
A Generative Model for Sampling High-Performance and Diverse Weights for Neural Networks

Lior Deutsch^{*1} Erik Nijkamp^{*2} Yu Yang²

Abstract

Recent work on mode connectivity in the loss landscape of deep neural networks has demonstrated that the locus of (sub-)optimal weight vectors lies on continuous paths. In this work, we train a neural network that serves as a hypernetwork, mapping a latent vector into high-performance (low-loss) weight vectors, generalizing recent findings of mode connectivity to higher dimensional manifolds. We formulate the training objective as a compromise between accuracy and diversity, where the diversity takes into account trivial symmetry transformations of the target network. We demonstrate how to reduce the number of parameters in the hypernetwork by parameter sharing. Once learned, the hypernetwork allows for a computationally efficient, ancestral sampling of neural network weights, which we recruit to form large ensembles. The improvement in classification accuracy obtained by this ensembling indicates that the generated manifold extends in dimensions other than directions implied by trivial symmetries. For computational efficiency, we distill an ensemble into a single classifier while retaining generalization.

1. Introduction

In recent years, generative methods such as Variational Autoencoders (VAEs) (Kingma & Welling, 2013) and Generative Adversarial Networks (GANs) (Goodfellow et al., 2014a) have shown impressive results in generating samples from complex and high-dimensional distributions. The training of these generative models is data-driven, and an essential requirement is the existence of a rich enough data set, which faithfully represents the underlying probability

distribution. A trained VAE decoder or GAN generator is a neural network which is operated by feeding it with an input (usually random), and it outputs an array of numbers, which represent, for example, pixel intensities of an image.

It is natural to suggest this idea be extended from images to neural networks, such that a trained generator would output numbers which represent weights for a neural network with a fixed target architecture and a fixed task, such as classifying a specific type of data. By “weights” we are referring to any trainable parameter of a neural network, including bias parameter. After all, just like images, a neural network is a structured array of numbers. Indeed, this was carried out in recent works (Krueger et al., 2017; Louizos & Welling, 2017). The term *hypernetwork* was coined (Ha et al., 2016; Krueger et al., 2017) for a neural network that acts like such a generator. Applications include ensemble creation and Bayesian inference. Hypernetworks can also serve as a tool for a researcher to explore the loss function surface.

A hypernetwork cannot be trained in a data-driven manner in the same sense as the aforementioned generative methods, since there is no rich data set of neural networks for each given target architecture and each task. Thus, there is no underlying probability distribution of neural networks. We, therefore, take a different approach, where we decide upon useful properties that such a distribution would have, and express them as loss function terms. The two useful properties are **accuracy** and **diversity**. The former implies that the generated neural networks achieve high accuracy in performing their intended tasks. The diversity property means that the hypernetwork could generate a big number of essentially different networks. Two networks are considered *essentially different* if their weights differ by more than just trivial symmetry transformations.

Previous work (Krueger et al., 2017; Louizos & Welling, 2017) are set up in a Bayesian context, require a probabilistic interpretation of the target neural network’s output, and harness Variational Inference (VI) (Hinton & Van Camp, 1993; Graves, 2011; Kingma, 2017) to approximate the posterior of the weights given the data and sample from it. Our method improves upon some aspects of this approach. First, we do not require a probabilistic interpretation of the target network, which allows us to apply our method

^{*}Equal contribution ¹Independent Researcher
²Department of Statistics, UCLA. Correspondence to: Lior Deutsch <sliorde@gmail.com>, Erik Nijkamp <erik.nijkamp@gmail.com>, Yu Yang <yy19970901@ucla.edu>.




Figure 1. The hypernetwork on a toy problem in two dimensions. (a) Points generated by hypernetwork, where the background shades are the values of the Gaussian mixture. The curve (composed of points) is the hypernetwork’s output for 400 uniformly spaced points z in the range $[-1, 1]$. The points were colored so that they can be easily matched with their z values in the other graphs in this figure. (b) The values of the Gaussian mixture (solid curve, left axis) along the path, and the distance of each point to its nearest neighbor (dashed curve, right axis).

to a broader range of neural network types. Second, our loss function has an explicit hyperparameter that can be tuned to balance between accuracy and diversity. Lastly, our formalism allows different forms for the diversity loss term, while VI forces the use of entropy. Our formalism has VI as a private case, for specific choices of the hyperparameter and the diversity loss term.

The main contributions of this work:

- We provide a simple template for hypernetwork loss functions, which have VI as a private case. The loss function is applicable to any type of neural network task - not only to networks with a probabilistic interpretation.
- We show how to make the measure of the diversity of the generated networks more meaningful by taking symmetry transformations into account.
- We describe a parameter sharing architecture which reduces the size of the hypernetwork.
- We demonstrate for the first time (to our knowledge) a hypernetwork that can generate all the weights of a deep neural network, in such a way that all weights are statistically dependent. We show that the set of generated networks lies on a highly non-trivial manifold in weights space.
- We distill an ensemble into a single neural network by stochastic approximation to improve computational

efficiency while retaining the accuracy of the ensemble.

2. Related Work

Recent papers have discussed the existence of connected one-dimensional regions in the space of weights of a neural network, with low loss values (Freeman & Bruna, 2016; Draxler et al., 2018; Garipov et al., 2018). (Freeman & Bruna, 2016) show this for half-rectified networks with two layers. (Draxler et al., 2018) and (Garipov et al., 2018) construct such paths using Nudged Elastic Bands and quadratic Bézier curves, respectively. It is natural to ask if we can capture higher dimensional manifolds as well.

Bayesian Hypernetworks (BHNs) (Krueger et al., 2017) and multiplicative normalizing flows with Gaussian posterior (MNFGs) (Louizos & Welling, 2017) both transform a random input into the weights of a target neural network. These works formalize the problem in a Bayesian setting and use VI to get approximate samples from the posterior of the target network’s weights. A key ingredient used by both is a normalizing flow (NF) (Rezende & Mohamed, 2015), which serves as a flexible invertible function approximator. A NF requires that the input size is equal to the output size, and thus scales badly with the target networks’ size.

To reduce the number of parameters in the hypernetwork, BHNs and MNFGs employ reparameterizations of the weights. In BHN, the chosen reparameterization is weight normalization (Salimans & Kingma, 2016), and the NF produces samples only of the norms of the filters. The remaining degrees of freedom are trained to be constant (non-random). MNFG models the weights as a diagonal Gaussian when conditioned on the NF, with trainable means and variances. The NF acts as scaling factors on the means, one scale factor per filter (in the case of convolutional layers). The sources of randomness are thus the NF input and the per-weight Gaussian. The weights for different layers are generated independently. This limits the diversity of generated networks, since each layer needs to learn how to generate weights which give good accuracy without having any information about the other layers’ weights.

In contrast, in the current work, we don’t use an NF, relying instead on MLPs and convolutions to create a flexible distribution. This makes it possible to use a small vector as the latent representation of the primary network, while generating all of its weights, without the need to use a restrictive model (e.g. Gaussian). The downside of this approach is that the MLPs are not invertible, which might cause the output manifold to have a lower dimension than the input, and which makes it necessary to use an approximation for the entropy of the output.

There is more prior art on using auxiliary neural networks for the task of obtaining the weights of a target network.

Here we discussed the two most relevant to our work. In Appendix A we review additional papers. Common to most of these papers is that there is no attempt for the generated weights to be diverse.

3. Methods

3.1. Hypernetworks

Let $T(x; \theta) : X \times \Theta \rightarrow Y$ be the neural network whose weights we want to generate, where X is the input domain, Y is the output domain and Θ is the set of trainable weight vectors of the network. We refer to T as the *target network architecture*, or more shortly as the *target network*. Let $\mathcal{L}(\theta|p_{\text{data}})$ be the loss function associated with the target network, where $p_{\text{data}}(x, y)$, defined on $X \times Y$, is the data distribution. The standard practice for obtaining useful weights θ is to minimize \mathcal{L} using backpropagation, where the gradients of \mathcal{L} are estimated using batches of samples from the training data set. The outcome of this process is a single optimal vector of weights θ^* . The set of possible outcomes may be very large, due to the prevalence of local minima and flat regions in the loss function surface (Choromanska et al., 2015; Dauphin et al., 2014) and due to symmetry transformations, such as permutation between filters and scaling of weights, which keep the network output unchanged.

The approach offered here is different. Instead of obtaining θ^* by directly minimizing \mathcal{L} , we obtain θ^* as the output of a *hypernetwork*, which is a generator neural network $G(z; \varphi) : Z \times \Phi \rightarrow \Theta$, where Z is some input domain and Φ is the space of parameters (To reduce confusion, we refer to θ as “weights” and to φ as “parameters”; We occasionally use the notation $G(z)$ as a shorthand for $G(z; \varphi)$). We will draw the values for z from a simple probability distribution p_{noise} . We train G by minimizing a loss function $L(\varphi|p_{\text{noise}}, p_{\text{data}})$ that depends on the combined network $T(x; G(z; \varphi))$.

3.2. Loss Function

3.2.1. TWO COMPONENTS: ACCURACY AND DIVERSITY

To benefit from having a trained hypernetwork, it is not enough that $T(x; G(z; \varphi^*))$ yields low values for \mathcal{L} . It is also important that the hypernetwork can generate essentially different networks (we define this in Section 3.2.3). This forces us to let the hypernetwork $G(z)$ generate also sub-optimal weights. The reason is that $G(z)$ is a continuous function, and therefore if its image contains two distinct global optima, it must also contain a continuous path in weight space between them, and this path must also pass through sub-optimal weights (assuming there is no constant-loss path connecting the two optima) (Freeman & Bruna, 2016). Thus, we train φ^* by minimizing over a loss function




Figure 2. Histograms of accuracies of generated networks. (a) Our hypernetwork, $\lambda = 10^5$. (b) Our hypernetwork, $\lambda = 10^3$. (c) MNFG.

which includes two terms:

$$L(\varphi|p_{\text{noise}}, p_{\text{data}}) = \lambda L_{\text{accuracy}}(\varphi|p_{\text{noise}}, p_{\text{data}}) + L_{\text{diversity}}(\varphi|p_{\text{noise}}), \quad (1)$$

where L_{accuracy} depends on \mathcal{L} . An obvious choice is

$$L_{\text{accuracy}}(\varphi|p_{\text{noise}}, p_{\text{data}}) = \mathbb{E}_{z \sim p_{\text{noise}}} \mathcal{L}(G(z; \varphi) | p_{\text{data}}).$$

$L_{\text{diversity}}$ should ensure that there is high diversity in the results of G as a function of z , obviating the risk of undergoing “mode collapse”. $\lambda > 0$ is a hyperparameter that balances the two losses (Strictly speaking, λ should not be called a hyperparameter, since it is not a variable such as the number of layers or a regularization coefficient which should be tuned so as to minimize the validation loss. λ is an essential part of the validation loss function itself. However, we do not make this distinction in the following).

For the diversity term, we will use (the negative of) the entropy of the generated weights. Other possibilities include the variance, or diversity terms that are used in texture synthesis (Li et al., 2017) and feature visualization (Olah et al., 2017). During training, we estimate the entropy of a minibatch of generated weights using a modified form of the Kozachenko-Leonenko estimator (Kozachenko & Leonenko, 1987; Kraskov et al., 2004) (see Appendix E).

The loss function in equation (1) is heuristically simple to justify. However, it is illuminating to see how it can also be obtained in another way, as the relaxation of an optimistic choice for the distribution of $G(z)$. This is described in Appendix B.

3.2.2. TAKING SYMMETRIES INTO ACCOUNT

An increase in entropy may not always translate to an increase in diversity, since the diversity that we are interested in is that of *essentially different* outputs. Two weight vectors $\theta_1, \theta_2 \in \Theta$ are considered *essentially different* if there is no trivial symmetry transformation that transforms θ_1 into a close proximity of θ_2 , where the proximity is measured by some metric on Θ . Symmetry transformations are functions $S : \Theta \rightarrow \Theta$ such that $T(x; \theta) = T(x; S(\theta))$ for all $x \in X$ and all $\theta \in \Theta$. The trivial symmetry transformations include any composition of the following:

- **Scaling** - If the target network is a feed-forward convolutional network with piecewise-linear activations such as ReLU (Jarrett et al., 2009; Glorot et al., 2011) or leaky-ReLU (Maas et al., 2013), then scaling the weights of one filter by a positive factor, while unscaling the weights of the corresponding channel in all filters in the next layer by the same factor, keeps the output of the network unchanged.
- **Logits’ bias** - If a network produces logit values which are fed into a softmax layer, then adding the same number to all logit values does not change the values of the softmax probabilities.
- **Permutation** - Permuting the filters in a layer, while performing the same permutation on the channels of the filters in the next layer.

A hypernetwork that generates weight vectors that differ only by a trivial symmetry transformation should not score high on diversity, even though it may have high entropy. To deal with this problem, we use *gauge fixing* (a term borrowed from theoretical physics), which breaks the symmetry by choosing only one representative from each equivalence class of trivial symmetry transformations. This can be realized by a function $\mathcal{G} : \Theta \rightarrow \Theta$ which transforms any weight vector to its equivalence class representative: $\mathcal{G}(\theta_1) = \mathcal{G}(\theta_2)$ if and only if θ_1 and θ_2 are related by a trivial symmetry transformation. Therefore, we use the following form for the entropy in the diversity term:

$$L_{\text{diversity}}(\varphi|p_{\text{noise}}) = -\mathbb{H}_{z \sim p_{\text{noise}}}[\mathcal{G}(G(z; \varphi))], \quad (2)$$

where \mathbb{H} is the entropy of its argument. We choose \mathcal{G} to break the symmetries in scaling and in logits’ bias. The former is broken by requiring $\sum_k (\theta_{l,i}[k])^2 = n_{l,i}$, where $\theta_{l,i}[k]$ is the k ’th element of the i ’th filter of the l ’th layer, and $n_{l,i}$ is the number of elements in the filter, including the bias term. The sum is over all elements of the filter. This constraint is applied to the layers $1 \leq l \leq m - 1$ where m is the number of layers, and for all filters i in the layer. It is imposed on all but the last layer, since this is the freedom




Figure 3. Examples of samples of filter slices. Each of the four figures contains 25 samples for one specific filter slice. (a) first layer. (b) first layer (different filter). (c) second layer. (d) third layer.

we have under the scaling symmetry transformation. To see this, we can take an arbitrary network, and repeat the following process, starting with $l = 1$ and then incrementing l by 1: we scale the filters of the l ’th layer to obey the constraint, and unscale the filters of the $l + 1$ layer accordingly to keep the output unchanged. This process cannot be performed when $l = m$, since for the last layer there is no next layer to do the unscaling on. The logits’ bias symmetry is broken by requiring $\sum_i \theta_{m,i}[\text{bias}] = 0$, where the summand is the bias term of the filter. Incorporating permutation symmetry is also possible, for example by lexicographically sorting filters in a layer. However, we decide to ignore this symmetry, due to implementation constraints. Note that the permutations form a discrete group, and it is harder for a continuous generator to fail by generating discrete transformations.

3.2.3. DIVERSITY MEASURES WHICH ARE NOT AFFECTED BY SYMMETRIES

An alternative approach for the diversity loss term is to measure diversity in the *results* of the generated networks, as opposed to diversity in their weights (this means that $L_{\text{diversity}}$ will depend on p_{data} and well as p_{noise}). This approach will inherently assign low diversity when symmetries are generated. Here are two ways to do this:

- Instead of calculating the entropy of the weight vector, calculate the entropy of the concatenation of the outputs of the target network for various inputs. This can be written informally as

$$\mathbb{H}_{z \sim p_{\text{noise}}}[\text{concat}_{x \sim p_{\text{data}}} T(x; G(z; \varphi))]. \quad (3)$$

- We expect high diversity to yield an increase in prediction accuracy for an ensemble of generated networks, compared to the individual networks. This can be translated into a loss function by aggregating the results of different generated networks (e.g by averaging), and applying the target network’s loss criteria on this aggregation.

3.2.4. THE CASE OF CLASSIFICATION, AND THE RELATION TO VARIATIONAL INFERENCE

In Section 4, we describe experiments where the target network performs classification. In this case, the outputs of the target network are probability distributions over the finite set of classes. Thus, $T(x; \theta)_i$ is the probability for the i ’th class for input x . We can rewrite this as $p(i|x; \theta)$. Typically, the loss function is taken to be the negative mean log likelihood:

$$\mathcal{L}(\theta|p_{\text{data}}) = -\mathbb{E}_{(x,y) \sim p_{\text{data}}} \log p(y|x; \theta), \quad (4)$$

where for simplicity we assumed here that the data set contains only deterministic distributions y , i.e. y is a one-hot encoding of a class, and we can identify between y and its class.

Combining equations (1 - 4) while ignoring the gauge fixing function and setting $\lambda = n$, where n is the size of the data set, we obtain:

$$\begin{aligned} L(\varphi|p_{\text{noise}}, p_{\text{data}}) = & \\ & - \mathbb{E}_{z \sim p_{\text{noise}}, (x,y) \sim p_{\text{data}}} n \log p(y|x; G(z; \varphi)) \\ & - \mathbb{H}_{z \sim p_{\text{noise}}} [G(z; \varphi)]. \end{aligned} \quad (5)$$

This is very similar to the VI objective (Krueger et al., 2017; Louizos & Welling, 2017), as we describe in Appendix C. However, equation (1) is more general in that it does not require the outputs of the target network to be probability distributions, it allows us to use diversity terms other than entropy, which could include gauge fixing, and it incorporates the hyperparameter λ that enables us to control the balance between accuracy and diversity (However, we believe that the appearance of λ is consistent with VI, as we explain in Appendix D).

3.3. Architecture

We assume that the target network T is convolutional with m layers, where within each layer the filters have the same size. $\theta_{l,i}$ are the weights for the i ’th filter of the l ’th layer. If the hypernetwork G were a fully connected multilayer perceptron (MLP), it would scale badly with the dimension of the weights of T . We therefore utilize parameter sharing in G by giving it a convolutional structure, see Fig. 4.




Figure 4. Architecture block diagram of the hypernetwork with latent variable z , extractor E , codes c , weight-generators W , and generated weights θ .




Figure 5. Examples of samples of filter slices for MNFG. Each of the three figures contains 25 samples for one specific filter slice. (a) first layer. (b) first layer (different filter). (c) second layer.

The input z for the hypernetwork is drawn from a prior distribution such as Uniform or Gaussian and is fed into a fully-connected sub-network E , which we call the *extractor*, whose output is a set of *codes* $c_{l,i}$, where $l = 1, \dots, m$. The code $c_{l,i}$ is a latent representation of $\theta_{l,i}$. The code $c_{l,i}$ is then fed into the *weight generator* W_l , which is another fully connected network which generates the weights $\theta_{l,i}$ for the target network. We emphasize that the same weight generator is re-used for all filters in a certain layer of T , thus reducing the overall number of parameters in the hypernetwork. We consider the weights at the input of a fully connected neuron as a filter, whose receptive field is the entire previous layer.

W_l can be seen as a convolutional non-linear filter with a receptive field that is a single code for the layer l . To allow

flexible use of high-level features, we treat all weights in the last layer of T as one “filter”, thus using W_m only once.

To ease the optimization of the hypernetwork G , we may introduce residual connections and batch-normalization in extractor network E and weight generators W_l .

4. Experiments

All code used to run the experiments can be found online¹. Please refer to the code for a full specification of hyperparameter values and implementation details.

4.1. Toy Problem

We start with a toy problem which is easy to visualize. Instead of generating weights for a neural network, we generate a two-dimensional vector. The goal is for the generated vectors have high values on a specified Gaussian mixture, and also to obtain high diversity. We take the input z to the hypernetwork to have dimension 1, to demonstrate the case where the hypernetworks output manifold has a lower dimensionality than the weight space. This problem does not have symmetries, so we do not include gauge fixing in the loss. Instead, we use a conventional ℓ_2 regularization loss. The hypernetwork is taken as an MLP whose hidden layers have sizes 30, 10 and 10. The final distribution learned by the hypernetwork is displayed in Fig. 1. We see that the one-dimensional distribution is supported on a path which passes through all peaks of the Gaussian mixture. Inevitably, the path must pass through regions with low values of the Gaussian mixture. However, in these regions, the density is lower.

4.2. MNIST

We take the target network architecture T to be a simple four layer convolutional network for classifying images of the MNIST data set (LeCun et al., 1998). The full target network specification is displayed in Appendix F. The total number of weights in the network is 20,018. This network can easily be trained to achieve an accuracy of over 99% on the validation set.

For the hypernetwork, we take the input vector z to be 300 dimensional, drawn from a uniform distribution. The extractor and weight generators have three layers each. The total number of parameters is 633,640. For a detailed specification, including training details, see Appendix G.

We compare our results to MNFG (Louizos & Welling, 2017). We used code that is available online², and modi-

¹<https://github.com/sliorde/generating-neural-networks-with-neural-networks>

²https://github.com/AMLab-Amsterdam/MNF_VBNN

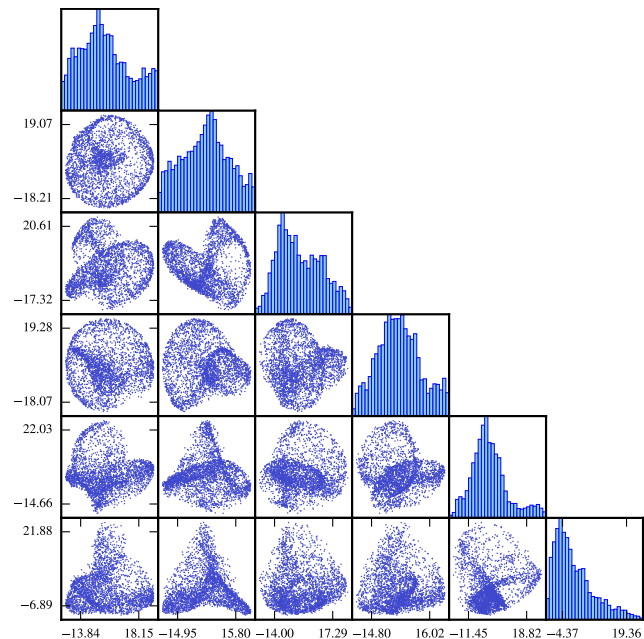


Figure 6. Scatter plots of the generated weights for a specific second layer filter, in PCA space. The (i, j) scatter plot has principal component i against principal component j .

fied it so that it generates the same target networks as our hypernetwork.

4.2.1. ACCURACY

The validation set accuracy of generated networks depends on the hyperparameter λ . Figs. 2a and 2b display histograms of the accuracies of the generated weights, for $\lambda = 10^5$ and $\lambda = 10^3$ respectively. We also show the corresponding histogram for MNFG in Fig. 2c. From now on, we will use $\lambda = 10^3$, which yields lower accuracies, but they are more comparable to the results of MNFG.

4.2.2. DIVERSITY

We explore the diversity in a few different ways. The histograms in Fig. 2 give an initial indication of diversity by showing that there is variance in the generated networks’ accuracies.

Visual Inspection. In Fig. 3 we show images of different samples of the generated filters. By visual inspection we see that different samples can result in different forms of filters. However, it is noticeable that there are some repeating patterns between samples. For comparison, we show corresponding images for MNFG in Fig. 5, where we gauged the filters generated by MNFG just as our own. We see that MNFG yields high diversity for most filters (e.g. Fig. 5a), but very low diversity for others (Fig. 5b), mainly in the first layer. We see this phenomena also in our generated filters

Dataset	Model	SGD			Hypernetwork			
		(a) 1 Sample	(b) 10 Samples		(c) w/o Entropy		(d) w/ Entropy	
			Mean	Majority	Mean	Majority	Mean	Majority
CIFAR-10	LeNet	77.36%	74.95%	80.60%	76.26%	76.47%	72.34%	78.21%
	ResNet-32	93.57%	91.03%	94.96%	88.20%	88.26%	83.61%	90.37%

Table 1. Test accuracies on the CIFAR-10 dataset for (a) SGD with decaying learning rate as reference and ensembles where (b) samples are drawn from a SGD dynamics, and, (c-d) samples drawn by ancestral sampling from the hypernetwork.




Figure 7. Accuracies along three paths, with end points z_1, z_2 sampled at random. The dashed lines are the direct paths, and the solid lines are the interpolated paths.

only to a lesser extent. We hypothesize that a possible mode of failure for a hypernetwork is when it concentrates much of its diversity in specific filters, while making sure that these filters get very small weights in the next layer, thereby effectively canceling these filters. Future work should consider this mode of failure in the diversity term of the loss function.

Scatter. One may wonder whether our method of weight generation is equivalent to trivially sampling from $\mathcal{N}(\theta_0, \Sigma)$, for some optimal weight vector θ_0 and constant Σ . To see that this is not the case, we view the scattering of the weight vectors using principle component analysis (PCA). This is shown in Fig. 10 and in Appendix J. We see that the hypernetwork learned to generate a distribution of weights on a non-trivial manifold, with prominent one-dimensional structures. Scatter graphs for MNFG are displayed in Appendix

K.

The construction of connected regions in weight space, with low accuracy loss values, was recently discussed in (Freeman & Bruna, 2016; Draxler et al., 2018; Garipov et al., 2018) for the case of one-dimensional regions. Here we see that this can be done also for higher dimensional manifolds.

Paths in Weight Space. For two given input vectors z_1 and z_2 , we define two paths which coincide in their endpoints: The direct path $\{G(z_1)t + G(z_2)(1-t) \mid t \in [0, 1]\}$, and the interpolated path $\{G(z_1t + z_2(1-t)) \mid t \in [0, 1]\}$. We expect high accuracy along the interpolated path. Whether the direct path has high accuracy depends on the nature of the generated manifold. For example, if the diversity is achieved only via random isotropic noise around a specific weight vector, the direct path would have high accuracy. In Fig. 1 we see that this is not the case. The analogous graph for MNFG, displayed in Appendix I, shows that for MNFG the direct path does give high accuracy.

Ensembles. We compare the accuracy of the generated networks with the accuracy of ensembles of generated networks. If the generated classifiers are sufficiently different, then combining them should yield a classifier with reduced variance (Friedman et al., 2001), and therefore lower error. We created 20 ensembles, each of size 200, and we take their majority vote as a classification rule. The result is that the average accuracy of the ensembles on the validation set was 99.14%, which is higher than typical results that we see in the histogram in Fig. 2. For MNFG, the same experiment gives an average ensemble accuracy of 99.28%.

Adversarial Examples. We show that using ensembles of generated networks can help reduce the sensitivity to adversarial examples (Szegedy et al., 2013; Goodfellow et al., 2014b). The experiment was conducted by following these steps: **1)** Use the hypernetwork to generate a weight vector θ . **2)** Sample a pair (x, y) of image and label from the validation set. **3)** Randomly pick a new label $y' \neq y$ to be the target class of the adversarial example. **4)** Use the fast gradient method (Goodfellow et al., 2014b) to generate adversarial examples. Do this for perturbation sizes ϵ in

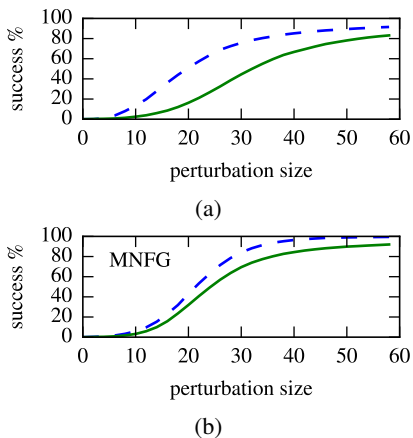


Figure 8. The success probability of adversarial examples created with the fast gradient method with a given perturbation size, against an ensemble of 100 generated networks. The dashed lines are the single classifier, and the solid lines are ensembles. (a) our hypernetwork. (b) MNFG.

the range 0 to 0.24, as fractions of the dynamic range of an image (8 bits of grayscale). **5** Test which values of ϵ yield adversarial examples that fool the classifier with weights θ . **6** Use the hypernetwork to generate an ensemble of 100 classifiers, and test for which values ϵ the adversarial examples created in step 4 fool the ensemble. We repeat this experiment over all images in the validation set. The results are shown in Fig. 000, together with results for MNFG. We see that the probability of success for an adversarial attack is reduced when using ensembles.

4.3. CIFAR-10

We empirically validate the diversity of the generated target networks by comparing various ensembling methods and the scalability of the hypernetwork by introducing skip connections. As such, we train the hypernetwork with LeNet (69, 656 parameters) and ResNet-32 (466, 906 parameters) on the CIFAR-10 (Krizhevsky et al., 2014) dataset. We refer Appendix H for the specifications of the hypernetworks.

4.3.1. ENSEMBLES

Table 1 depicts the maximum test accuracy for target networks trained by Stochastic Gradient Descent (SGD) and our hypernetwork. That is, (a) we train both models by SGD with decaying learning-rate to achieve close-to state-of-the-art-performance, (b) form an ensemble of 10 models by treating the SGD with constant learning-rate as a sampler (Mandt et al., 2017) from the weight posterior and report mean and majority vote accuracy, (c) an ensemble of 10 models sampled from the hypernetwork trained without entropy term, (d) and with entropy term (3) as ablation study.

Ablation. The result indicates that the hypernetwork can learn diverse weights and scale to larger target networks. For SGD ensembles of independent chains, as expected, the diversity in the weights leads to consistently highest test accuracy. For LeNets, the hypernetwork appears to learn a patch of the weight manifold when including the entropy term as objective and a sampled ensemble outperforms the accuracy of SGD on a single model by 0.85 on CIFAR-10. Including the entropy term decreases the mean accuracy while increasing the ensemble accuracy. For ResNet-32, the entropy term increases diversity in the weights, but the ensemble does not achieve state-of-the-art performance. We conjecture with further fine-tuning the performance on ResNet-32 can be improved.

4.4. Distillation

It is desirable to distill an ensemble of target networks into a single classifier for computational efficiency while retaining generalization. Polyak-Ruppert averaging (Polyak, 1990; Ruppert, 1988) is in the form of averaging iterates $\{\theta_t\}$ of a dynamics such that $\bar{\theta}_t = \frac{1}{t} \sum_{i=0}^{t-1} \theta_i$. We propose to substitute iterates of the trajectory by ancestral samples $\{\theta_i = G(z_i), z_i \sim \mathcal{N}(0, \sigma I)\}_{i=1}^n$. For LeNet on CIFAR-100, the ensemble accuracy of 45.42% drops to 45.41% under Polyak averaging. For ResNet-32 on CIFAR-10, the accuracy drops from 90.02% to 89.46%.

5. Discussion

In this work, we have shown how a hypernetwork can be trained to generate accurate and diverse weight vectors. Important directions of further inquiry are: Is there a certain gauge which yields better training? What are the performances of other diversity terms? How should diversity be evaluated? What are good methods for initializing parameters for the hypernetwork?

In answering the latter question, one should notice that popular methods (for example, (Glorot & Bengio, 2010; He et al., 2015)), use the fan_{in} and fan_{out} of a unit, but in the case of a weight generator sub-network in our architecture, we may want to take into account also the fan_{in} and fan_{out} of the generated filter.

It is critical to find architectures for hypernetworks which scale better with the size of the target network. Ideally, hypernetworks will have fewer parameters than their target networks, and therefore could be used as a compressed version of the target network.

In the future, applications of the proposed hypernetwork like transfer learning or generating sparse weight vectors for target networks should be explored.

References

- Github repository, https://github.com/AMLab-Amsterdam/MNF_VBNN.
- Github repository, <https://github.com/sliorde/generating-neural-networks-with-neural-networks>.
- Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M. W., Pfau, D., Schaul, T., and de Freitas, N. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, pp. 3981–3989, 2016.
- Bertinetto, L., Henriques, J. F., Valmadre, J., Torr, P., and Vedaldi, A. Learning feed-forward one-shot learners. In *Advances in Neural Information Processing Systems*, pp. 523–531, 2016.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pp. 1613–1622, 2015.
- Chang, O. and Lipson, H. Neural network quine. *arXiv preprint arXiv:1803.05859*, 2018.
- Charzyńska, A. and Gambin, A. Improvement of the k-nn entropy estimator with applications in systems biology. *Entropy*, 18(1):13, 2015.
- Choromanska, A., Henaff, M., Mathieu, M., Arous, G. B., and LeCun, Y. The loss surfaces of multilayer networks. In *Artificial Intelligence and Statistics*, pp. 192–204, 2015.
- Dauphin, Y. N., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., and Bengio, Y. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27*, pp. 2933–2941. Curran Associates, Inc., 2014.
- De Brabandere, B., Jia, X., Tuytelaars, T., and Van Gool, L. Dynamic filter networks. In *Neural Information Processing Systems (NIPS)*, 2016.
- Denil, M., Shakibi, B., Dinh, L., Ranzato, M. A., and de Freitas, N. Predicting parameters in deep learning. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 26*, pp. 2148–2156. Curran Associates, Inc., 2013.
- Draxler, F., Veschgini, K., Salmhofer, M., and Hamprecht, F. A. Essentially no barriers in neural network energy landscape. *arXiv preprint arXiv:1803.00885*, 2018.
- Freeman, C. D. and Bruna, J. Topology and geometry of half-rectified network optimization. *arXiv preprint arXiv:1611.01540*, 2016.
- Friedman, J., Hastie, T., and Tibshirani, R. *The elements of statistical learning*, volume 1. Springer series in statistics New York, 2001.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059, 2016.
- Garipov, T., Izmailov, P., Podoprikin, D., Vetrov, D. P., and Wilson, A. G. Loss surfaces, mode connectivity, and fast ensembling of dnns. *arXiv preprint arXiv:1802.10026*, 2018.
- Glorot, X. and Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 249–256, 2010.
- Glorot, X., Bordes, A., and Bengio, Y. Deep sparse rectifier neural networks. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 315–323, 2011.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27*, pp. 2672–2680. Curran Associates, Inc., 2014a.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014b.
- Graves, A. Practical variational inference for neural networks. In *Advances in Neural Information Processing Systems*, pp. 2348–2356, 2011.
- Ha, D., Dai, A., and Le, Q. V. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- He, K., Zhang, X., Ren, S., and Sun, J. 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.
- Hinton, G. E. and Van Camp, D. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pp. 5–13. ACM, 1993.

- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pp. 448–456, 2015.
- Jarrett, K., Kavukcuoglu, K., LeCun, Y., et al. What is the best multi-stage architecture for object recognition? In *Computer Vision, 2009 IEEE 12th International Conference on*, pp. 2146–2153. IEEE, 2009.
- Kingma, D. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. *Variational inference & deep learning: A new synthesis*. PhD thesis, University of Amsterdam, 2017.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kozachenko, L. and Leonenko, N. N. Sample estimate of the entropy of a random vector. *Problemy Peredachi Informatsii*, 23(2):9–16, 1987.
- Kraskov, A., Stögbauer, H., and Grassberger, P. Estimating mutual information. *Physical review E*, 69(6):066138, 2004.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Krizhevsky, A., Nair, V., and Hinton, G. The cifar-10 dataset. *online: <http://www.cs.toronto.edu/kriz/cifar.html>*, 2014.
- Krueger, D., Huang, C.-W., Islam, R., Turner, R., Lacoste, A., and Courville, A. Bayesian hypernetworks. *arXiv preprint arXiv:1710.04759*, 2017.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Li, K. and Malik, J. Learning to optimize. *arXiv preprint arXiv:1606.01885*, 2016.
- Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., and Yang, M.-H. Diversified texture synthesis with feed-forward networks. *arXiv preprint arXiv:1703.01664*, 2017.
- Lorraine, J. and Duvenaud, D. Stochastic hyperparameter optimization through hypernetworks. *arXiv preprint arXiv:1802.09419*, 2018.
- Louizos, C. and Welling, M. Multiplicative normalizing flows for variational bayesian neural networks. *arXiv preprint arXiv:1703.01961*, 2017.
- Maas, A. L., Hannun, A. Y., and Ng, A. Y. Rectifier nonlinearities improve neural network acoustic models. In *Proc. ICML*, volume 30, 2013.
- Mandt, S., Hoffman, M. D., and Blei, D. M. Stochastic gradient descent as approximate bayesian inference. *The Journal of Machine Learning Research*, 18(1):4873–4907, 2017.
- Marsh, C. Introduction to continuous entropy. 2013.
- Neal, R. M. *Bayesian learning for neural networks*. PhD thesis, University of Toronto, 1995.
- Olah, C., Mordvintsev, A., and Schubert, L. Feature visualization. *Distill*, 2017. doi: 10.23915/distill.00007. <https://distill.pub/2017/feature-visualization>.
- Polyak, B. T. New stochastic approximation type procedures. *Automat. i Telemekh*, 7(98-107):2, 1990.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- Ruppert, D. Efficient estimations from a slowly convergent robbins-monro process. Technical report, Cornell University Operations Research and Industrial Engineering, 1988.
- Salimans, T. and Kingma, D. P. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 901–909. Curran Associates, Inc., 2016.
- Schmidhuber, J. A ‘self-referential’ weight matrix. In *ICANN’93*, pp. 446–450. Springer, 1993.
- Schmidhuber, J. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Learning*, 4(1), 2008.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *Journal of machine learning research*, 15(1):1929–1958, 2014.
- Stanley, K. O., D’Ambrosio, D. B., and Gauci, J. A hypercube-based encoding for evolving large-scale neural networks. *Artificial life*, 15(2):185–212, 2009.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

Welling, M. and Teh, Y. W. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 681–688, 2011.

Zoph, B. and Le, Q. V. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.

A. Additional Related Work

There are many methods of using auxiliary neural networks for the task of obtaining the weights of a target network. An early example is fast weights (Schmidhuber, 2008), where the target network is trained together with an auxiliary memory controller(MC) network, whose goal is to drive changes in the weights of the target network. Although the overall architecture of this approach is quite similar to hypernetworks as presented here, the approach differs in two important aspects: **(a)** Fast weights are designed specifically as alternatives to recurrent networks for temporal sequence processing, and therefore the MC can never be decoupled from the target network, even after the MC has done a computation; **(b)** The input to the MC is just the input to the target network, and there is no attempt to generate diverse iid samples of the target network weights.

Under the paradigm of *learning to learn*, the works (Andrychowicz et al., 2016; Li & Malik, 2016) have trained auxiliary neural networks to act as optimizers of a target network. The optimizers apply all the updates to the target network’s weights during its training. As opposed to the method presented here, these trained optimizers can generate weights to the target network only by receiving a long sequence of batches of training examples. On the other hand, these optimizers can generalize to various loss functions and problem instances.

In (Bertinetto et al., 2016), the auxiliary network is used for one-shot learning: It receives as input a single training example, and produces weights for the target network. This is different from hypernetworks, where the input is a latent representation of the target network, and the generated weights are not seen as a generalization from a single training example. A related method is dynamic filter networks (De Brabandere et al., 2016), where the auxiliary network is fed with the same input as the target network, or with a related input (such as previous frames of a video). The goal is to make the target network more adaptive to the instantaneous input or task, rather than to generate diverse versions of the target network which are on an equal footing.

It has been shown (Denil et al., 2013) that for common machine learning tasks, there is a redundancy in the raw representation of the weights of several neural network models. (Ha et al., 2016) exploits this fact to train a small neural network that generates the weights for a larger target network. This has the advantage of reduced storage size in memory, and can also be seen as a means of regularization. However, the weight generating network in (Ha et al., 2016) does not have a controllable input, and therefore it cannot be used to generate diverse random samples of weights.

In (Lorraine & Duvenaud, 2018), a hypernetwork is presented whose input are hyperparameter values, and it gener-

ates weights for the target network, which correspond to the hyperparameter.

In HyperNEAT (Stanley et al., 2009) an auxiliary neural network is evolved using a genetic algorithm. This auxiliary network encodes the weights of the target network in the following way: The neurons of the target network are assigned coordinates on a grid. The weight between every pair of neurons is given by the output of the auxiliary network whose input are the coordinates of the two neurons. Generating the weights for the entire target network requires reapplication of the auxiliary network to all pairs of neurons. The goal of HyperNEAT is to generate networks with large scale that exhibit connectivity patterns which can be described as functions on low dimensions. There is no emphasis on generating a diverse set of such networks.

In (Zoph & Le, 2016), the auxiliary neural network is a recurrent network, trained via reinforcement learning to generate an optimal target architecture. The weights of the target network are obtained by standard training, not by generation.

Neural networks that generate their own weights have been discussed in (Schmidhuber, 1993; Chang & Lipson, 2018). These may have interesting implications, such as the ability of a network to self-introspect or self-replicate.

There are also Bayesian approaches for weight generation, which do not use a neural network as a generator. Markov Chain Monte Carlo (Neal, 1995; Welling & Teh, 2011) use a Markov chain in the weight space with equilibrium distribution that equals the required posterior. In (Graves, 2011; Blundell et al., 2015), the posterior is approximated as a diagonal Gaussian with trainable means and variances, and the VI (Hinton & Van Camp, 1993; Kingma, 2017) objective is used. Another VI method is (Gal & Ghahramani, 2016), which approximates the posterior of the weights as proportional to Bernoulli variables, which is equivalent to the dropout regularization technique (Srivastava et al., 2014).

B. Another Way to Obtain the Hypernetwork Loss Function

Here we demonstrate how the loss function (1) arises from another consideration. We denote by $p_\varphi(\theta)$ the probability distribution over Θ of $\theta = G(z; \varphi)$. We also denote by $p(\theta|p_{\text{data}})$ the required distribution over Θ , which we would like $p_\varphi(\theta)$ to be equal to. An optimistic choice for $p(\theta|p_{\text{data}})$ would be the following:

$$p(\theta|p_{\text{data}}) = \begin{cases} \frac{1}{Z} & \text{if } \mathcal{L}(\theta|p_{\text{data}}) = \min_{\theta'} \mathcal{L}(\theta'|p_{\text{data}}) \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where Z is a normalization constant. In other words, this distribution chooses only from the global optima, with equal probabilities. We can relax this optimistic form by turning it into a Gibbs distribution:

$$p(\theta|p_{\text{data}}) = \frac{\exp(-\lambda \mathcal{L}(\theta|p_{\text{data}}))}{Z}, \quad (7)$$

The hyperparameter $\lambda > 0$ controls how close this distribution is to the optimistic form, which is recovered for $\lambda \rightarrow \infty$. The relaxation is required so that there is higher connectivity between the mode regions of $p(\theta|p_{\text{data}})$. This will make it possible for $p_\varphi(\theta)$ to become close to $p(\theta|p_{\text{data}})$. Achieving this can be done by minimizing a loss function which is the Kullback-Leibler divergence between them:

$$L = D_{\text{KL}}(p_\varphi(\theta) \| p(\theta|p_{\text{data}})) = \int p_\varphi(\theta) \log \left(\frac{p_\varphi(\theta)}{p(\theta|p_{\text{data}})} \right) d\theta \quad (8)$$

(This integral can be defined also when $p_\varphi(\theta)$ is supported on a low dimensional manifold. In this case, $p_\varphi(\theta)$ can be written as a product of a delta function distribution and a finite distribution, where the delta serves as the restriction to the low dimensional manifold. The integration is understood to be only over this manifold, using the finite component of $p_\varphi(\theta)$. It can be seen from (6) that $p(\theta|p_{\text{data}})$ never vanishes on this manifold). Inserting (7) into (8) we get:

$$L = \lambda \mathbb{E}_{\theta \sim p_\varphi} \mathcal{L}(\theta|p_{\text{data}}) + \mathbb{E}_{\theta \sim p_\varphi} \log p_\varphi(\theta) + \log Z. \quad (9)$$

Noticing that $\log Z$ is a constant (does not depend on p_φ), we see that (9) is of the form of (1), where the diversity term is taken as the negation of the entropy.

C. Relation to Variational Inference

The objective function for hypernetworks using VI (Krueger et al., 2017; Louizos & Welling, 2017) is:

$$\tilde{L}_{\text{VI}}(\varphi|p_{\text{noise}}, D) = -\mathbb{E}_{z \sim p_{\text{noise}}}[\log p(D|G(z; \varphi)) - \log p_{\text{prior}}(G(z; \varphi)) + \log p(G(z; \varphi))], \quad (10)$$

where $D = \{(x_1, y_1), \dots, (x_n, y_n)\}$ is a data set of size n , p_{prior} is a prior. The last term on the right hand side of equation (10) is the probability distribution induced by the hypernetwork. The first term is the probability computed by the target network. Assuming iid samples and conditional independence, we have:

$$\begin{aligned} \log p(D|G(z; \varphi)) &\equiv \log p(y_1, \dots, y_n | x_1, \dots, x_n; G(z; \varphi)) \\ &= \sum_{i=1}^n \log p(y_i | x_i; G(z; \varphi)). \end{aligned} \quad (11)$$

We see that $\log p(D|G(z; \varphi))$ is an unbiased estimate of $n\mathbb{E}_{(x,y) \sim p_{\text{data}}} \log p(y|x; G(z; \varphi))$. Writing the objective in terms of “true” (as opposed to estimated) terms, we get:

$$L_{\text{VI}}(\varphi|p_{\text{noise}}, D) = -\mathbb{E}_{z \sim p_{\text{noise}}, (x,y) \sim p_{\text{data}}}[n \log p(y|x; G(z; \varphi)) - \log p_{\text{prior}}(G(z; \varphi)) + \log p(G(z; \varphi))], \quad (12)$$

The last term in this equation is the differential entropy of the generated weights. We see that equation (12) differs from equation (5) only in the presence of the prior term. However, we note that differential entropy is not a correct generalization of Shannon entropy, and it has some undesired properties (Marsh, 2013). Generalizing from the discrete case to the continuous case requires either binning of the sample space, or using a reference probability distribution. The later approach yields the relative entropy (which is identical to the Kullback-Leibler divergence). Therefore, equation (12) can be seen as equivalent to equation (5) if the entropy \mathbb{H} in equation (5) is measured with respect to a reference distribution p_{prior} .

D. The Hyperparameter λ in Variational Inference

In this section, we explain why the hyperparameter λ that appears in equation 1 could have also appeared in the objective for hypernetworks derived from VI (Krueger et al., 2017; Louizos & Welling, 2017).

We look first at the non-Bayesian case. Denote by $p_\theta(y|x)$ the probability distribution computed by the target network. $p_\theta(y|x)$ is a function approximator, with weights θ . In other words, $p_\theta(y|x)$ represents a family of models parametrized by θ . A main idea in machine learning is to use a flexible function approximator, such that some weights θ^* yield good approximations to the “true” probability distribution (This is in contrast with other fields, such as physics, where a minimal model is preferred, derived from an underlying theory and assumptions). For predictive applications, it does not matter what the ontological interpretation of the model $p_{\theta^*}(y|x)$ is, so long as it gives good approximations.

In the Bayesian case, the weights are considered random variables, with a prior distribution $p(\theta)$. To emphasize this, we promote the subscript θ in $p_\theta(y|x)$ to the argument of the target network: $p(y|x; \theta)$. Notice that this is not a function approximator anymore - it is one single specific model that describes how θ and x are combined to form the distribution over y (The distributions $p(\theta)$ and $p(y|x; \theta)$ jointly form a generative model). There are no parameters to optimize so as to get a “better” model. Moreover, this model is observed only for one (unknown) sample of θ . What is the ontological status of this model? If we treat this model as “true”, then we can infer the posterior distribution using Bayes’ law: $\log p(\theta|D) = \log p(D|\theta) + \log p(\theta) - \log Z$ for a normalizing factor Z (D is the data set, see Appendix C). But usually the model $p(y|x; \theta)$ was not chosen on the basis of a theory or assumptions about the data, and for almost all values of θ it is meaningless to speak of the validity of $p(y|x; \theta)$ as a model. It follows that we should treat the model as a mere approximation. When inferring the posterior, we should take into consideration the amount of trust that we have in the model. One way of doing this is by adding the hyperparameter λ to Bayes’ law $\log p(\theta|D) = \lambda \log p(D|\theta) + \log p(\theta) - \log Z$, and the hyperparameter λ would also appear in the VI objective.

E. Entropy Estimation

To backpropagate through the entropy term, using mini-batches of samples of z , we will need to use an estimator for differential entropy. We use the Kozachenko-Leonenko estimator (Kozachenko & Leonenko, 1987; Kraskov et al., 2004). For a set of samples $\theta_1, \theta_2, \dots, \theta_N$, the estimator is:

$$\hat{H} = \psi(N) + \frac{d}{N} \sum_{i=1}^N \log(\epsilon_i), \quad (13)$$

where ψ is the digamma function, d is the dimension of the samples, which we take to be the dimension of z , and ϵ_i is the distance from θ_i to its nearest neighbor in the set of samples. The original estimator takes d to be the dimension of θ . However, as explained in Appendix B, we are interested in an estimation of the differential entropy only on the manifold defined by $G(z)$, stripping away the delta functions that restrict θ to the manifold. We are using the Euclidean metric on the space Θ , although a better estimator would use the metric induced on the lower dimensional manifold. We don't do this due to the complications that it introduces. We omit terms in the definition of the estimator which are constants that do not matter for optimization (The reason that we did keep the terms $\psi(N)$ and $\frac{d}{N}$ is that N represents the batch size during training, but it also represents the size of the validation set during validation. If we want the entropy units to be comparable between the two, we need to keep the dependence on N , even though it does not matter for the optimization). This estimator is biased (Charzyńska & Gambin, 2015) but consistent in the mean square. In the case of the LeNet and ResNet-32 trials, we estimate the entropy term based on distances between on a subset of random samples for computational efficiency (and do not evaluate the full Cartesian product of pair-wise distances).

F. Target Network Architectures

Table 2. The target network architecture for MNIST. Note that all layers also have bias parameters. See (Krizhevsky et al., 2012; Simonyan & Zisserman, 2014) for definitions of the various terms.

Layer name	Layer components	Number of weights	Output size
input image			$28 \times 28 \times 1$
layer1	convolution: 5×5 , stride: 1×1 , padding: 'SAME', number of filters: 32	832	$14 \times 14 \times 32$
	activation: ReLU		
layer2	max pool: 2×2 , stride: 2×2 , padding: 'SAME'	12,816	$7 \times 7 \times 16$
	convolution: 5×5 , stride: 1×1 , padding: 'SAME', number of filters: 16		
	activation: ReLU		
layer3	fully-connected, number of filters: 8	6,280	8
	activation: ReLU		
layer4	fully-connected, number of filters: 10	90	10
	softmax		

Table 3. LeNet

Layer name	Layer components	Number of weights	Output size
input image			$32 \times 32 \times 3$
layer1	convolution: 5×5 , stride: 1×1 , padding: 'SAME', number of filters: 6	456	$14 \times 14 \times 6$
	activation: ReLU		
	max pool: 2×2 , stride: 2×2		
layer2	convolution: 5×5 , stride: 1×1 , padding: 'SAME', number of filters: 16	2,416	$5 \times 5 \times 16$
	activation: ReLU		
	max pool: 2×2 , stride: 2×2		
layer3	fully-connected, number of filters: 120	48,120	120
	activation: ReLU		
layer4	fully-connected, number of filters: 84	10,164	84
	activation: ReLU		
layer5	fully-connected, number of filters: 10	850	10
	softmax		

Table 4. ResNet-32

Layer name	Layer components	Number of weights	Output size
input image			$32 \times 32 \times 3$
layer1	convolution: 3×3 , stride: 1×1 , padding: 'SAME', number of filters: 16	448	$32 \times 32 \times 16$
	activation: ReLU		
layer2	convolution: 3×3 , stride: 1×1 , padding: 'SAME', number of filters: 16	$2,320 \times 10 =$ 23,200	$32 \times 32 \times 16$
	activation: ReLU		
layer3	convolution: number of filters: 32	$4,640 +$ $9,248 + 544 +$ $9,248 \times 8 =$ 88,416	$16 \times 16 \times 32$
	activation: ReLU		
layer4	convolution: number of filters: 64	$18,496 +$ $36,928 +$ $2112 +$ $36,928 \times 8 =$ 352,960	$1 \times 1 \times 64$
	activation: ReLU		
	avg pool: $8 \times$ 8 , stride: 1×1 , padding: 0		
layer5	fully-connected, number of filters: 10	650	10
	softmax		

G. MNIST Hypernetwork Architecture and Training

The specifications of the architecture for the hypernetwork are given in Table 5. We choose the input z to the network to be a 300 dimensional vector, drawn from a uniform distribution over $[-1, 1]^{300}$. The code sizes for the weight generators are chosen to be 15 for all four layers. The extractor and the weight generators are all MLPs with leaky-ReLU activations (Maas et al., 2013). Batch normalization (Ioffe & Szegedy, 2015) is employed in all layers besides the output layers of each sub-network. We do not use bias parameters in the hypernetwork, since our empirical evidence suggests (albeit inconclusively) that this helps for diversity. The total number of parameters in G is 633,640.

We trained the hypernetwork using the Adam optimizer (Kingma & Ba, 2014). The gradients of the loss were estimated using minibatches of 32 samples of z , and 32 images per sample of z (we use different images for each noise sample). We found that this relatively large batch size is a good operating point in terms of the tradeoff between estimator variance and learning rate. We trained on a total of 13,000 minibatches, and this took about 30 minutes on a single NVIDIA Tesla K80 GPU.

Table 5. The layer structures for each of the fully-connected sub-networks of the hypernetwork. The layer structures are in the format (input size) \rightarrow (first layer size) \rightarrow (second layer size) \rightarrow (output size). The sub-networks do not include bias terms. The number of parameters shown here does not include the batch normalization parameters.

Sub-net	Layer structure	Parameters
E	$300 \rightarrow 300 \rightarrow 300 \rightarrow 855$ $= 15 \cdot (32 + 16 + 8 + 1)$	436,500
W_1	$15 \rightarrow 40 \rightarrow 40 \rightarrow 26$ $= 5 \cdot 5 + 1$	3240
W_2	$15 \rightarrow 100 \rightarrow 100 \rightarrow 801$ $= 5 \cdot 5 \cdot 32 + 1$	91,600
W_3	$15 \rightarrow 100 \rightarrow 100 \rightarrow 785$ $= ((28^2)/(4^2)) \cdot 16 + 1$	90,000
W_4	$15 \rightarrow 60 \rightarrow 60 \rightarrow 90$ $= (8 + 1) \cdot 10$	9,900

H. CIFAR Hypernetwork Architecture and Training

The specifications of the architecture for the hypernetwork are given in Table 6 and 7 for the LeNet and ResNet-32 target network, respectively. We choose the input z to the network to be a 100 and 500 dimensional vector for LeNet and ResNet-32, respectively, drawn from a standard normal distribution. The code sizes are determined by multiplying the number of filters in a given layer by a constant (here 8).

We trained the hypernetwork using the Adam optimizer (Kingma & Ba, 2014) with a learning rate of $1e-4$ and $\beta = (0.9, 0.999)$. The gradients of the loss were estimated using minibatches of 10 samples of z , and 512 images per sample of z (we use different images for each noise sample).

For ResNet-32, we employed skip connections in the extractor and weight generators to ease optimization. Further, the weight generators output scale γ and shift β parameters of the batch-normalization operations (Ioffe & Szegedy, 2015). In order to apply batch-normalization in inference mode, we first compute moving averages of mean and variance statistics with one additional full pass over the training data.

Table 6. LeNet

Sub-net	Layer structure	Parameters
E	$100 \rightarrow 200 \rightarrow 200 \rightarrow 236$ $= 8 \cdot (6 + 16 + 120 + 84 + 10)$	107,200
W_1	$8 \rightarrow 400 \rightarrow 400 \rightarrow 76$ $= 5 \cdot 5 \cdot 3 + 1$	193,600
W_2	$8 \rightarrow 400 \rightarrow 400 \rightarrow 151$ $= 5 \cdot 5 \cdot 6 + 1$	223,600
W_3	$8 \rightarrow 400 \rightarrow 400 \rightarrow 401$ $= 5 \cdot 5 \cdot 16 + 1$	323,600
W_4	$8 \rightarrow 400 \rightarrow 400 \rightarrow 10, 164$ $= (120 + 1) \cdot 84$	4,228,800
W_5	$8 \rightarrow 400 \rightarrow 400 \rightarrow 850$ $= (84 + 1) \cdot 10$	503,200

Table 7. ResNet-32

Sub-net	Layer structure	Parameters
E	$500 \rightarrow 1000 \rightarrow 1000 \rightarrow 1104$ $= 8 \cdot (16 + 16 + 32 + 64 + 10)$	2,935,200
W_1	$8 \rightarrow 800 \rightarrow 800 \rightarrow 10$ $= 3 \cdot 3 \cdot 3 + 1$	654,400
W_2	$8 \rightarrow 800 \rightarrow 800 \rightarrow 145$ $= 3 \cdot 3 \cdot 16 + 1$	762,400
W_3	$8 \rightarrow 800 \rightarrow 800 \rightarrow 145$ $= 3 \cdot 3 \cdot 16 + 1$	762,400
W_4	$8 \rightarrow 800 \rightarrow 800 \rightarrow 289$ $= 3 \cdot 3 \cdot 32 + 1$	887,600
W_5	$8 \rightarrow 800 \rightarrow 800 \rightarrow 650$ $= (64 + 1) \cdot 10$	1,166,400

I. Paths in Weight Space - MNFG




Figure 9. Accuracies along paths, with end points z_1, z_2 sampled at random, **for MNFG**. Each graph corresponds to a different sampled pair of endpoints. The dashed lines are the direct paths, and the solid lines are the interpolated paths.

J. PCA Scatter Plots

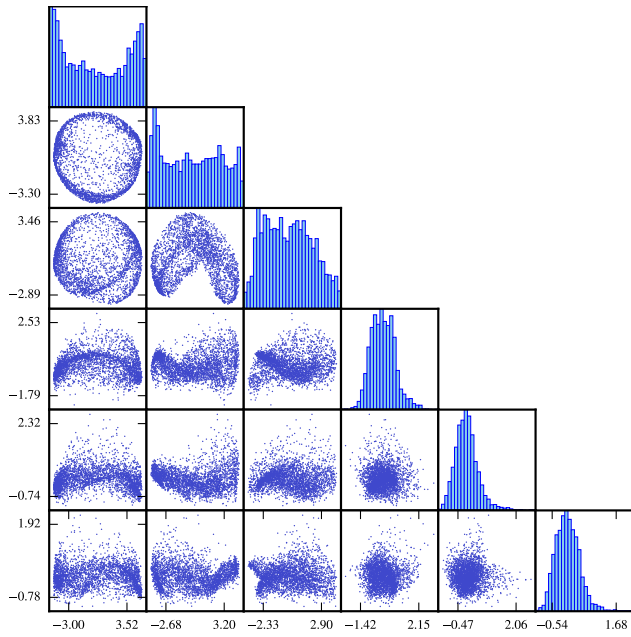


Figure 10. Scatter plots of the generated weights for a specific first layer filter, in PCA space. The (i, j) scatter plot has principal component i against principal component j .

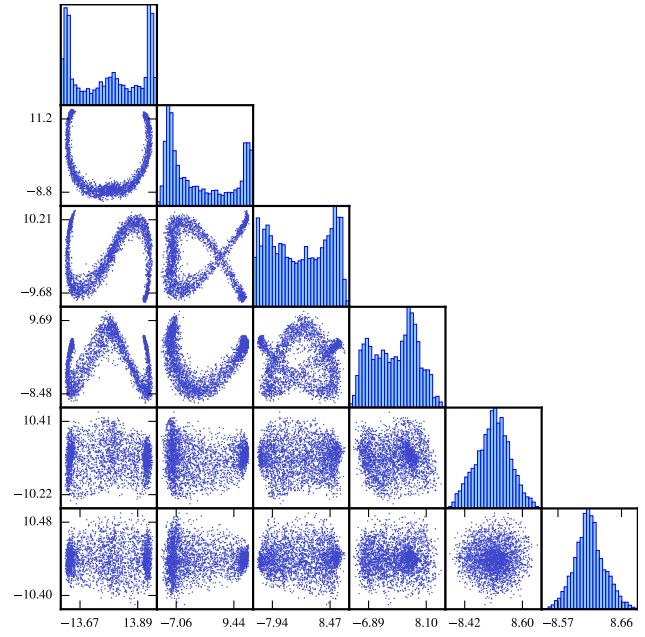


Figure 12. Scatter plots of the generated weights for the entire first layer, in PCA space. The (i, j) scatter plot has principal component i against principal component j .

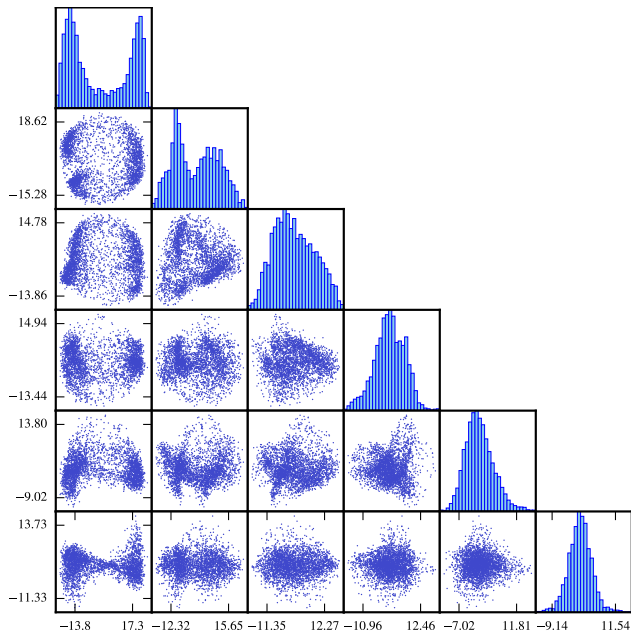


Figure 11. Scatter plots of the generated weights for a specific third layer filter, in PCA space. The (i, j) scatter plot has principal component i against principal component j .




Figure 13. Scatter plots of the generated weights for the entire second layer, in PCA space. The (i, j) scatter plot has principal component i against principal component j .

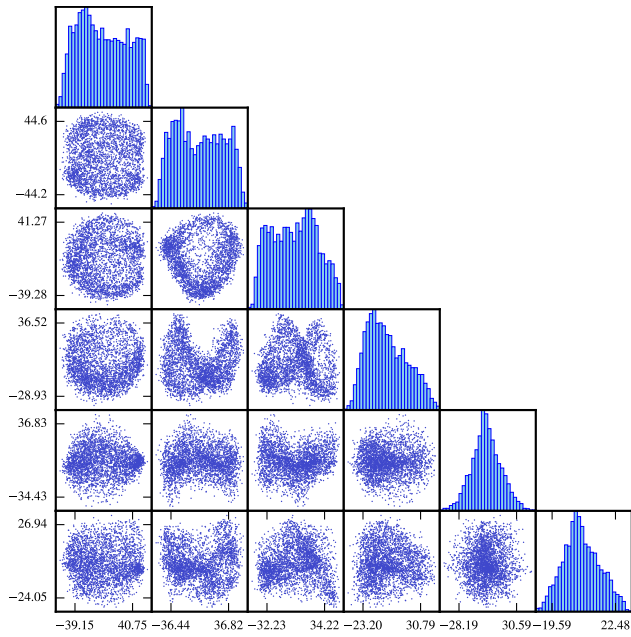


Figure 14. Scatter plots of the generated weights for the entire third layer, in PCA space. The (i, j) scatter plot has principal component i against principal component j .




Figure 16. Scatter plots of the generated weights for the entire target network, in PCA space. The (i, j) scatter plot has principal component i against principal component j .

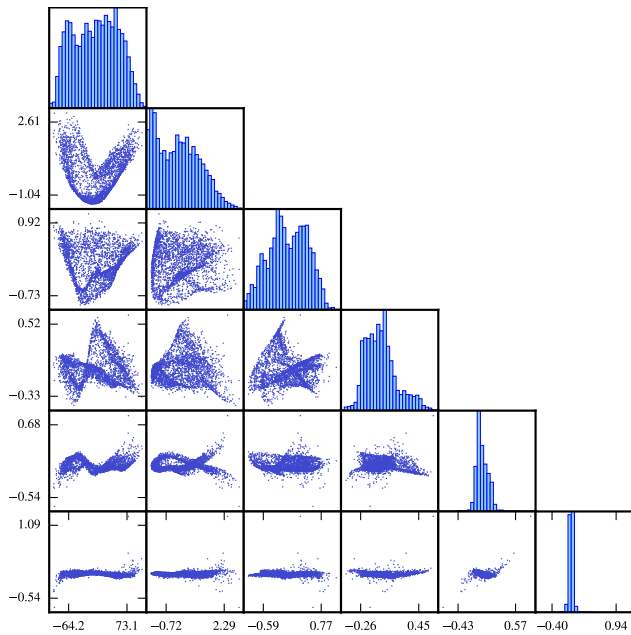


Figure 15. Scatter plots of the generated weights for the entire fourth layer, in PCA space. The (i, j) scatter plot has principal component i against principal component j .

K. PCA Scatter Plots - MNFG




Figure 17. Scatter plots of the generated weights for the entire first layer, in PCA space, for MNFG. The (i, j) scatter plot has principal component i against principal component j .




Figure 18. Scatter plots of the generated weights for the entire second layer, in PCA space, for MNFG. The (i, j) scatter plot has principal component i against principal component j .




Figure 19. Scatter plots of the generated weights for the entire third layer, in PCA space, for MNFG. The (i, j) scatter plot has principal component i against principal component j .




Figure 20. Scatter plots of the generated weights for the entire fourth layer, in PCA space, for MNFG. The (i, j) scatter plot has principal component i against principal component j .




Figure 21. Scatter plots of the generated weights for the entire target network, in PCA space, for MNFG. The (i, j) scatter plot has principal component i against principal component j .