model parameters. The weights and activations refer to those within the residual layers of the model, which are the targets for binarization.

	Precision		Modelling loss		# Parameters	% Binary	Memory cost
	Weights	Activations	CIFAR	ImageNet (32×32)			
ResNet VAE	32-bit	32-bit	3.45	4.25	56M	0%	255 MB
	1-bit	32-bit	3.60	4.47	56M	97.1%	13 MB
	1-bit	1-bit	3.73	4.58	56M	97.1%	13 MB

<i>increased width</i>	1-bit	32-bit	3.56	-	96M	97.7%	20 MB
	1-bit	1-bit	3.68	-	96M	97.7%	20 MB

<i>no residual</i>	N.A.	N.A.	3.78	-	1.6M	0%	6 MB

Flow++	32-bit	32-bit	3.21	4.05	34M	0%	129 MB
	1-bit	32-bit	3.29	4.18	34M	90.1%	14 MB
	1-bit	1-bit	3.43	4.30	34M	90.1%	14 MB

<i>no residual</i>	N.A.	N.A.	3.54	-	2.2M	0%	9 MB

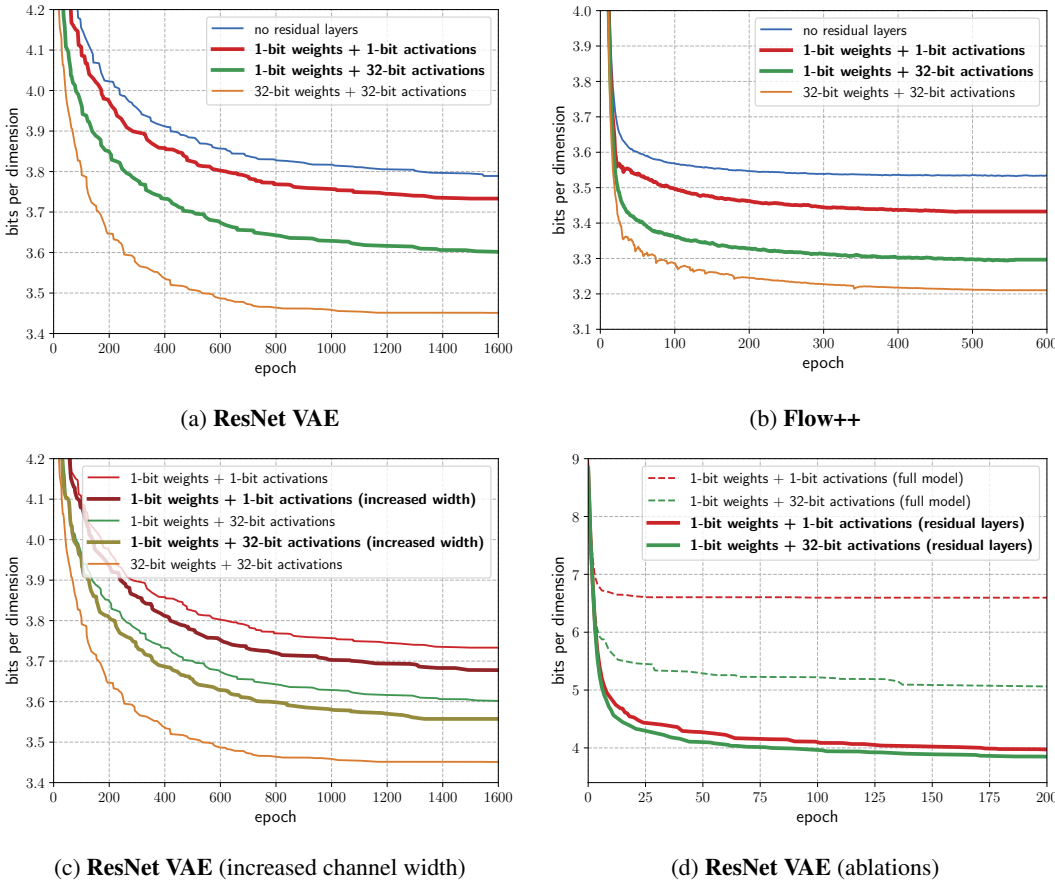


Figure 2: Test loss values during training of the ResNet VAE and Flow++ models on the CIFAR dataset. Subfigures (a) and (b): models with binary weights and either binary or real-valued activations. Compared to the model with real-valued weights and activations, and a baseline with the residual layers set to the identity. Subfigures (c) and (d): the effect of increasing the width of the residual channels, and ablations.

Proof-of-concept implementations of these binary kernels have been developed (Courbariaux et al., 2016; Rastegari et al., 2016; Pedersoli et al., 2018). However, there is no implementation that is readily usable as a substitute for the existing kernels in frameworks such as PyTorch and Tensorflow.

7 CONCLUSION

We have shown that it is possible to implement state-of-the-art deep generative models using binary neural networks. We proposed using a fast binary weight normalization procedure, and shown that it is necessary to binarize only the residual layers of the model to maintain modelling performance. We demonstrated this by binarizing two state-of-the-art models, the ResNet VAE and the Flow++ model, reducing the computational cost massively. We hope this insight into the possible trade-off between modelling performance and computational cost will stimulate further research into the efficiency of deep generative models.