

---

# Subspace Inference for Bayesian Deep Learning

---

Pavel Izmailov<sup>\*1</sup>, Wesley J. Maddox<sup>\*1</sup>, Polina Kirichenko<sup>\*1</sup>, Timur Garipov<sup>\*4</sup>  
Dmitry Vetrov<sup>2,3</sup>, Andrew Gordon Wilson<sup>1</sup>

<sup>1</sup>Cornell University, <sup>2</sup>Higher School of Economics, <sup>3</sup>Samsung-HSE Laboratory,  
<sup>4</sup>Samsung AI Center in Moscow

## Abstract

Bayesian inference was once a gold standard for learning with neural networks, providing accurate full predictive distributions and well-calibrated uncertainty. However, scaling Bayesian inference techniques to deep neural networks is challenging due to the high dimensionality of the parameter space. In this paper, we construct low-dimensional subspaces of parameter space, such as the first principal components of the stochastic gradient descent (SGD) trajectory, which contain diverse sets of high performing models. In these subspaces, we are able to apply elliptical slice sampling and variational inference, which struggle in the full parameter space. We show that Bayesian model averaging over the induced posterior in these subspaces produces accurate predictions and well-calibrated predictive uncertainty for both regression and image classification.

## 1 INTRODUCTION

Bayesian methods were once the state-of-the-art approach for inference with neural networks (MacKay, 2003; Neal, 1996a). However, the parameter spaces for modern deep neural networks are extremely high dimensional, posing challenges to standard Bayesian inference procedures.

In this paper, we propose a different approach to approximate Bayesian inference in deep learning models: we design a low-dimensional subspace  $\mathcal{S}$  of the weight space and perform posterior inference over the parameters within this subspace. We call this approach *Subspace Inference (SI)*.<sup>1</sup>

<sup>\*</sup>Equal contribution.

<sup>1</sup>PyTorch code is available at <https://github.com/wjmaddox/drbytes>.

It is our contention that the subspace  $\mathcal{S}$  can be chosen to contain a diverse variety of representations, corresponding to different high quality predictions, over which Bayesian model averaging leads to accuracy gains and well-calibrated uncertainties.

In Figure 1, we visualize the samples from the approximate posterior and the corresponding predictive distributions in performing subspace inference for a ten-dimensional random subspace, and a rich two-dimensional subspace containing a low-loss curve between two independently trained SGD solutions (see Garipov et al., 2018) on a synthetic one-dimensional regression problem. As we can see, the predictive distribution corresponding to a random subspace does not capture a diverse set of possible functions required for greater uncertainty away from the data, but sampling from the posterior in the rich curve subspace provides meaningful uncertainty over functions.

Our paper is structured as follows. We begin with a discussion of related work in Section 2. In Section 3, we describe the proposed method for inference in low-dimensional subspaces of the parameter space. In Section 4, we discuss possible choices of the low-dimensional subspaces. In particular, we consider random subspaces, subspaces corresponding to the first principal components of the SGD trajectory (Maddox et al., 2019), and subspaces containing low-loss curves between independently trained solutions (Garipov et al., 2018).

We analyze the effects of using different subspaces and approximate inference methods, by visualizing uncertainty on a regression problem in Section 5.1. We then apply the proposed method to a range of UCI regression datasets in Section 5.2, as well as CIFAR-10 and CIFAR-100 classification problems in Section 5.3, achieving consistently strong performance in terms of both test accuracy and likelihood. Although the dimensionality of the weight space for modern neural networks is extraordinarily large, we show that surprisingly low dimensional subspaces contain

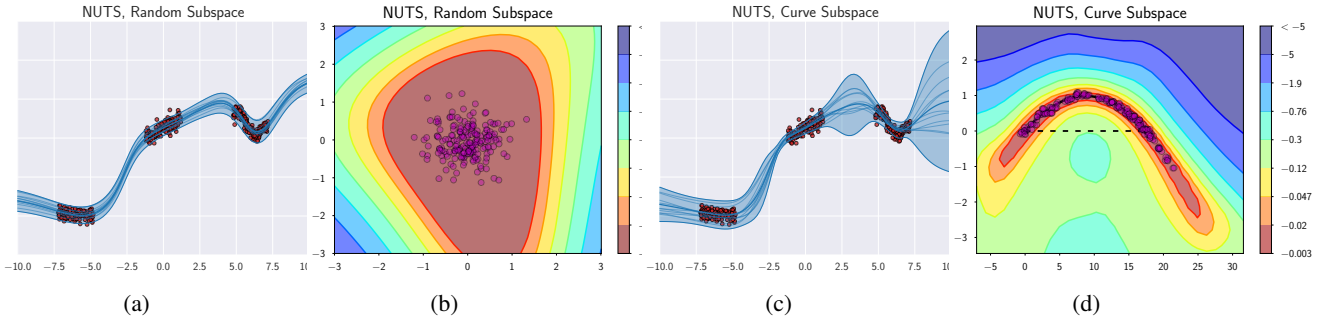


Figure 1: Predictive distribution and samples in the parameter space for subspace inference on a synthetic regression problem in a random subspace (a, b) and subspace containing a near-constant loss (log posterior) curve between two independently trained solutions (c, d) (see Garipov et al., 2018, for details). On the plots (a, c) data points are shown with red circles, the shaded region represents the  $3\sigma$ -region of the predictive distribution at each point, the predictive mean is shown with a thick blue line and sample trajectories are shown with thin blue lines. Panels (b, d) show the contour plots of the posterior log-density within the corresponding subspace; magenta circles represent samples from the posterior in the subspace. In the rich subspace containing the near-constant loss curve, the samples produce better uncertainty estimates and more diverse trajectories. We use a small fully-connected network with 4 hidden layers. See Section 5.1 for more details.

a rich diversity of representations. For example, we can construct 5 dimensional subspaces where Bayesian model averaging leads to notable performance gains on a 36 million dimensional WideResNet trained on CIFAR-100.

We summarize subspace inference in Algorithm 1. We note that this procedure uses three modular steps: (1) construct a subspace; (2) posterior inference in the subspace; and (3) form a Bayesian model average. Different design choices are possible for each step. For example, choices for the subspace include a random subspace, a PCA subspace, or a mode connected subspace. Many other choices are also possible. For posterior inference, one can use deterministic approximations over the parameters in the subspace, such as a variational method, or MCMC.

---

#### Algorithm 1 Bayesian Subspace Inference

---

**Input:** data,  $\mathcal{D}$ ; model,  $\mathcal{M}$ ;

1. Construct subspace, i.e. using Alg. 2, Section 4
  2. Posterior inference within subspace (Section 3)
  3. Form a Bayesian model average (Section 3.2, 3.3)
- 

## 2 RELATED WORK

Maddox et al. (2019) proposed SWAG, which forms an approximate Gaussian posterior over neural network weights, with a mean and low rank plus diagonal covariance matrix formed from a partial trajectory of the SGD iterates with a modified learning rate schedule. SWAG provides scalable Bayesian model averaging, with compelling accuracy and calibration results on CIFAR and ImageNet. The low-rank part of the SWAG covariance defines a distribution over a low-dimensional subspace

spanned by the first principal components of the SGD iterates.

Silva and Kalaitzis (2015) consider the related problem of Bayesian inference using projected methods for constrained latent variable models, with applications to probabilistic PCA (Roweis, 1998; Bishop, 1999).

Pradier et al. (2018) propose to perform variational inference (VI) in a subspace formed by an auto-encoder trained on a set of models generated from fast geometric ensembling (Garipov et al., 2018); this approach requires training several models and fitting an auto-encoder, leading to limited scalability.

Similarly, Karaletsos et al. (2018) propose to use a meta-prior in a low-dimensional space to perform variational inference for BNNs. This approach can be viewed as a generalization of hyper-networks (Ha et al., 2017). Alternatively, both Titsias (2017) and Krueger et al. (2017) propose Bayesian versions of hyper-networks to store meta-models of parameters.

Patra and Dunson (2018) provide theoretical guarantees for Bayesian inference in the setting of constrained posteriors. Their method samples from the unconstrained posterior before using a mapping into the constrained parameter space. In their setting, the constraints are chosen a priori; on the other hand, we choose the constraints (e.g. the subspace) after performing unconstrained inference via SGD.

Bayesian coresets (Huggins et al., 2016) use a weighted combination of the full dataset and Bayesian compressed regression (Guhaniyogi and Dunson, 2015) uses random projections of the data inputs in linear regression settings;

both are designed for the purpose of efficient inference, but unlike our subspace inference, these methods operate solely in data space, rather than in parameter space.

### 3 INFERENCE WITHIN A SUBSPACE

In this section we discuss how to perform Bayesian inference within a given subspace of a neural network. In Section 4 we will propose approaches for effectively constructing such subspaces.

#### 3.1 MODEL DEFINITION

We consider a model,  $\mathcal{M}$ , with weight parameters  $w \in \mathbb{R}^d$ . The model has an associated likelihood for the dataset,  $\mathcal{D}$ , given by  $p_{\mathcal{M}}(\mathcal{D}|w)$ .

We perform inference in a  $K$ -dimensional subspace  $\mathcal{S}$  defined by

$$\begin{aligned} \mathcal{S} &= \{w|w = \hat{w} + z_1v_1 + \dots z_Kv_K\} \\ &= \{w|w = \hat{w} + Pz\}, \end{aligned} \quad (1)$$

where  $P = (v_1^\top, \dots, v_K^\top) \in \mathbb{R}^{d \times K}$ ,  $\hat{w} \in \mathbb{R}^d$ ,  $z = (z_1, \dots, z_K)^\top \in \mathbb{R}^K$ . With a fixed  $\hat{w}$  and projection matrix  $P$ , which assign the subspace, the free parameters of the model, over which we perform inference, are now simply  $z \in \mathbb{R}^K$ . We describe choices for  $\hat{w}$  and  $P$  in Section 4.




Figure 2: Illustration of subspace  $\mathcal{S}$  with shift vector  $\hat{w}$  and basis vectors  $v_1, v_2$ , with a contour plot of the posterior log-density over parameters  $z$ .

The new model has the likelihood function:

$$p(\mathcal{D}|z) = p_{\mathcal{M}}(\mathcal{D}|w = \hat{w} + Pz), \quad (2)$$

where the right-hand side represents the likelihood for the model,  $\mathcal{M}$ , with parameters  $\hat{w} + Pz$  and data  $\mathcal{D}$ . We can then perform Bayesian inference over the low-dimensional subspace parameters  $z$ . We illustrate the subspace parameterization as well as the posterior log-density over parameters  $z$  in Figure 2.

We emphasize that the new model (2) is *not* a reparameterization of the original model, as the mapping from the full parameter space to the subspace is not invertible. For this reason, we consider the subspace model parameterized by  $z$  as a *different model* that shares many functional properties with the original model  $\mathcal{M}$  (see Section A.1 for an extended discussion).

#### 3.2 BAYESIAN MODEL AVERAGING

We can sample from an induced posterior over the weight parameters in the original space by first sampling from the posterior over the parameters in the subspace  $\tilde{z} \sim p(z|\mathcal{D})$ , using an approximate inference method of choice, and then transforming those samples into the original space as  $\tilde{w} = \hat{w} + P\tilde{z}$ .

To perform Bayesian model averaging on new test data points,  $\mathcal{D}^*$ , we can compute a Monte Carlo estimate of the integral

$$\begin{aligned} p(\mathcal{D}^*|\mathcal{D}) &= \int p_{\mathcal{M}}(\mathcal{D}^*|\tilde{w} = \hat{w} + Pz)p(z|\mathcal{D})dz \quad (3) \\ &\approx \frac{1}{J} \sum_{j=1}^J p_{\mathcal{M}}(\mathcal{D}^*|\tilde{w} = \hat{w} + P\tilde{z}_j), \tilde{z}_j \sim p(z|\mathcal{D}). \end{aligned}$$

Using the Monte Carlo estimate of the integral in (3) produces mixtures of Gaussian predictive distributions for regression tasks with Gaussian likelihoods, and categorical distributions for classification tasks.

#### 3.3 APPROXIMATE INFERENCE PROCEDURES

Our goal is to approximate the posterior  $p(z|\mathcal{D})$  over the free parameters  $z$  in the subspace  $\mathcal{S}$ , in order to perform a Bayesian model average. As we can set the number of parameters,  $K$ , to be much smaller than the dimensionality  $d$  of the full parameter space, performing Bayesian inference becomes considerably more tractable in the subspace. We can make use of a wide range of approximate inference procedures, even if we are working with a large modern neural network.

In particular, we can use powerful and exact full-batch MCMC methods to approximately sample from  $p(z|\mathcal{D})$ , such as Hamiltonian Monte Carlo (HMC) (Neal et al., 2011) or elliptical slice sampling (ESS) (Murray et al., 2010). Alternatively, we can perform a deterministic approximation  $q(z|\mathcal{D}) \approx p(z|\mathcal{D})$ , for example using Laplace or a variational approach, and then sample from  $q$ . The low dimensionality of the problem allows us to choose very flexible variational families such as RealNVP (Dinh et al., 2017) to approximate the posterior.

Ultimately, the inference procedure is an experimental design choice, and we are free to use a wide range of approximate inference techniques.

### 3.4 PRIOR CHOICE

There is a significant practical difference between *Bayesian model averaging* (Section 3.2) and standard training (regularized maximum likelihood estimation) for a range of priors  $p(z)$ , including vague priors. However, the exact specification of the prior itself,  $p(z)$ , if sufficiently diffuse, is not crucial for good performance. Moreover, what matters is not the prior over parameters in isolation, but how this prior interacts with the functional form of the model. The neural network itself induces a structured prior distribution over functions, when combined with a vague prior over its parameters. We discuss reasonable choices of priors for various subspaces in Section 4.

### 3.5 PREVENTING POSTERIOR CONCENTRATION WITH FIXED TEMPERATURE POSTERiors

In the model proposed in Section 3.1, there are only  $K \ll N$  parameters as opposed to  $d \gg N$  parameters in the full weight space, while the number of observed data points  $N$  is constant. In this setting, the posterior can overly concentrate around the maximum likelihood estimate (MLE), becoming too constrained by the data, leading to overconfident uncertainty estimates.

To address the issue of premature posterior concentration in the subspace, we propose to introduce a temperature hyperparameter that scales the likelihood. In particular, we use the *tempered* posterior:

$$p_T(z|\mathcal{D}) \propto \underbrace{p(\mathcal{D}|z)}_{\text{likelihood}}^{1/T} \underbrace{p(z)}_{\text{prior}}. \quad (4)$$

When  $T = 1$  the true posterior is recovered, and as  $T \rightarrow \infty$ , the tempered posterior  $p_T(z|\mathcal{D})$  approaches the prior  $p(z)$ .

The temperature  $T$  is a hyper-parameter that can be determined through cross-validation. We study the effect of temperature on the performance of subspace inference in Section F.1. When the temperature is close to 1 the posterior concentrates around the MLE and subspace inference fails to improve upon maximum likelihood training. As  $T$  becomes large, subspace inference produces increasingly less confident predictions. In Section F.1, good performance can be achieved with a broad range of  $T$ .

Tempered posteriors are often used in Bayesian inference algorithms to enhance multi-modal explorations (e.g.,

Geyer and Thompson, 1995; Neal, 1996b). Similarly, Watanabe (2013) uses a tempered posterior to recover an expected generalization error of Bayesian models.

## 4 SUBSPACE CONSTRUCTION

In the previous section we showed how to perform inference in a given subspace  $\mathcal{S}$ . We now discuss various ways to construct  $\mathcal{S}$ .

### 4.1 RANDOM SUBSPACES

To construct a simple random subspace,  $\mathcal{S}$ , we draw  $K$  random  $v_1, \dots, v_K \sim \mathcal{N}(0, I_p)$  in the weight space. We then rescale each of the vectors to have norm 1. Random subspaces require only drawing  $Kp$  random normal numbers and so are quick to generate and form, but contain little information about the model. In related work, Li et al. (2018a) train networks from scratch in a random subspace without a shift vector, requiring projections into much higher dimensions than are considered in this paper.

We use the weights of a network pre-trained with stochastic weight averaging (SWA) (Izmailov et al., 2018) as the shift vector  $\hat{w} = w_{\text{SWA}}$ . In particular, we run SGD with a high constant learning rate from a pre-trained solution, and form the average  $w_{\text{SWA}} = \frac{1}{T} \sum_i w_i$  from the SGD iterates  $w_i$ .

Since the log likelihoods as a function of neural network parameters for random subspaces appear approximately quadratic (Izmailov et al., 2018) a reasonable prior for  $p(z)$  is  $\mathcal{N}(0, \sigma^2 I_K)$ .

### 4.2 PCA OF THE SGD TRAJECTORY

Intuitively, we want the subspace  $\mathcal{S}$  over which we perform inference to (1) contain a diverse set of models that produce meaningfully different predictions and (2) be cheap to construct. Garipov et al. (2018) and Izmailov et al. (2018) argue that the subspace spanned by the SGD trajectory satisfies both (1) and (2). They run SGD starting from a pre-trained solution with a high constant learning rate and then ensemble predictions or average the weights of the iterates. Further, Maddox et al. (2019) showed fitting the SGD iterates with a Gaussian distribution with a low-rank plus diagonal covariance for scalable Bayesian model averaging provides well-calibrated uncertainty estimates. Finally, Li et al. (2018b) and Maddox et al. (2019) used the first few PCA components of the SGD trajectory for loss surface visualization. These observations motivate inference *directly* in the subspace spanned by the SGD trajectory.

We propose to use the first few PCA components  $v_i$  of

the SGD trajectory to define the basis of the subspace. As in Izmailov et al. (2018), we run SGD with a high constant learning rate from a pre-trained solution and capture snapshots  $w_i$  of weights at the end of each of  $T$  epochs. We store the deviations  $a_i = w_{\text{SWA}} - w_i$  for the last  $M$  epochs. The number  $M$  here is determined by the amount of memory we can use.<sup>2</sup> We then run PCA based on randomized SVD (Halko et al., 2011)<sup>3</sup> on the matrix  $A$  comprised of vectors  $a_1, \dots, a_M$  and use the first  $K$  principal components  $v_1, \dots, v_K$  to define the subspace (1). Like for the random subspace, we use the SWA solution (Izmailov et al., 2018) for the shift vector,  $\hat{w}$ . We summarize this procedure in Algorithm 2.

---

**Algorithm 2** Subspace Construction with PCA

---

$w_0$ : pretrained weights;  $\eta$ : learning rate;  $T$ : number of steps;  $c$ : moment update frequency;  $M$ : maximum number of columns in deviation matrix;  $K$ : rank of PCA approximation;  $P$ : projection matrix for subspace

$w_{\text{SWA}} \leftarrow w_0$  {Initialize mean}

**for**  $i \leftarrow 1, 2, \dots, T$  **do**

$w_i \leftarrow w_{i-1} - \eta \nabla_w \mathcal{L}(w_{i-1})$  {SGD update}

**if**  $\text{MOD}(i, c) = 0$  **then**

$n \leftarrow i/c$  {Number of models}

$w_{\text{SWA}} \leftarrow \frac{nw_{\text{SWA}} + w_i}{n+1}$  {Update mean}

**if**  $\text{NUM\_COLS}(A) = M$  **then**

$\text{REMOVE\_COL}(A[:, 1])$

$\text{APPEND\_COL}(A, w_i - \bar{w})$  {Store deviation}

$U, S, V^\top \leftarrow \text{SVD}(A)$  {Truncated SVD}

**return**  $\hat{w} = w_{\text{SWA}}, P = SV^\top$

---

Maddox et al. (2019) showed empirically that the log likelihood in the subspace looks locally approximately quadratic, and so a reasonable choice of prior  $p(z)$  is  $\mathcal{N}(0, I)$ , when scaling PCA vectors  $v_i$  to have norms proportional to the singular values of the matrix  $A$  as in Algorithm 2.

**Relationship to Eigenvalues of the Hessian** Li et al. (2018b) and Gur-Ari et al. (2019) argue that the first principal components of the SGD trajectory correspond to the top eigenvectors of the Hessian of the loss, and that these eigenvectors change slowly during training. This observation suggests that these principal components captures many of the sharp directions of the loss surface, corresponding to large Hessian eigenvalues. We expect then that our PCA subspace should include variation in

<sup>2</sup>We use  $M = 20$  in our experiments. To side-step any memory issues, we could use any online PCA technique instead, such as frequent directions (Ghashami et al., 2016).

<sup>3</sup>Implemented in `sklearn.decomposition.TruncatedSVD`.

the type of functions that it contains. See Appendix C for more details as well as a computation of Hessian and Fisher eigenvalues through a GPU accelerated Lanczos algorithm (Gardner et al., 2018).

### 4.3 CURVE SUBSPACES

Garipov et al. (2018) proposed a method to find paths of near-constant low loss (and consequently high posterior density) in the weight space between converged SGD solutions starting from different random initializations. These curves lie in 2-dimensional subspaces of the weight space. We visualize the loss surface in such a space for a synthetic regression problem in Figure 1 (d). This *curve subspace* provides an example of a rich subspace containing diverse high performing models, and stress-tests the inference procedure for effectively exploring a highly non-Gaussian distribution.

To parameterize the curve subspace we set  $\hat{w} = (w_1 + w_0)/2$ ,  $v_1 = (w_0 - \hat{w})/\|w_0 - \hat{w}\|$ ,  $v_2 = (w_{1/2} - \hat{w})/\|w_{1/2} - \hat{w}\|$ , where  $w_0$  and  $w_1$  are the endpoints, and  $w_{1/2}$  is the midpoint of the curve.

In this case, the posterior in the subspace is clearly non-Gaussian. However, a vague but centred standard Gaussian prior  $\mathcal{N}(0, I)$  is reasonable as a simple choice with our parameterization of the curve subspace.

### 4.4 COMPUTATIONAL COST OF SUBSPACE CONSTRUCTION

We consider the cost of constructing each of the subspaces described in Section 4.1-4.3. We note that constructing any subspace in our approach is a *one-time* computation.

The random subspace (Section 4.1) is virtually free to construct, as it only requires sampling  $K$  independent Gaussian vectors.

To construct the PCA subspace (Section 4.2), we run SVD on the deviation matrix  $A$ , which is a *one-time* computation and is very fast. In particular, exact SVD takes  $\mathcal{O}(\min(M^2d, d^2M))$ , while randomized SVD takes  $\mathcal{O}(Md \log K + (M+d)K^2)$  (see Section 1.4.1 of Halko et al. (2011)). For our largest examples, while the number of parameters in the model  $d$  is on the order of  $10^7$ ,  $M$  is on the order of  $10^2$ ; thus, taking the exact SVD is linear in the number of parameters. For example, using standard hardware (a Dell XPS-15 laptop, with Intel core i7, 16Gb RAM), it takes 4 minutes to perform *exact* SVD on our largest model, WideResNet on CIFAR-100, with 36 million parameters. By comparison, it takes approximately eight hours on an NVIDIA 1080Ti GPU to train the same WideResNet on CIFAR-100 to completion.

The curve subspace (Section 4.3) is the most expensive

to construct, as it requires pre-training the two solutions corresponding to the endpoints of the curve, and then running the curve-finding procedure in Garipov et al. (2018), which in total is roughly  $3\times$  the cost of training a single DNN of the same architecture.

Constructing the subspace is in general very fast and readily applicable to large deep networks, with minimal overhead compared to standard training.

## 5 EXPERIMENTS

We evaluate subspace inference empirically using the random, PCA, and curve subspace construction methods discussed in Section 4, in conjunction with the approximate posterior inference methods discussed in Section 3.3. In particular, we experiment with the No-U-Turn-Sampler (NUTS) (Hoffman and Gelman, 2014), Elliptical Slice Sampling (ESS) (Murray et al., 2010) and Variational Inference with fully factorized Gaussian approximation family (VI) and Real-valued Non-Volume Preserving flow (RealNVP) (Dinh et al., 2017) family. Section B contains more details on each of the approximate inference methods. We use the priors for  $p(z)$  discussed in Section 4.

We show that approximate Bayesian inference within a subspace provides good predictive uncertainties on regression problems, first visually in Section 5.1 and then quantitatively on UCI datasets in Section 5.2.

We then apply subspace inference to large-scale image classification on CIFAR-10 and CIFAR-100 and obtain results competitive with state-of-the-art scalable Bayesian deep learning methods.

### 5.1 VISUALIZING REGRESSION UNCERTAINTY

We want predicted uncertainty to grow as we move away from the data. Far away from the data there are many possible functions that are consistent with our observations and thus there should be greater uncertainty. However, this intuitive behaviour is difficult to achieve with Bayesian neural networks (Foong et al., 2019). Further, log-likelihoods on benchmark datasets (see Section 5.2) do not necessarily test this behaviour, where overconfident methods can obtain better likelihoods (Yao et al., 2019).

We use a fully-connected architecture with hidden layers that have  $[200, 50, 50, 50]$  neurons respectively. The network takes two inputs,  $x$  and  $x^2$  (the redundancy helps with training), and outputs a single real value  $y = f(x)$ . To generate the data we set the weights of the network with this same architecture randomly, and evaluate the predictions  $f(x)$  for 400 points sampled uniformly in in-

tervals  $[-7.2, -4.8]$ ,  $[-1.2, 1.2]$ ,  $[4.8, 7.2]$ . We add Gaussian noise to the outputs  $y = f(x) + \epsilon(x)$ . We show the data with red circles in Figure 3.

We train an SWA solution (Izmailov et al., 2018), and construct three subspaces: a ten-dimensional random subspace, ten-dimensional PCA-subspace and a two-dimensional curve subspace (see Section 4). We then run each of the inference methods listed in Section B in each of the subspaces. We visualize the predictive distributions in the observed space for each combination of method and subspace in Figure 8 and the posterior density overlaid by samples of the subspace parameters in Figure 9.

In Figure 9 the shape of the posterior in random and PCA subspaces is close to Gaussian, and all approximate inference methods produce reasonable samples. In the mode connecting curve subspace the posterior has a more complex shape. The variational methods were unable to produce a reasonable fit. The simple variational approach is highly constrained in its Gaussian representation of the posterior. However, RealNVP in principle has the flexibility to fit many posterior approximating distributions, but perhaps lacks the inductive biases to easily fit these types of curvy distributions in practice, especially when trained with the variational ELBO. On the other hand, certain MCMC methods, such as elliptical slice sampling, can effectively navigate these distributions.

In the top row of Figure 3 we visualize the predictive distributions for elliptical slice sampling in each of the subspaces. In order to represent this uncertainty, the posterior in the subspace must assign mass to settings of the weights that give rise to models making a diversity of predictions in these regions. In the random subspace, the predictive uncertainty does not significantly grow far away from the data, suggesting that it does not capture a diversity of models. On the other hand, the PCA subspace captures a diverse collection of models, with uncertainty growing away from the data. Finally, the predictive distribution for the curve subspace is the most adaptive, suggesting that it contains the greatest variety of models corresponding to weights with high posterior probabilities.

In the bottom row of Figure 3 we visualize the predictive distributions for simple variational inference applied in the original parameter space (as opposed to a subspace), SWAG (Maddox et al., 2019), and a Gaussian process with an RBF kernel. The SWAG predictive distribution is similar to the predictive distribution of ESS in the PCA subspace, as both methods attempt to approximate the posterior in the subspace containing the principal components of the SGD trajectory; however, ESS operates directly in the subspace, and is not constrained to a Gaussian representation. Variational inference interestingly

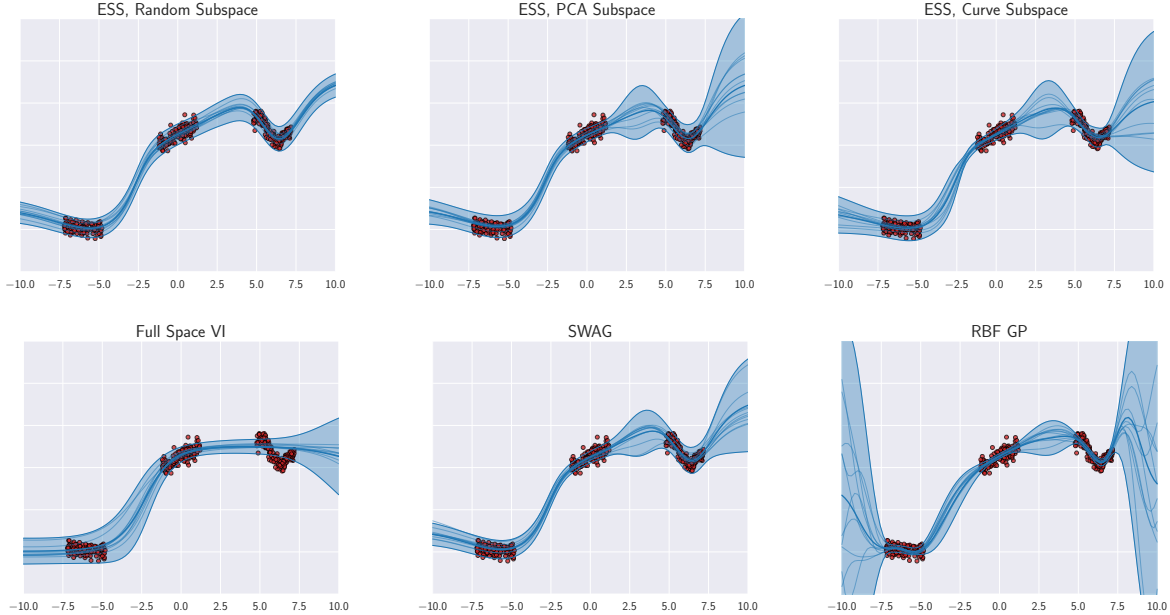


Figure 3: Predictive distribution for visualizing uncertainty in regression. Data (red circles), predictive mean (dark blue line), sample posterior functions (light blue lines),  $\pm 3$  standard deviations about the mean (shaded region). Elliptical slice sampling (ESS) with either a PCA or curve subspace provides uncertainty that intuitively grows away from the data, unlike variational inference in the full parameter space or ESS with a random subspace. This intuitive behaviour matches a GP with an RBF kernel, except the GP is less confident for extrapolation.

appears to *underfit* the data, failing to adapt the posterior mean to structured variations in the observed points, and provides relatively homogenous uncertainty estimates, except very far from the data. We hypothesize that underfitting happens because the KL divergence term between approximate posterior and prior distributions overpowers the data fit term in the variational lower bound objective, as we have many more parameters than data points in this problem (see Blundell et al. (2015) for details about VI). In the bottom right panel of Figure 3 we also consider a Gaussian process (GP) with an RBF kernel — presently the gold standard for regression uncertainty. The GP provides a reasonable predictive distribution in a neighbourhood of the data, but arguably becomes underconfident for extrapolation, with uncertainty quickly blowing up away from the data.

## 5.2 UCI REGRESSION

We next compare our subspace inference methods on UCI regression tasks to a variety of methods for approximate Bayesian inference with neural networks.<sup>4</sup> To measure performance we compute Gaussian test likelihood (details in Appendix E.1.1). We follow convention and parameterize our neural network models so that for an input  $x$

they produce two outputs, predictive mean  $\mu(x, w)$  and predictive variance  $\sigma(x, w)$ . For all datasets we tune the temperature  $T$  in (2) by maximizing the average likelihood over 3 random validations splits (see Appendix F.1 for a discussion of the effect of temperature).

### 5.2.1 Large UCI Regression Datasets

We experiment with 5 large regression datasets from UCI: *elevators*, *keggdirected*, *keggundirected*, *pol*, *protein* and *skillcraft*. We follow the experimental framework of Wilson et al. (2016).

On all datasets except *skillcraft* we use a feedforward network with five hidden layers of sizes [1000, 1000, 500, 50, 2], ReLU activations and two outputs  $\mu(x, w)$  and  $\sigma_w(x)$  parameterizing predictive mean and variance. On *skillcraft*, we use a smaller architecture [1000, 500, 50, 2] like in Wilson et al. (2016). We additionally learn a global noise variance  $s^2$ , so that the predictive variance at  $x$  is  $\sigma^2(x, w, s) = s^2 + \sigma_w^2(x)$ , where  $\sigma_w^2(x)$  is the variance output from the final layer of the network. We use softplus parameterizations for both  $\sigma_w$  and  $s$  to ensure positiveness of the variance, initializing the global variance at  $s^2 = 1$  (the total variance in the dataset).

We compare subspace inference to deep kernel learning (DKL) (Wilson et al., 2016) with a spectral mixture base kernel (Wilson and Adams, 2013), SGD trained

<sup>4</sup>We use the pre-processing from [https://github.com/hughsalimbeni/bayesian\\_benchmarks](https://github.com/hughsalimbeni/bayesian_benchmarks).




Figure 4: Test log-likelihoods for subspace inference and baselines on UCI regression datasets. Subspace inference with PCA achieves as good or better test log-likelihoods compared to SGD and SWAG, and is competitive with DVI (which does not immediately scale to larger problems or networks). We report mean over 20 random splits of the data  $\pm$  one standard deviation. For consistency with the literature, we report normalized log-likelihoods on large datasets (elevators, pol, protein, skillcraft, keggD, keggU; see Section 5.2.1) and unnormalized log-likelihoods on small datasets (naval, concrete, yacht, boston, energy; see Section 5.2.2).

networks with a Bayesian final layer as in Riquelme et al. (2018) and two approximate Gaussian processes: orthogonally decoupled variational Gaussian processes (OrthVGP) (Salimbeni et al., 2018) and Fastfood approximate kernels (FF)<sup>5</sup> (Yang et al., 2015). We show the test log-likelihoods and RMSEs in Appendix Tables 5 and 6, and the 95% credible intervals in Appendix Table 7. We summarize the test log-likelihoods in Figure 4. Subspace inference outperforms SGD by a large margin on elevators, pol and protein, and is competitive on the other datasets. Compared to SWAG, subspace inference typically improves results by a small margin.

Finally, we plot the coverage of the 95% predictive intervals in Figure 10. Again, subspace inference performs at least as well as the SGD baseline, and has substantially better calibration on elevators and protein.

### 5.2.2 Small UCI Regression Datasets

We compare subspace inference to the state-of-the-art approximate BNN inference methods including deterministic variational inference (DVI) (Wu et al., 2019), deep GPs (DGP) with two GP layers trained via expectation propagation (Bui et al., 2016), and variational inference with the re-parameterization trick (Kingma and Welling, 2013). We follow the set-up of Bui et al. (2016) and use a fully-connected network with a single hidden layer with 50 units. We present the test log-likelihoods, RMSEs and test calibration results in Appendix Tables 2, 3 and 4. We additionally visualize test log-likelihoods in Figure 4 and calibrations in Appendix Figure 10.

Our first observation is that SGD provides a surprisingly strong baseline, sometimes outperforming DVI. The subspace methods outperform SGD and DVI on naval, concrete and yacht and are competitive on the other datasets.

### 5.3 IMAGE CLASSIFICATION

Next, we test the proposed method on state-of-the-art convolutional networks on CIFAR datasets. Similarly to Section 5.1, we construct five-dimensional random and five-dimensional PCA subspaces around a trained SWA solution. We also construct a two-dimensional curve subspace by connecting our SWA solution to another independently trained SWA solution. We use the value  $T = 5000$  for temperature in all the image classification experiments with random and PCA subspaces, and  $T = 10000$  in the curve subspace (Section 3.5), from cross-validation using VGG-16 on CIFAR-100.

We visualize the samples from ESS in each of the subspaces in Figure 5. For VI we also visualize the  $3\sigma$ -region of the closed-form approximate posterior in the random and PCA subspaces. As we can see, ESS is able to capture the shape of the posterior in each of the subspaces.

In the PCA subspace we also visualize the SWAG approximate posterior distribution. SWAG overestimates the variance along the first principle component (horizontal axis in Figure 5b) of the SGD trajectory<sup>6</sup>. VI in the subspace is able to provide a better fit for the posterior distribution, as it is directly approximating the posterior while SWAG is approximating the SGD trajectory.

We report the accuracy and negative log-likelihood for each of the subspaces in Table 1. Going from random, to PCA, to curve subspaces provides progressively better results, due to increasing diversity and quality of models within each subspace. In the remaining experiments we use the PCA subspace, as it generally provides good performance at a much lower computational cost than the curve subspace.

We next apply ESS and simple VI in the PCA subspace on VGG-16, PreResNet-164 and WideResNet28x10 on

<sup>5</sup>Results for FF are from Wilson et al. (2016).

<sup>6</sup>See also Figure 5 in Maddox et al. (2019).



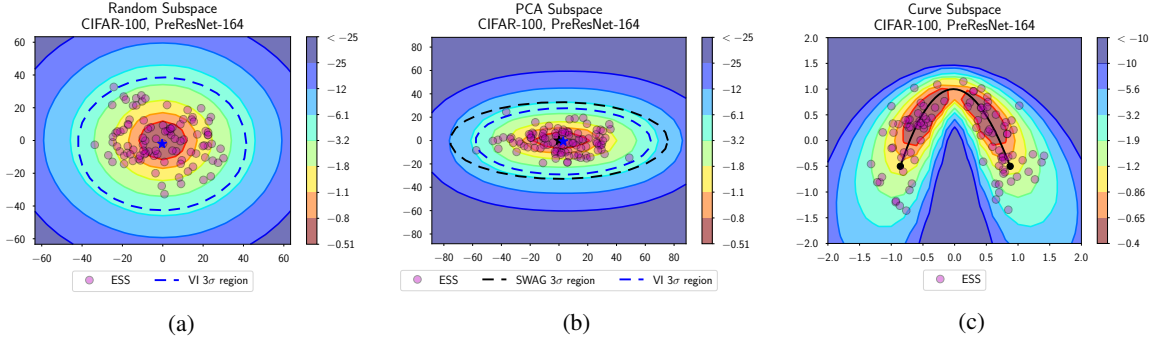


Figure 5: Posterior log-density surfaces, ESS samples (shown with magenta circles), and VI approximation posterior distribution ( $3\sigma$ -region shown with blue dashed line) in (a) random, (b) PCA and (c) curve subspaces for PreResNet-164 on CIFAR-100. In panel (b) the dashed black line shows the  $3\sigma$ -region of the SWAG predictive distribution.

Table 1: Negative log-likelihood and Accuracy for PreResNet-164 for 10-dimensional random, 10-dimensional PCA, and 2-dimensional curve subspaces. We report mean and stdev over 3 independent runs.

	Random	PCA	Curve
NLL	$0.6858 \pm 0.0052$	$0.6652 \pm 0.004$	$0.6493 \pm 0.01$
Accuracy (%)	$80.17 \pm 0.03$	$80.54 \pm 0.13$	$81.55 \pm 0.26$

CIFAR-10 and CIFAR-100. We report the results in Appendix Tables 8, 9. Subspace inference is competitive with SWAG and consistently outperforms most of the other baselines, including MC-dropout (Gal and Ghahramani, 2016), temperature scaling (Guo et al., 2017) and KFAC-Laplace Ritter et al. (2018).

## 6 CONCLUSION

Bayesian methods were once a gold standard for inference with neural networks. Indeed, a Bayesian model average provides an entirely different mechanism for making predictions than standard training, with benefits in accuracy and uncertainty representation. However, efficient and practically useful Bayesian inference has remained a critical challenge for the exceptionally high dimensional parameter spaces in modern deep neural networks. In this paper, we have developed a subspace inference approach to side-step the dimensionality challenges in Bayesian deep learning: by constructing subspaces where the neural network has sufficient variability, we can easily apply standard approximate Bayesian inference methods, with strong practical results.

In particular, we have demonstrated that simple affine subspaces constructed from the principal components of the SGD trajectory contain enough variability for practically effective Bayesian model averaging, often outperforming full parameter space inference techniques. Such subspaces can be surprisingly low dimensional (e.g.,

5 dimensions), and combined with approximate inference methods, such as slice sampling, which are good at automatically exploring posterior densities, but would be entirely intractable in the original parameter space. We further improve performance with subspaces constructed from low-loss curves connecting independently trained solutions (Garipov et al., 2018), with additional computation. Subspace inference is particularly effective at representing growing and shrinking uncertainty as we move away and towards data, which has been a particular challenge for Bayesian deep learning methods.

Crucially, our approach is modular and easily adaptable for exploring both different subspaces and approximate inference techniques, enabling fast exploration for new problems and architectures. There are many exciting directions for future work. It could be possible to explicitly train subspaces to select high variability in the functional outputs of the model, while retaining quality solutions. One could also develop Bayesian PCA approaches (Minka, 2001) to automatically select the dimensionality of the subspace. Moreover, the intuitive behaviour of the regression uncertainty for subspace inference could be leveraged in a number of active learning problems, such as Bayesian optimization and probabilistic model-based reinforcement learning, helping to address the current issues of input dimensionality in these settings.

The ability to perform approximate Bayesian inference in low-dimensional subspaces of deep neural networks is a step towards, scalable, modular, and interpretable Bayesian deep learning.

## Acknowledgements

WJM, PI, PK and AGW were supported by an Amazon Research Award, Facebook Research, and NSF IIS-1563887. WJM was additionally supported by an NSF Graduate Research Fellowship under Grant No. DGE-1650441.

## References

- Bingham, E., Chen, J. P., Jankowiak, M., Obermeyer, F., Pradhan, N., Karaletsos, T., Singh, R., Szerlip, P., Horsfall, P., and Goodman, N. D. (2018). Pyro: Deep universal probabilistic programming. *Journal of Machine Learning Research*.
- Bishop, C. M. (1999). Bayesian PCA. In *Advances in neural information processing systems*, pages 382–388.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. (2015). Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424*.
- Bui, T., Hernández-Lobato, D., Hernandez-Lobato, J., Li, Y., and Turner, R. (2016). Deep gaussian processes for regression using approximate expectation propagation. In *International Conference on Machine Learning*, pages 1472–1481.
- Demmel, J. W. (1997). *Applied Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. (2017). Density estimation using real nvp. In *International Conference on Learning Representations*.
- Foong, A. Y. K., Li, Y., Hernández-Lobato, J. M., and Turner, R. E. (2019). ‘In-Between’ Uncertainty in Bayesian Neural Networks. *arXiv e-prints*, page arXiv:1906.11537.
- Gal, Y. and Ghahramani, Z. (2016). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pages 1050–1059.
- Gardner, J., Pleiss, G., Weinberger, K. Q., Bindel, D., and Wilson, A. G. (2018). Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In *Advances in Neural Information Processing Systems*, pages 7587–7597.
- Garipov, T., Izmailov, P., Podoprikin, D., Vetrov, D. P., and Wilson, A. G. (2018). Loss surfaces, mode connectivity, and fast ensembling of dnns. In *Advances in Neural Information Processing Systems*, volume abs/1802.10026.
- Geyer, C. J. and Thompson, E. A. (1995). Annealing markov chain monte carlo with applications to ancestral inference. *Journal of the American Statistical Association*, 90(431):909–920.
- Ghashami, M., Liberty, E., Phillips, J. M., and Woodruff, D. P. (2016). Frequent directions: Simple and deterministic matrix sketching. *SIAM Journal on Computing*, 45(5):1762–1792.
- Guhaniyogi, R. and Dunson, D. B. (2015). Bayesian compressed regression. *Journal of the American Statistical Association*, 110(512):1500–1514.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. (2017). On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1321–1330. JMLR. org.
- Gur-Ari, G., Roberts, D. A., and Dyer, E. (2019). Gradient descent happens in a tiny subspace. <https://openreview.net/forum?id=ByeTHsAqtX>.
- Ha, D., Dai, A., and Le, Q. V. (2017). Hypernetworks. In *International Conference on Learning Representations*. arXiv preprint arXiv:1609.09106.
- Halko, N., Martinsson, P.-G., and Tropp, J. A. (2011). Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288.
- Hoffman, M. D., Blei, D. M., Wang, C., and Paisley, J. (2013). Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347.
- Hoffman, M. D. and Gelman, A. (2014). The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *Journal of Machine Learning Research*, 15(1):1593–1623.
- Huggins, J., Campbell, T., and Broderick, T. (2016). Coresets for scalable bayesian logistic regression. In *Advances in Neural Information Processing Systems*, pages 4080–4088.
- Izmailov, P., Podoprikin, D., Garipov, T., Vetrov, D., and Wilson, A. G. (2018). Averaging weights leads to wider optima and better generalization. In *Uncertainty in Artificial Intelligence*.
- Karaletsos, T., Dayan, P., and Ghahramani, Z. (2018). Probabilistic meta-representations of neural networks. arXiv preprint arXiv:1810.00555.
- Kingma, D. P., Salimans, T., and Welling, M. (2015). Variational dropout and the local reparameterization trick. In *Advances in Neural Information Processing Systems*, pages 2575–2583.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Krueger, D., Huang, C.-W., Islam, R., Turner, R., Lacoste, A., and Courville, A. (2017). Bayesian hypernetworks. *arXiv preprint arXiv:1710.04759*.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, pages 6402–6413.

- Li, C., Farkhoor, H., Liu, R., and Yosinski, J. (2018a). Measuring the intrinsic dimension of objective landscapes. In *International Conference on Learning Representations*. arXiv preprint arXiv:1804.08838.
- Li, H., Xu, Z., Taylor, G., Studer, C., and Goldstein, T. (2018b). Visualizing the Loss Landscape of Neural Nets. In *Advances in Neural Information Processing Systems*. arXiv: 1712.09913.
- MacKay, D. J. (2003). *Information theory, inference and learning algorithms*. Cambridge university press.
- Maddox, W., Garipov, T., Izmailov, P., Vetrov, D., and Wilson, A. G. (2019). A simple baseline for bayesian uncertainty in deep learning. *arXiv preprint arXiv:1902.02476*.
- Minka, T. P. (2001). Automatic choice of dimensionality for pca. In *Advances in neural information processing systems*, pages 598–604.
- Murray, I., Prescott Adams, R., and MacKay, D. J. (2010). Elliptical slice sampling. In *Artificial Intelligence and Statistics*.
- Neal, R. M. (1996a). *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media.
- Neal, R. M. (1996b). Sampling from multimodal distributions using tempered transitions. *Statistics and computing*, 6(4):353–366.
- Neal, R. M. et al. (2003). Slice sampling. *The annals of statistics*, 31(3):705–767.
- Neal, R. M. et al. (2011). Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2.
- Patra, S. and Dunson, D. B. (2018). Constrained bayesian inference through posterior projections. *arXiv preprint arXiv:1812.05741*.
- Pradier, M. F., Pan, W., Yao, J., Ghosh, S., and Doshi-velez, F. (2018). Projected BNNs: Avoiding weight-space pathologies by learning latent representations of neural network weights. *arXiv preprints*, page arXiv:1811.07006.
- Riquelme, C., Tucker, G., and Snoek, J. (2018). Deep bayesian bandits showdown. In *International Conference on Learning Representations*.
- Ritter, H., Botev, A., and Barber, D. (2018). A scalable laplace approximation for neural networks. In *International Conference on Learning Representations*.
- Roweis, S. T. (1998). EM algorithms for PCA and SPCA. In *Advances in neural information processing systems*, pages 626–632.
- Salimbeni, H., Cheng, C.-A., Boots, B., and Deisenroth, M. (2018). Orthogonally decoupled variational gaussian processes. In *Advances in Neural Information Processing Systems*, pages 8725–8734.
- Silva, R. and Kalaitzis, A. (2015). Bayesian inference via projections. *Statistics and Computing*, 25(4):739–753.
- Titsias, M. K. (2017). Learning model reparametrizations: Implicit variational inference by fitting mcmc distributions. *arXiv preprint arXiv:1708.01529*.
- Watanabe, S. (2007). Almost all learning machines are singular. In *2007 IEEE Symposium on Foundations of Computational Intelligence*, pages 383–388. IEEE.
- Watanabe, S. (2013). A widely applicable bayesian information criterion. *Journal of Machine Learning Research*, 14(Mar):867–897.
- Wilson, A. and Adams, R. (2013). Gaussian process kernels for pattern discovery and extrapolation. In *International Conference on Machine Learning*, pages 1067–1075.
- Wilson, A. G., Hu, Z., Salakhutdinov, R., and Xing, E. P. (2016). Deep kernel learning. In *Artificial Intelligence and Statistics*, pages 370–378.
- Wu, A., Nowozin, S., Meeds, E., Turner, R. E., Hernández-Lobato, J. M., and Gaunt, A. L. (2019). Deterministic Variational Inference for Robust Bayesian Neural Networks. In *International Conference on Learning Representations*. arXiv:1810.03958.
- Yang, Z., Wilson, A., Smola, A., and Song, L. (2015). A la carte—learning fast kernels. In *Artificial Intelligence and Statistics*, pages 1098–1106.
- Yao, J., Pan, W., Ghosh, S., and Doshi-Velez, F. (2019). Quality of Uncertainty Quantification for Bayesian Neural Network Inference. *arXiv preprints*, page arXiv:1906.09686.