

---

# On permutation symmetries in Bayesian neural network posteriors: a variational perspective

---

**Simone Rossi**  
Stellantis, France

**Ankit Singh**  
Stellantis, India

**Thomas Hannagan**  
Stellantis, France

## Abstract

The elusive nature of gradient-based optimization in neural networks is tied to their loss landscape geometry, which is poorly understood. However recent work has brought solid evidence that there is essentially no loss barrier between the local solutions of gradient descent, once accounting for weight-permutations that leave the network’s computation unchanged. This raises questions for approximate inference in Bayesian neural networks (BNNs), where we are interested in marginalizing over multiple points in the loss landscape. In this work, we first extend the formalism of marginalized loss barrier and solution interpolation to BNNs, before proposing a matching algorithm to search for linearly connected solutions. This is achieved by aligning the distributions of two independent approximate Bayesian solutions with respect to permutation matrices. We build on the results of Ainsworth et al. (2023), reframing the problem as a combinatorial optimization one, using an approximation to the sum of bilinear assignment problem. We then experiment on a variety of architectures and datasets, finding nearly zero marginalized loss barriers for linearly connected solutions.

## 1 Introduction

Throughout the last decade, deep neural networks (DNNs) have achieved significant success in a wide range of practical applications, becoming the fundamental ingredient for e.g., computer vision [e.g., 57, 25, 41, 64], language models [e.g., 24, 82, 17] and generative models [e.g., 54, 96, 97, 98, 102, 34]. Despite recent important advancements, understanding the loss landscape of DNNs is still challenging. The characterization of its highly non-convex nature, its relation with architectural choices like depth and width and the connection with optimization and generalization are just some of the problems which have been the focus of extensive research in the last few years [e.g., 76, 26, 36, 32, 2, 30, 38, 79]. It is well known, for example, that one of the fundamental characteristics of deep neural networks is their ability to learn hierarchical features, and in this regards deeper networks seem to be exponentially more expressive than shallower models [e.g., 3, 5, 7, 8, 18, 109], leading the loss landscape to have many optima due to symmetries and over-parameterization [76, 26, 111, 94]. At the same time, the role of the depth of a model in relation with its width is far less understood [80], despite wide neural networks exhibiting important theoretical properties in their infinite limit behavior [e.g., 73, 23, 29, 53, 48, 37, 20].

Two notions that have been useful to shed light on the geometry of loss landscapes are that of *loss barriers* and *mode connectivity* [36, 26]. The mode connectivity hypothesis states that given two points in the landscape, there exists a path connecting them such that the loss is constant or near constant (or, said differently, the loss barrier is null). We refer to *linear mode connectivity* when the path connecting the two solutions is linear [32]. Recently, evidence has surfaced that stochastic gradient descent (SGD) solutions to the loss minimization problem can be linearly connected. Indeed, Entezari et al. [30] discuss the role of permutation symmetries from a loss connectivity viewpoint, conjecturing the possibility that mode connectivity is actually linear once accounting for all permutation invariances.

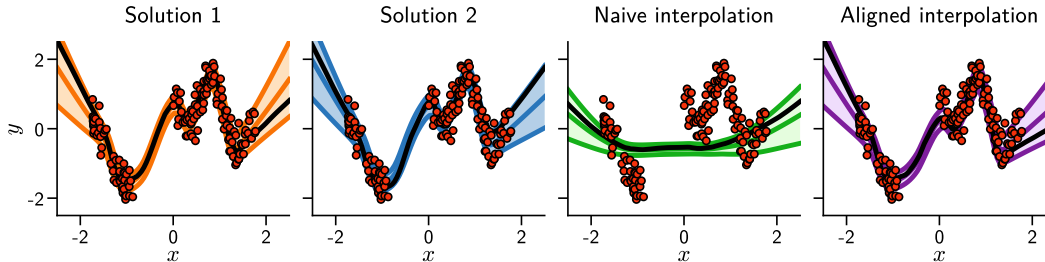


Figure 1: **Permutation symmetries for regression on the Snelson dataset [95].** (Left) Two different solutions, with similar function-space behavior (showing  $\mu \pm 2\sigma$ ). (Right) Functions obtained when using two different strategies to interpolate between solutions. When we align the solutions, by taking into account the permutation symmetries, we retain the capability to model the data, indicating that there are solutions that once linearly interpolated exhibit no loss in performance. Note that, in weight space, the solution found with alignment is neither equal to solution 1 nor 2 (see the black curve which is a single function sample with fixed randomness).

Additionally, [2] gathers compelling empirical evidence across several network architectures and tasks, that under such a permutation symmetry the loss landscape often contains a single, nearly convex basin.

In this work, we are taking a different perspective on this analysis. We are interested in the Bayesian treatment of neural networks, which results in a natural form of regularization and allows to reason about uncertainty in the predictions [101, 73, 66]. Bayesian inference for deep neural networks is notoriously challenging, as we wish to marginalize over multi-modal distributions with high dimensionality [47]. For this reason, there are various ways to approximate the posterior, involving techniques like variational inference [39, 14, 35, 62, 77], Markov chain Monte Carlo (MCMC) methods [75, 72, 27], possibly with stochastic gradients [19, 112, 33, 105, 70] and the Laplace approximation [68, 68, 84]. Indeed, fundamentally the Bayesian posterior and the loss landscapes are tightly interconnected: (i) solutions to the loss minimization problem are equivalent to maximum-a-posteriori (MAP) solutions, (ii) the loss landscape is equivalent to the un-normalized negative log-posterior. While in theory, given a dataset, the posterior is unique and the solution is global, many approximations will only explore local properties of the true posterior<sup>1</sup>. It’s worth noting that a posterior over the parameters of the neural network induces a posterior on the functions generated by the model. Permutation symmetries play an important role in the geometry of the weight-space posterior, which are generally not reflected in function-space. While it is possible to carry out inference directly in function space, this poses a number of challenges [99, 71, 89, 61]. Fig. 1 illustrates this situation for a regression task on the Snelson dataset using a 3-layer DNN: on the left we compare two (approximate) solutions which have different weight-space behavior but similar function-space behavior. Notably, when we interpolate these two solutions (Fig. 1 on the right), we completely lose all capability of modeling the data. However, when we account for permutation symmetries in the posterior, we end up with solutions that once interpolated are still good approximations. This suggests that for any weight-space distribution, there exists a class of solutions which are functionally equivalent and linearly connected. This example motivates an informal generic conjecture:

*Solutions of approximate Bayesian inference for neural networks are linearly connected after accounting for functionally equivalent permutations.*

While being similar to the one in [30, 2], if this conjecture was to hold true for approximate Bayesian neural networks (BNNs) it would represent an important step in further characterizing the properties of the Bayesian posterior and the effect of various approximations. We purposely leave the previous conjecture broadly open regarding the choice of the approximation method to allow for a more general discussion. More specifically, in this paper we will analyze and focus our discussion on the variational inference framework, making a more specific conjecture:

**Conjecture 1.** *Variational inference solutions for approximate Bayesian inference in neural networks are linearly connected after accounting for functionally equivalent permutations.*

<sup>1</sup>By local properties, we mean that despite theoretical convergence guarantees of many methods like variational inference and MCMC, in practice the true posterior for deep neural networks is still highly elusive; see for instance the empirical convergence analysis of Hamiltonian Monte Carlo (HMC) in [47].

**Contributions.** With this work, we aim at studying the linear connectivity properties of approximate solutions to the Bayesian inference problem and we make several contributions. (i) We extend the formalism of loss barrier and solution interpolation to BNNs. (ii) For the variational inference setting, propose a matching algorithm to search for linearly connected solutions by aligning the distributions of two independent solutions with respect to permutation matrices. Inspired by [2], we frame the problem as a combinatorial optimization problem using approximation to the linear sum assignment problem. (iii) We then experiment on a variety of architectures and datasets, finding nearly zero-loss barriers for linearly connected solutions. In Fig. 2 we present a sneak-peek and a visualization of our findings, where we show that after weight distribution alignment we can find a permutation map  $P$  of the solution  $q_1$  such that it can be linearly connected through high density regions to  $q_0$ .

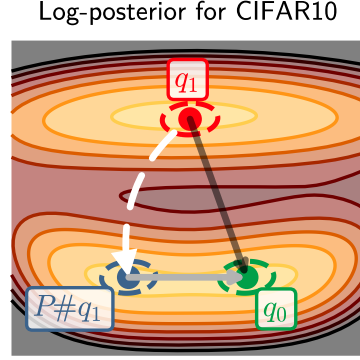


Figure 2: **Permutations in multi-modal posterior.** Log-posterior for MLP/CIFAR10, showing the two solutions ( $q_0$  and  $q_1$ ) for which we can find a permutation map such that  $P\#q_1$  can be linearly connected to  $q_0$  with low barrier (brighter regions).

## 2 Preliminaries on Bayesian deep learning

In this section, we review some basic notations on BNNs and we review stochastic variational inference (SVI), which is the main approximation method that we analyze in this paper. Let’s consider a generic multilayer perceptron (MLP) with  $L$  layers, where the output of the  $l$ -th layer  $f_l(\theta_l, \mathbf{x})$  is a vector-valued function of the previous layer output  $f_{l-1}$  as follows,

$$f_l(\theta_l, \mathbf{x}) = \mathbf{W}_l a(f_{l-1}(\theta_{l-1}, \mathbf{x})) + \mathbf{b}_l \quad (1)$$

where  $a(\cdot)$  is a non-linearity,  $\mathbf{W}_l$  is a  $D_l \times D_{l-1}$  weight matrix and  $\mathbf{b}_l$  the corresponding bias vector. We shall refer to the parameters of the layer  $l$  as  $\theta_l = \{\mathbf{b}_l, \mathbf{W}_l\}$ , and the union of all trainable parameters as  $\theta = \{\theta_l\}_{l=1}^L$ .

The objective of using Bayesian inference on deep neural networks [67, 65] involves inferring a posterior distribution over the parameters of the neural network given the available dataset  $\{\mathbf{X}, \mathbf{Y}\} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ . This requires choosing a likelihood and a prior [73, 74, 103]:

$$p(\theta | \mathbf{Y}, \mathbf{X}) = Z^{-1} p(\mathbf{Y} | \theta, \mathbf{X}) p(\theta) \quad (2)$$

where the normalization constant  $Z$  is the marginal likelihood  $p(\mathbf{Y} | \mathbf{X})$ . As usually done, we assume that the likelihood factorizes over observations, i.e.  $p(\mathbf{Y} | \theta, \mathbf{X}) = \prod_{i=1}^N p(\mathbf{y}_i | \theta, \mathbf{x}_i)$ .

Bayesian deep learning is intractable due to the non-conjugacy likelihood-prior and thus we don’t have access to closed form solutions. Variational inference (VI) is a common technique to handle intractable Bayesian neural networks [12, 43, 39, 50]. Let  $\mathcal{P}(\mathbb{R}^d)$  be the space of probability measures on  $\mathbb{R}^d$ ; VI reframes the inference problem into an optimization one, commonly by introducing a parameterized distribution  $q(\theta) \in \mathcal{P}(\mathbb{R}^d)$  which is optimized to minimize the Kullback-Leibler (KL) divergence with respect to the true posterior  $p(\theta | \mathbf{Y}, \mathbf{X})$ . In practice, this involves the maximization of the evidence lower bound (ELBO) defined as

$$\mathcal{L}_{\text{ELBO}}(q) \stackrel{\text{def}}{=} \int \log p(\mathbf{Y} | \theta, \mathbf{X}) dq(\theta) - \text{KL}[q(\theta) \| p(\theta)] \quad (3)$$

whose gradients can be unbiasedly estimated with mini-batches of data [43] and the reparameterization trick [54, 55]. Despite its simple formulation, the optimization of the ELBO hides several challenges, like the initialization of the variational parameters [87], the effects of over-parameterization on the quality of the approximation [88, 44, 58]. Here we are interested in how different solutions to Eq. (3) relate to each other in terms of loss barrier, which we will define formally in the following section.

## 3 Loss barriers

In the context of Bayesian inference, we are interested in the loss computed by marginalization of the model parameters with respect to (an approximation of) the posterior. As such, we use the predictive

likelihood, a proper scoring method for probabilistic models [83], defined as

$$\log p(\mathbf{y}^* | \mathbf{x}^*) = \log \int p(\mathbf{y}^* | \boldsymbol{\theta}, \mathbf{x}^*) p(\boldsymbol{\theta} | \mathbf{Y}, \mathbf{X}) d\boldsymbol{\theta} \approx \log \int p(\mathbf{y}^* | \boldsymbol{\theta}, \mathbf{x}^*) dq(\boldsymbol{\theta}) \quad (4)$$

where  $q(\boldsymbol{\theta})$  is an approximation of the true posterior (parametric or otherwise),  $\{\mathbf{x}^*, \mathbf{y}^*\}$  are respectively the input and its corresponding label a data point under evaluation. To keep the notation uncluttered for the remaining of the paper, we write the predictive likelihood computed for a set of points  $\{\mathbf{x}_i^*, \mathbf{y}_i^*\}_{i=1}^N$  as a functional  $\mathcal{L} : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ , defined as

$$\mathcal{L}(q) \stackrel{\text{def}}{=} \sum_{i=1}^N \log \int p(\mathbf{y}_i^* | \boldsymbol{\theta}, \mathbf{x}_i^*) dq(\boldsymbol{\theta}) \quad (5)$$

Let's assume two models trained with VI with two different initializations, random seeds, and batch ordering. Variational inference in the classic inverse sense  $\text{KL}[q \| p]$  is mode seeking, thus we expect the two runs to converge to different solutions, say  $q_0$  and  $q_1$ . To test the loss barrier as we interpolate between the two solutions we need to decide on the interpolation rule. We decide to interpolate the solutions following the Wasserstein geodesics between  $q_0$  and  $q_1$ . First, let's start with a few definitions. Let  $q \in \mathcal{P}(\mathbb{R}^d)$  be a probability measure on  $\mathbb{R}^d$  and  $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$  a measurable map; we denote  $T_{\#}q$  the *push-forward measure* of  $q$  through  $T$ . Now we can introduce the *Wasserstein geodesics*, as follows.

**Definition 1.** *The Wasserstein geodesics between  $q_0$  and  $q_1$  is defined as the path*

$$q_\tau = ((1 - \tau)\text{Id} + \tau T_{q_0}^{q_1})_{\#} q_0, \quad \tau \in [0, 1] \quad (6)$$

where  $\text{Id}$  is the identity map and  $T_{q_0}^{q_1}$  is the optimal transport map between  $q_0$  and  $q_1$ , which for Brenier's theorem [16], is unique.

While we could interpolate using a mixture of the two solutions, we argue that this choice is trivial and does not fully give us a picture of the underlying loss landscape. Indeed, Eq. (6) is fundamentally different from a naive mixture path  $\tilde{q}_\tau = (1 - \tau)q_0 + \tau q_1$ . In case of Gaussian distributions, when  $q_0 = \mathcal{N}(\mathbf{m}_0, \mathbf{S}_0)$  and  $q_1 = \mathcal{N}(\mathbf{m}_1, \mathbf{S}_1)$ ,  $q_\tau$  is Gaussian as well [100] with mean and covariance computed as follows:

$$\begin{aligned} \mathbf{m}_\tau &= (1 - \tau)\mathbf{m}_1 + \tau\mathbf{m}_2 \\ \mathbf{S}_\tau &= \mathbf{S}_1^{-1/2} \left( (1 - \tau)\mathbf{S}_1 + \tau \left( \mathbf{S}_1^{1/2} \mathbf{S}_2 \mathbf{S}_1^{1/2} \right)^{1/2} \right)^2 \mathbf{S}_1^{-1/2} \end{aligned} \quad (7)$$

which simplifies even further when the covariances are diagonal.

Now, we can define convexity along Wasserstein geodesics [4] as follows.

**Definition 2.** *Let  $\mathcal{L} : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ ,  $\mathcal{L}$  is  $\lambda$  geodesics convex with  $\lambda > 0$  if for any  $q_0, q_1 \in \mathcal{P}(\mathbb{R}^d)$  it holds that*

$$\mathcal{L}(q_\tau) \leq (1 - \tau)\mathcal{L}(q_0) + \tau\mathcal{L}(q_1) - \frac{\lambda\tau(1 - \tau)}{2} \mathcal{W}_2^2(q_0, q_1) \quad (8)$$

where  $\mathcal{W}_2^2(q_0, q_1)$  is the Wasserstein distance defined as [104, 51, 52]

$$\mathcal{W}_2^2(q_1, q_0) = \inf_{\gamma \in \Pi(q_1, q_0)} \int \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0\|_2^2 d\gamma(\boldsymbol{\theta}_1, \boldsymbol{\theta}_0) \quad (9)$$

with  $\Pi(\cdot, \cdot)$  being the space of measure with  $q_0$  and  $q_1$  as marginals.

While mathematically proving the geodesics convexity of the predictive likelihood for arbitrary architectures and densities is currently beyond the scope of this work, we can empirically define a proxy using the *functional loss barrier*, defined as follows.

**Definition 3.** *The functional loss barrier along the Wasserstein geodesics from  $q_0$  and  $q_1$  is defined as the highest difference between the marginal loss computed when interpolating two solutions  $q_0$  and  $q_1$  and the linear interpolation of the loss at  $q_0$  and  $q_1$ :*

$$\mathcal{B}(q_0, q_1) = \max_{\tau} \mathcal{L}(q_\tau) - ((1 - \tau)\mathcal{L}(q_0) + \tau\mathcal{L}(q_1)) \quad (10)$$

where  $q_\tau$  follows the definition in Eq. (6).

This definition is a more general than the ones in [2, 30, 32] but we can recover [30] by assuming delta posteriors  $q_i = \delta(\theta - \theta_i)$  and we can further recover [2, 32] by also assuming  $\mathcal{L}(q_0) = \mathcal{L}(q_1)$ .

**A comment on mixtures.** In previous paragraphs, we argued that the mixture of distributions is not sufficient to capture the underlying complex geometry of the posterior. Now, we want to better illustrate this choice with a simple example. In Fig. 3 we plot the test likelihood with two interpolation strategies between two solutions (MLP on CIFAR10): the Wasserstein geodesics and the mixture. With mixtures, we see that the likelihood is pretty much constant during the interpolation, but this is very miss-leading: we don't see barriers not because they don't exist, but because the mixture simply re-weights the distributions, without continuously transporting mass in the parameter space.

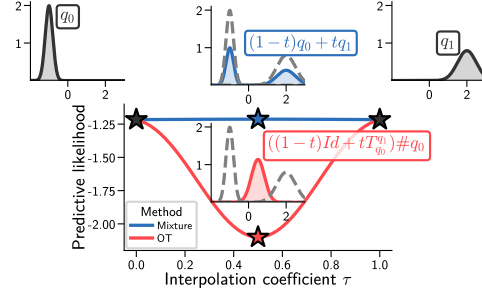


Figure 3: **Wasserstein geodesics and mixtures.** Test likelihood for mixture and the Wasserstein geodesics interpolation. Solutions are MLPs trained on CIFAR10.

#### 4 Aligning distributions by looking for permutation symmetries

In this section, we formalize the algorithm that aligns the solutions of Bayesian inference through permutation symmetries of weight matrices and biases. Let  $\mathbb{S}(d)$  be the set of valid  $d \times d$  permutation matrices. Given a generic distribution  $q(\theta)$ , we can apply a permutation matrix  $P \in \mathbb{S}(D_l)$  to a hidden layer output at layer  $l$ , and if we define  $\theta'$  to be equivalent to  $\theta$  with the exception of

$$W'_l = PW_l, \quad b'_l = Pb_l, \quad W'_{l+1} = W_{l+1}P^\top, \quad (11)$$

then  $P\#q$  is the equivalent push-forward distribution for  $\theta'$ , where  $P$  is the associated permutation map. Let us define the distribution over the functional output of the model as

$$q(\mathbf{f}(\theta, \cdot)) = \int \delta(\mathbf{f}(\theta, \cdot) - \mathbf{f}(\hat{\theta}, \cdot)) d q(\hat{\theta}), \quad (12)$$

and, equivalently, the distribution on the function using the permuted parameters as

$$q(\mathbf{f}(\theta', \cdot)) = \int \delta(\mathbf{f}(\theta', \cdot) - \mathbf{f}(\hat{\theta}', \cdot)) d P\#q(\hat{\theta}'), \quad (13)$$

where in both cases  $\delta(\cdot)$  is the Dirac function. Then, it is simple to verify that the two models are functionally equivalent for any inputs,

$$q(\mathbf{f}(\theta, \cdot)) = q(\mathbf{f}(\theta', \cdot)). \quad (14)$$

This implies that for any weight-space distribution  $q$ , there exists a class of functionally equivalent solutions  $P\#q$ , in the sense of Eq. (14). These same considerations can be easily extended to other layers, by considering multiple permutation matrices  $P_l$ . For our analysis, given two solutions  $q_0$  and  $q_1$  we are interested in finding the permuted distribution  $P\#q_1$ , functionally equivalent to  $q_1$ , in such a way that once interpolating using Eq. (6) we observe similar performance to  $q_0$  and  $q_1$ . Formally, we can write

$$\arg \min_P \mathcal{D}(q_1(\mathbf{f}(\theta', \cdot)), q_0(\mathbf{f}(\theta, \cdot))) = \arg \min_P \mathcal{D}(P\#q_1(\theta), q_0(\theta)), \quad (15)$$

where  $\mathcal{D}$  is a generic measure of discrepancy.<sup>2</sup>

##### 4.1 Problem setup for permutation of vectors

We start from a single vector of parameters, disregarding for the moment the functional equivalence constraint. We will extend these results to matrices and multiple layers later. In practice considering

<sup>2</sup>To be formally correct, the l.h.s. is a discrepancy defined on stochastic processes while the r.h.s. is defined on random vectors. Under mild assumptions on the distribution on the parameters and the architectures,  $\mathcal{D}$  is well defined in both cases [see e.g., 81].

Gaussian distributions, we know that if  $q = \mathcal{N}(\mathbf{m}, \mathbf{S})$ , then  $P_{\#}q = \mathcal{N}(\mathbf{P}\mathbf{m}, \mathbf{P}\mathbf{S}\mathbf{P}^\top)$  and if  $q = \mathcal{N}(\mathbf{m}, \text{diag}(\mathbf{s}^2))$  then  $P_{\#}q = \mathcal{N}(\mathbf{P}\mathbf{m}, \text{diag}(\mathbf{P}\mathbf{s}^2))$  [11]. With the KL divergence  $\text{KL}[P_{\#}q_1 \parallel q_0]$  it's easy to verify that it leads to just a distance between means, disregarding any covariance information. While certainly this represents a valid choice, we argue that we can find a better solution by using the Wasserstein distance. For Gaussian measures, the Wasserstein distance has analytic solution:

$$\begin{aligned} \mathcal{W}_2^2(q_1, q_0) &= \|\mathbf{m}_0 - \mathbf{m}_1\|_2^2 + \text{Tr}\left(\mathbf{S}_1 + \mathbf{S}_0 - 2\left(\mathbf{S}_1^{1/2}\mathbf{S}_0\mathbf{S}_1^{1/2}\right)^{1/2}\right) = \\ &= \|\mathbf{m}_0 - \mathbf{m}_1\|_2^2 + \|\mathbf{S}_0 - \mathbf{S}_1\|_F^2, \end{aligned} \quad (16)$$

where  $\|\cdot\|_F$  denotes the Frobenius norm,  $\|\mathbf{A}\| = \sum_{ij} a_{ij}^2$ , and where the second line is valid only if the covariances commute ( $\mathbf{S}_1\mathbf{S}_0 = \mathbf{S}_0\mathbf{S}_1$ ). In our case, then, we can simplify as follows:

$$\mathcal{W}_2^2(P_{\#}q_1, q_0) = \|\mathbf{m}_0 - \mathbf{P}\mathbf{m}_1\|_2^2 + \|\mathbf{s}_0 - \mathbf{P}\mathbf{s}_1\|_2^2. \quad (17)$$

To summarize, the problem now can be written as:

$$\arg \min_{\mathbf{P} \in \mathbb{S}(d)} \|\mathbf{m}_0 - \mathbf{P}\mathbf{m}_1\|_2^2 + \|\mathbf{s}_0 - \mathbf{P}\mathbf{s}_1\|_2^2 = \arg \max_{\mathbf{P} \in \mathbb{S}(d)} \langle \mathbf{P} | \mathbf{m}_0\mathbf{m}_1^\top + \mathbf{s}_0\mathbf{s}_1^\top \rangle_F, \quad (18)$$

where the expression  $\langle \mathbf{A} | \mathbf{B} \rangle_F$  is the Frobenius inner product,  $\langle \mathbf{A} | \mathbf{B} \rangle_F = \sum_{ij} A_{ij}B_{ij}$ . Note that the r.h.s. of Eq. (18) is a valid instantiation of the linear assignment problem (LAP) [10], which can be solved in polynomial time.

## 4.2 From vectors to neural network parameters

Finally, we need to take into account that we have multiple layers and weight matrices, and that we are trying to find functionally equivalent solutions. For this, we decide to explicitly change our main objective by enforcing the functional equivalence constraint as follows:

$$\arg \min_{\{P_i\}} \mathcal{W}_2^2\left(P_{1\#}q_1^{(1)}, q_0^{(1)}\right) + \mathcal{W}_2^2\left((P_2 \circ P_1^\top)_{\#}q_1^{(2)}, q_0^{(2)}\right) + \dots + \mathcal{W}_2^2\left((P_{L-1}^\top)_{\#}q_1^{(L)}, q_0^{(L)}\right),$$

where the notation  $(P_l \circ P_{l-1}^\top)$  represents the composition of the two permutation maps applied to rows and columns of the random weight matrices. More conveniently, this can be rewritten in terms of means and standard deviations. To leave the notation uncluttered, let's collect the means and the standard deviations for the layer  $l$  in  $\mathbf{M}^{(l)}$  and  $\mathbf{S}^{(l)}$ , which are now both  $D_l \times D_{l-1}$  matrices, so that  $q^{(l)} = \prod_{ij} \mathcal{N}(M_{ij}^{(l)}, S_{ij}^{(l)})$ . Now we can write,

$$\begin{aligned} \arg \max_{\{P_i\}_{i=1}^L} &\left\langle \mathbf{M}_0^{(1)} | \mathbf{P}_1 \mathbf{M}_1^{(1)} \right\rangle_F + \left\langle \mathbf{S}_0^{(1)} | \mathbf{P}_1 \mathbf{S}_1^{(1)} \right\rangle_F + \left\langle \mathbf{M}_0^{(2)} | \mathbf{P}_2 \mathbf{M}_1^{(2)} \mathbf{P}_1^\top \right\rangle_F + \left\langle \mathbf{S}_0^{(2)} | \mathbf{P}_2 \mathbf{S}_1^{(2)} \mathbf{P}_1^\top \right\rangle_F + \\ &+ \dots + \left\langle \mathbf{M}_0^{(L)} | \mathbf{M}_1^{(L)} \mathbf{P}_{L-1}^\top \right\rangle_F + \left\langle \mathbf{S}_0^{(L)} | \mathbf{S}_1^{(L)} \mathbf{P}_{L-1}^\top \right\rangle_F. \end{aligned}$$

This optimization problem is more challenging than the one presented in Eq. (18): we are interested in finding permutation matrices to be applied concurrently to rows and columns of both means and standard deviations. This class of problems, also known as sum of bilinear assignment problems (SOLAP), is NP-hard and no polynomial-time solutions exist. For this reason, we propose to use the setup in Ainsworth et al. [2] by extending it to our problem. In particular, by fixing all matrices with the exception of  $\mathbf{P}_l$ , we observe that also in our case the problem can be reduced to a classic LAP.

$$\begin{aligned} &\arg \max_{\mathbf{P}_l} \left\langle \mathbf{M}_0^{(l)} | \mathbf{P}_l \mathbf{M}_1^{(l)} \mathbf{P}_{l-1}^\top \right\rangle_F + \left\langle \mathbf{M}_0^{(l+1)} | \mathbf{P}_{(l+1)} \mathbf{M}_1^{(l+1)} \mathbf{P}_l^\top \right\rangle_F + \\ &\quad \left\langle \mathbf{S}_0^{(l)} | \mathbf{P}_l \mathbf{S}_1^{(l)} \mathbf{P}_{l-1}^\top \right\rangle_F + \left\langle \mathbf{S}_0^{(l+1)} | \mathbf{P}_{(l+1)} \mathbf{S}_1^{(l+1)} \mathbf{P}_l^\top \right\rangle_F = \\ &= \arg \max_{\mathbf{P}_l} \left\langle \mathbf{P}_l | \mathbf{M}_0^{(l)} \mathbf{P}_{l-1} \left(\mathbf{M}_1^{(l)}\right)^\top + \left(\mathbf{M}_0^{(l+1)}\right)^\top \mathbf{P}_{l+1} \mathbf{M}_1^{(l+1)} + \right. \\ &\quad \left. \mathbf{S}_0^{(l)} \mathbf{P}_{l-1} \left(\mathbf{S}_1^{(l)}\right)^\top + \left(\mathbf{S}_0^{(l+1)}\right)^\top \mathbf{P}_{l+1} \mathbf{S}_1^{(l+1)} \right\rangle_F. \end{aligned} \quad (19)$$

As discussed in [2], going through each layer, and greedily selecting its best  $\mathbf{P}_l$ , leads to a coordinate descent algorithm which guarantees to end in finite time. We present a pseudo-code in Algorithm 1.

---

**Algorithm 1:** Algorithm to align variational inference solutions

---

**Data:** Variational inference solutions  $q_0$  and  $q_1$ **Result:** Permutation matrices  $P_i$ 

```
1  $P_i \leftarrow I, \forall i \in \{1, \dots, L\};$   
2 while not converged do  
3   for  $i \in \text{RandomPerm}(1, \dots, L - 1)$  do  
4      $P_i = \arg \max \langle P_i | \star \rangle$ , where  $\star$  is the r.h.s. of Eq. (19)
```

---

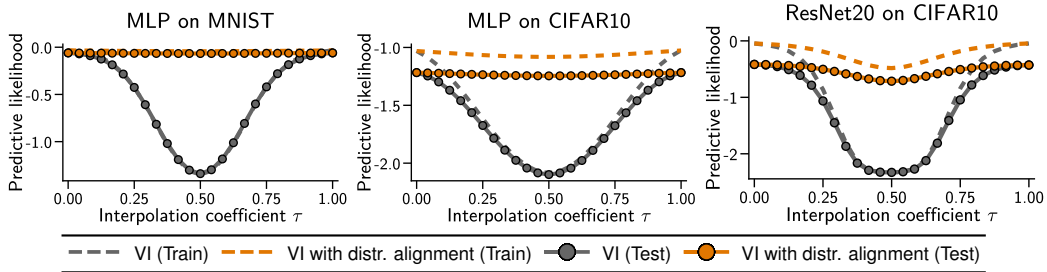


Figure 4: **Zero barrier solutions.** Comparison of loss barriers for standard VI (gray) and VI with alignment (orange). While loss barriers always appear between two solutions in the standard VI approach, in the case of VI with alignment there is no noticeable loss barrier for MLPs and a nearly-zero loss barrier for ResNet20.

## 5 Experiments

Now, we present some supporting evidence to [Conjecture 1](#). We start by training two replicas of BNNs with variational inference (we refer to the Appendix for additional details on the experimental setup). We then compute the marginalized barrier as  $\mathcal{B}(q_0, q_1) = \max_{\tau} \mathcal{L}(q_{\tau}) - ((1 - \tau)\mathcal{L}(q_0) + \tau\mathcal{L}(q_1))$  where  $\mathcal{L}(\cdot)$  is the predictive likelihood and  $\tau \in [0, 1]$ , from which we take 25 evenly distributed points. In particular, we seek to understand what happens to the VI solutions first for the naive interpolation from  $q_0$  and  $q_1$ , and then for the interpolation after aligning  $q_0$  and  $P_{\#}q_1$ . We experiment with MLPs with three layers and ResNet20 [41] with various widths on MNIST [60], Fashion-MNIST [108] and CIFAR10 [56]. All models are trained without data augmentation [106] and with filter response normalization (FRN) layers instead of BatchNorm. Finally, we set the prior to be Gaussian  $\mathcal{N}(\mathbf{0}, \alpha^2 \mathbf{I})$ , with the flexibility of choosing the variance.

### 5.1 Low-barrier interpolations

Fig. 4 shows the results with and without alignment. We see that regardless of the dataset and the model used, the performance degrades significantly when we move between the two solutions with the naive interpolation, showing the existence of barriers in the predictive likelihood for Gaussian VI solutions. However, with the alignment proposed in § 4 and Algorithm 1, we recover zero barrier solutions for MLPs on both MNIST and CIFAR10, and nearly-zero barrier for ResNet20 on CIFAR10. This holds both for the train and test splits, with quantifiably smaller barriers in the test set.

In Fig. 5 we study the effect of the width of a neural network in relation to the loss barrier by taking an MLP and a ResNet20 with an increasing number of hidden features. We see that wider models generally provide lower barriers: for MLPs this holds true with and without alignment, while for the ResNet20 this is happening only after alignment. This extends some previous analysis done on loss-optimized networks. Specifically, Entezari et al. [30] show that barriers seem to have a double descent trend, while Ainsworth et al. [2] discuss that low barrier solutions after accounting for symmetries are easier to find in wider networks.

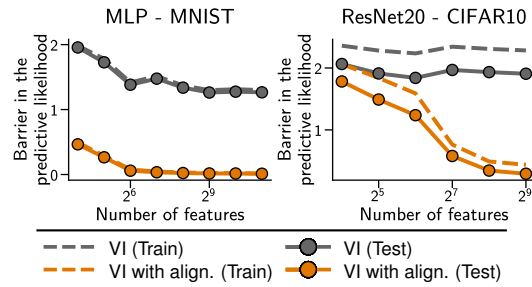


Figure 5: **Effect of width.** After distribution alignment, wider models exhibit lower likelihood barrier.

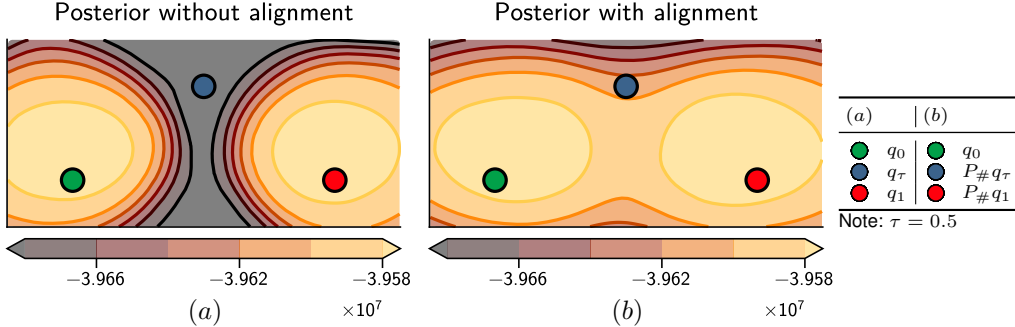


Figure 6: **Posterior density visualization.** Analysis of the log-posterior computed for ResNet20. Samples from  $q_0$  and  $q_1$  are connected by lower density regions, while  $q_0$  and  $P_{\#}q_1$  are not.

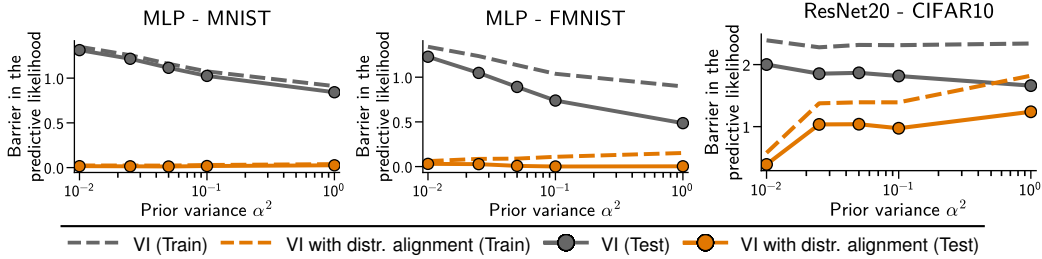


Figure 7: **Effect of prior variance.** After distribution alignment, prior variance has low effect in finding zero-loss barriers, while with naive interpolation we see a decreasing trend the higher the variance is.

We speculate that this might be due to the limiting behavior of Bayesian neural networks, which makes the posterior landscape Gaussian-like [48, 40]. While this does not fully explain the phenomenon observed, the existing connections between BNNs and non-parametric models, like Gaussian Processes (GPs) [83] and deep Gaussian processes (DGPs) [22, 20, 28, 93], can provide additional insights on the role of symmetries in weight space [80].

Finally, as an additional check, we analyze the log-posterior with and without alignment by projecting the density into two dimensional slices, following the setup in [47, 36]. We study the two dimensional subspace of the parameter space supported by the hyperplane  $H = \{\theta \in \mathbb{R}^d \mid \theta = a\theta_a + b\theta_b + (1 - a - b)\theta_c\}$ , where  $a, b \in \mathbb{R}$  and  $\theta_a, \theta_b$  and  $\theta_c$  are the samples either from  $q_0, q_1$  and  $q_\tau$  without alignment or from  $q_0, P_{\#}q_1$  and  $P_{\#}q_\tau$  with alignment. With this configuration, all three samples always lie on this hyper-plane. In Fig. 6, we present the visualization of ResNet20 trained on CIFAR10. We see that the samples from  $q_0$  and  $P_{\#}q_1$  are connected by higher density regions than the ones between  $q_0$  and  $q_1$ . This is in line with the results in Fig. 4, where we see that the loss barrier is lower after alignment.

## 5.2 Analyzing the effect of the prior and testing the cold posterior effect

In all previous experiments we used a Gaussian prior  $\mathcal{N}(\mathbf{0}, \alpha^2 \mathbf{I})$  with fixed  $\alpha^2$ ; now we study the effect of a varying prior variance  $\alpha^2$  on the behavior of the loss barriers. We experiment this on a MLP trained on MNIST and Fashion-MNIST and on a ResNet20 (width x8) on CIFAR10. We report the results in Fig. 7. We can appreciate two behaviors: with alignment, there is no measurable effect of using different variances in finding zero-barrier solutions; on the contrary, without alignment we see that naive VI solutions are easier to interpolate with lower barrier when the prior is more diffused. At the same time, we see that higher variances produce bigger gaps between train barriers and test barriers. We speculate that this is due to overfitting happening with more relaxed priors, which makes low-barrier (but low-likelihood) solutions easier to find.

Additionally, several previous works have analyzed the effect of tempering the posterior in BNNs [106, 112, 110, 47]. Specifically, we are interested in the distribution  $p_T(\theta \mid \mathbf{Y}) \propto (p(\mathbf{Y} \mid \theta, \mathbf{X})p(\theta))^{1/T}$ , where  $T$  is known as the temperature. Note that starting from the above definition, we can write an equivalent ELBO for VI which takes into account  $T$ . For  $T < 1$ , we have cold posteriors, which are



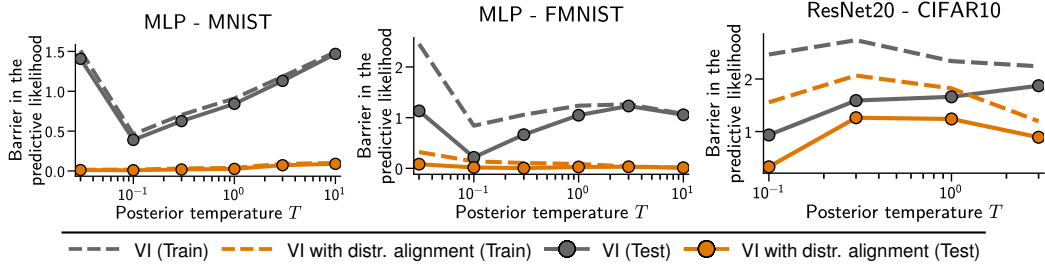


Figure 8: **Effect of temperature.** After alignment, cold posteriors makes barriers marginally closer to zero

sharper than the true Bayesian posteriors, while for  $T > 1$  we have warm posterior, which are more diffused. In Fig. 8 we see that barriers for cold posteriors with alignment are marginally closer to zero than for warm posteriors. Note that cold temperatures concentrate the distribution around the MAP, which motivates a further comparison with a non-Bayesian approach.

### 5.3 Comparison with SGD solutions

Motivated by the results with cold posteriors, we also compare the behavior of barriers for VI versus MAP solutions obtained via SGD. Note that using the MAP solution we end up with the same setup of weight matching than in [2]. In Fig. 9 we report this comparison, which is carried out on ResNet20 for various width multipliers. For narrow models, we see that VI with alignment and MAP with weight matching both behave in a similar manner. Interestingly, as the model becomes wider, MAP solutions achieve marginally lower barriers than VI. We speculate that this might be due to the simple Gaussian parameterization for the approximate posterior, which doesn't completely capture the local geometry of the true posterior.

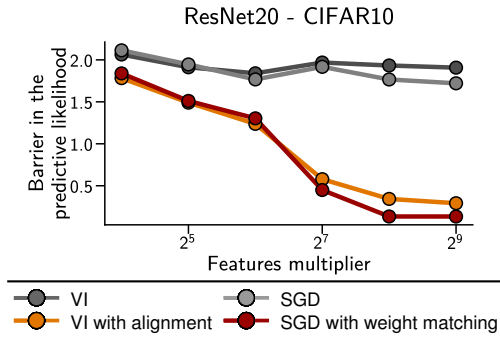


Figure 9: **VI versus SGD.** VI and SGD behave equivalently after permutation alignment in narrow models. For wider models, SGD solutions reach lower barriers than VI.

### 5.4 Effect of normalization layers and data augmentation

We conclude this section with a discussion on the effects of normalization layers and data augmentation on the loss barriers for Bayesian neural networks.

Different normalization strategies can affect the overall geometry of the problem [30, 2]. As discussed in [49], interpolating with BatchNorm layers [45] is pathological due to *variance collapse* of the feature representation in hidden layers. Additionally, batch-dependent normalization layers don't have a clear Bayesian interpretation, since the likelihood cannot factorize. For our experiments we choose to use the FRN layer, as done in previous works [e.g., 47]. Note that FRN layers are invariant to permutation units and therefore can be aligned without problems. To test the *variance collapse* behavior, we analyze the variance of activations following the instructions in [49, §3.1], with the sole difference that the activations are marginalized w.r.t. samples from the posterior. In Fig. 10 we can see that there isn't a pathological variance collapse after alignment. Finally, in Fig. 11 we compare another normalization layer, the LayerNorm (LN) [6]. Note that LN is also batch-independent, it has a clear Bayesian interpretation and it is invariant to permutation units. Indeed, we see that both normalization layers can be aligned, and LN exhibits lower barriers than FRN.

Finally, in all previous experiments we skipped data augmentation, because the random augmentations introduce stochasticity which lacks a proper Bayesian interpretation in the formulation of the likelihood function (e.g. re-weighting of the likelihood due to the increase of the effective sample size [77]). Additionally, data augmentation can contribute to spurious effects difficult to disentangle (e.g., cold posterior effect [106]). Nonetheless, during the development of the method we didn't make an

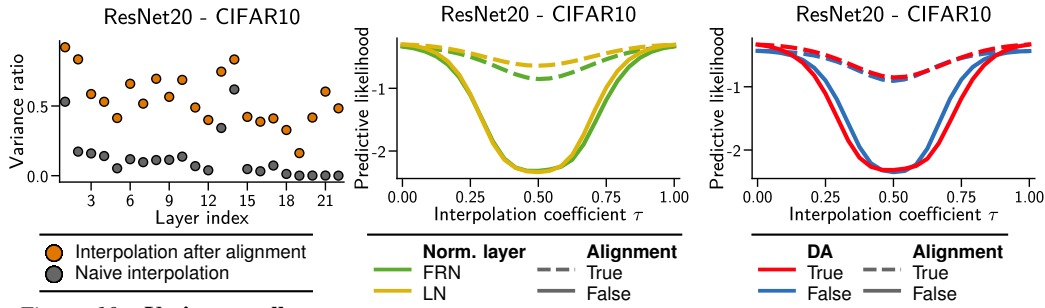


Figure 10: **Variance collapse.** Variance collapse is present for the naive interpolation, but it is reduced after alignment.

Figure 11: **Effect of normalization layers.** Both FRN and LN can be aligned, and LN exhibits lower barriers.

Figure 12: **Effect of data augmentation.** With and without DA, we are able to recover similar low barrier solutions.

assumption on data augmentation and in Fig. 12 we experiment both with and without augmentation, showing that we are still able to recover similar low barrier solutions in both cases. Having said that, we advocate caution when using data augmentation in Bayesian neural networks, as it changes the shape of the posterior.

## 6 Related work

In earlier sections of the paper, we already briefly discussed and reviewed relevant works on mode connectivity, symmetries in the loss landscape and connection with gradient-based optimization methods. Here we discuss some relevant works on the connection to Bayesian deep learning. The work of Garipov et al. [36] sparked several contributions on exploiting mode connectivity for ensembling models, which is akin to Bayesian model averaging. For example, in [46] the authors propose to ensemble models using curve subspaces to construct low-dimensional subspaces of parameter space. These curve subspaces are, among others, the non-linear paths connecting low-loss modes (and consequently high-posterior density) in weight space. In [31], the authors attempt to explain the effectiveness of deep ensembles [59], concluding that it is partially due to the diversity of the SGD solutions in parameter space induced by random initialization. More recently, in [9] the authors reason about mode connecting volumes, which are multi-dimensional manifolds of low loss that connect many independently trained models. These mode connecting volumes form the basis for an efficient method for building simplicial complexes for ensembling. Here, we want to highlight that these works have not taken into account the permutation symmetries. Finally, in a concurrent submission, [107] proposes an algorithm to remove symmetries in MCMC chains for tanh networks.

## 7 Conclusions

By studying the effect of permutation symmetries, which are ubiquitous in neural networks, it is possible to analyze the fundamental geometric properties of loss landscapes like (linear) mode connectivity and loss barriers. While previously this was done on loss-optimized networks [2, 30, 32], in this work we have extended the analysis to Bayesian neural networks. We have studied the linear connectivity properties of approximate Bayesian solutions and we have proposed a matching algorithm (Algorithm 1) to search for linearly connected solutions, by aligning the distributions of two independent VI solutions with respect to permutation matrices. We have empirically validated our framework on a variety of experiments, showing that we can find zero barrier linearly-connected solutions for BNNs trained with VI, on shallow models as well as on deep convolutional networks. This brings evidence for Conjecture 1 regarding the linear connectivity of variational inference approximations for BNNs. Furthermore, we have studied the effect of various design hyper-parameters, like width, prior and temperature, and observed complex patterns of behavior, which would require additional research. In particular, the experiments raise questions regarding the relation between linear mode connectivity and the generalization of BNNs, as well as the role of width with respect to limiting behaviors of non-parametric models like GPs and DGPs.

## Acknowledgments and Disclosure of Funding

The authors want to thank David Bertrand and the AIAO team at Stellantis for their work on the software and hardware infrastructure used for this work.

## References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015. Software available from tensorflow.org.
- [2] S. Ainsworth, J. Hayase, and S. Srinivasa. Git Re-Basin: Merging Models modulo Permutation Symmetries. In *The Eleventh International Conference on Learning Representations*, 2023.
- [3] Z. Allen-Zhu and Y. Li. What Can ResNet Learn Efficiently, Going Beyond Kernels? In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [4] L. Ambrosio, N. Gigli, and G. Savare. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics. ETH Zürich. Birkhäuser Basel, 2008. ISBN 9783764387228.
- [5] S. Arora, S. S. Du, W. Hu, Z. Li, R. R. Salakhutdinov, and R. Wang. On Exact Computation with an Infinitely Wide Neural Net. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [6] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer Normalization, 2016.
- [7] Y. Bai and J. D. Lee. Beyond Linearization: On Quadratic and Higher-Order Approximation of Wide Neural Networks. In *International Conference on Learning Representations*, 2020.
- [8] Y. Bengio, O. Delalleau, and N. Le Roux. The Curse of Dimensionality for Local Kernel Machines. Technical Report TR-1258, March 2005.
- [9] G. Benton, W. Maddox, S. Lotfi, and A. G. G. Wilson. Loss Surface Simplexes for Mode Connecting Volumes and Fast Ensembling. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 769–779. PMLR, 18–24 Jul 2021.
- [10] D. Bertsekas. *Network Optimization: Continuous and Discrete Methods*. Athena scientific optimization and computation series. Athena Scientific, 1998. ISBN 9788865290279.
- [11] C. M. Bishop. *Pattern recognition and machine learning*. Springer, 1st ed. 2006. Corr. 2nd printing 2011 edition, 2006. ISBN 0387310738.
- [12] D. Blei, S. Mohamed, and R. Ranganath. Variational inference: Foundations and modern methods. 2016.
- [13] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational Inference: A Review for Statisticians. 112(518):859–877, 2017.
- [14] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight Uncertainty in Neural Network. In F. Bach and D. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1613–1622. PMLR, 2015.
- [15] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. JAX: Composable Transformations of Python+NumPy Programs, 2018.
- [16] Y. Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on Pure and Applied Mathematics*, 44(4):375–417, 1991.

- [17] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language Models are Few-Shot Learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.
- [18] M. Chen, Y. Bai, J. D. Lee, T. Zhao, H. Wang, C. Xiong, and R. Socher. Towards Understanding Hierarchical Learning: Benefits of Neural Representations. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 22134–22145. Curran Associates, Inc., 2020.
- [19] T. Chen, E. Fox, and C. Guestrin. Stochastic Gradient Hamiltonian Monte Carlo. In E. P. Xing and T. Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, Proceedings of Machine Learning Research, pages 1683–1691. PMLR, 2014.
- [20] K. Cutajar, E. V. Bonilla, P. Michiardi, and M. Filippone. Random feature expansions for deep Gaussian processes. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 884–893. PMLR, 2017-08.
- [21] M. Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- [22] A. C. Damianou and N. D. Lawrence. Deep Gaussian Processes. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2013*, volume 31 of *Proceedings of Machine Learning Research*, pages 207–215. JMLR.org, 2013.
- [23] A. G. de G. Matthews, J. Hron, M. Rowland, R. E. Turner, and Z. Ghahramani. Gaussian Process Behaviour in Wide Deep Neural Networks. In *Proceedings of the 6th International Conference on Learning Representations*, 2018.
- [24] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- [25] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*, 2021.
- [26] F. Draxler, K. Veschgini, M. Salmhofer, and F. Hamprecht. Essentially No Barriers in Neural Network Energy Landscape. In J. Dy and A. Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1309–1318. PMLR, 10–15 Jul 2018.
- [27] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. 195(2):216 – 222, 1987.
- [28] M. M. Dunlop, M. A. Girolami, A. M. Stuart, and A. L. Teckentrup. How Deep Are Deep Gaussian Processes? 19(1):2100–2145, 2018-01.
- [29] V. Dutoit, J. Hensman, M. van der Wilk, C. H. Ek, Z. Ghahramani, and N. Durrande. Deep Neural Networks as Point Estimates for Deep Gaussian Processes. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 9443–9455. Curran Associates, Inc., 2021.
- [30] R. Entezari, H. Sedghi, O. Saukh, and B. Neyshabur. The Role of Permutation Invariance in Linear Mode Connectivity of Neural Networks. In *International Conference on Learning Representations*, 2022.
- [31] S. Fort, C. H. Hu, and B. Lakshminarayanan. Deep Ensembles: A Loss Landscape Perspective, 2020.

- [32] J. Frankle, G. K. Dziugaite, D. Roy, and M. Carbin. Linear Mode Connectivity and the Lottery Ticket Hypothesis. In H. D. III and A. Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3259–3269. PMLR, 13–18 Jul 2020.
- [33] G. Franzese, D. Milios, M. Filippone, and P. Michiardi. Revisiting the Effects of Stochasticity for Hamiltonian Samplers. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 6744–6778. PMLR, 17–23 Jul 2022.
- [34] G. Franzese, S. Rossi, L. Yang, A. Finamore, D. Rossi, M. Filippone, and P. Michiardi. How Much Is Enough? A Study on Diffusion Times in Score-Based Generative Models. *Entropy*, 25(4), 2023.
- [35] Y. Gal and Z. Ghahramani. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In M. F. Balcan and K. Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1050–1059. PMLR, 2016.
- [36] T. Garipov, P. Izmailov, D. Podoprikin, D. P. Vetrov, and A. G. Wilson. Loss Surfaces, Mode Connectivity, and Fast Ensembling of DNNs. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 8789–8798. Curran Associates, Inc., 2018.
- [37] A. Garriga-Alonso, C. E. Rasmussen, and L. Aitchison. Deep Convolutional Networks as shallow Gaussian Processes. In *International Conference on Learning Representations*, 2019.
- [38] C. Godfrey, D. Brown, T. Emerson, and H. Kvinge. On the Symmetries of Deep Learning Models and their Internal Representations. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [39] A. Graves. Practical Variational Inference for Neural Networks. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 2348–2356. Curran Associates, Inc., 2011.
- [40] B. He, B. Lakshminarayanan, and Y. W. Teh. Bayesian Deep Ensembles via the Neural Tangent Kernel. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33*, 2020.
- [41] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778, 2016.
- [42] G. E. Hinton and D. van Camp. Keeping the Neural Networks Simple by Minimizing the Description Length of the Weights. In *Proceedings of the Sixth Annual Conference on Computational Learning Theory*, 1993. ISBN 0897916115.
- [43] M. D. Hoffman, D. M. Blei, C. Wang, and J. W. Paisley. Stochastic Variational Inference. 14 (1):1303–1347, 2013.
- [44] T. Huix, S. Majewski, A. Durmus, E. Moulines, and A. Korba. Variational Inference of overparameterized Bayesian Neural Networks: a theoretical and empirical study, 2022.
- [45] S. Ioffe and C. Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In F. Bach and D. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 448–456, Lille, France, 07–09 Jul 2015. PMLR.
- [46] P. Izmailov, W. Maddox, P. Kirichenko, T. Garipov, D. P. Vetrov, and A. G. Wilson. Subspace Inference for Bayesian Deep Learning. In *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI 2019, Tel Aviv, Israel, July 22-25, 2019*, volume 115 of *Proceedings of Machine Learning Research*, pages 1169–1179. AUAI Press, 2019.
- [47] P. Izmailov, S. Vikram, M. D. Hoffman, and A. G. Wilson. What are bayesian neural network posteriors really like? In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 4629–4640. PMLR, 2021.

- [48] A. Jacot, F. Gabriel, and C. Hongler. Neural Tangent Kernel: Convergence and Generalization in Neural Networks. In *Advances in Neural Information Processing Systems*, volume 31, pages 8571–8580. Curran Associates, Inc., 2018.
- [49] K. Jordan, H. Sedghi, O. Saukh, R. Entezari, and B. Neyshabur. REPAIR: REnormalizing Permuted Activations for Interpolation Repair. In *The Eleventh International Conference on Learning Representations*, 2023.
- [50] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An Introduction to Variational Methods for Graphical Models. 37(2):183–233, 1999-11-01.
- [51] L. V. Kantorovich. On the transfer of masses. 37:227–229, 1942.
- [52] L. V. Kantorovich. On a problem of Monge. 3:225–226, 1948.
- [53] M. E. E. Khan, A. Immer, E. Abedi, and M. Korzepa. Approximate Inference Turns Deep Networks into Gaussian Processes. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [54] D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*, 2014.
- [55] D. P. Kingma, T. Salimans, and M. Welling. Variational Dropout and the Local Reparameterization Trick. In *Advances in Neural Information Processing Systems 28*, pages 2575–2583. Curran Associates, Inc., 2015.
- [56] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- [57] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012.
- [58] R. Kurle, R. Herbrich, T. Januschowski, B. Wang, and J. Gasthaus. On the detrimental effect of invariances in the likelihood for variational inference. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [59] B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6402–6413. Curran Associates, Inc., 2017.
- [60] Y. LeCun, C. Cortes, and C. Burges. MNIST handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- [61] J. Lee, J. Sohl-dickstein, J. Pennington, R. Novak, S. Schoenholz, and Y. Bahri. Deep Neural Networks as Gaussian Processes. In *International Conference on Learning Representations*, 2018.
- [62] Y. Li and Y. Gal. Dropout Inference in Bayesian Neural Networks with Alpha-divergences. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2052–2061. PMLR, 2017.
- [63] Q. Liu and D. Wang. Stein Variational Gradient Descent: A General Purpose Bayesian Inference Algorithm. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 2378–2386. Curran Associates, Inc., 2016.
- [64] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie. A ConvNet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [65] D. J. C. MacKay. Bayesian Model Comparison and Backprop Nets. In J. E. Moody, S. J. Hanson, and R. Lippmann, editors, *Advances in Neural Information Processing Systems 4*, pages 839–846. Morgan Kaufmann, 1991.
- [66] D. J. C. MacKay. A Practical Bayesian Framework for Backpropagation Networks. 4(3): 448–472, 1992-05.

- [67] D. J. C. Mackay. Bayesian methods for backpropagation networks. In E. Domany, J. L. van Hemmen, and K. Schulten, editors, *Models of Neural Networks III*, chapter 6, pages 211–254. Springer, 1994.
- [68] D. J. C. MacKay. Choice of Basis for Laplace Approximation. 33(1):77–86, 1998.
- [69] C. J. Maddison, A. Mnih, and Y. W. Teh. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. In *Proceedings of the 5th International Conference on Learning Representations*. OpenReview.net, 2017.
- [70] S. Mandt, M. D. Hoffman, and D. M. Blei. Stochastic Gradient Descent as Approximate Bayesian Inference. *Journal of Machine Learning Research*, 18(134):1–35, 2017.
- [71] A. G. d. G. Matthews, J. Hensman, R. Turner, and Z. Ghahramani. On Sparse Variational Methods and the Kullback-Leibler Divergence between Stochastic Processes. In A. Gretton and C. C. Robert, editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 231–239, Cadiz, Spain, 09–11 May 2016. PMLR.
- [72] N. Metropolis and S. Ulam. The monte carlo method. 44(247):335–341, 1949. PMID: 18139350.
- [73] R. M. Neal. *Bayesian Learning for Neural Networks*. PhD thesis, 1994.
- [74] R. M. Neal. Priors for infinite networks. Technical report crg-tr-94-1, University of Toronto, 1994.
- [75] R. M. Neal. *MCMC Using Hamiltonian Dynamics*, chapter 5. CRC Press, 2011.
- [76] B. Neyshabur, Z. Li, S. Bhojanapalli, Y. LeCun, and N. Srebro. The role of over-parametrization in generalization of neural networks. In *International Conference on Learning Representations*, 2019.
- [77] K. Osawa, S. Swaroop, M. E. E. Khan, A. Jain, R. Eschenhagen, R. E. Turner, and R. Yokota. Practical Deep Learning with Bayesian Principles. In *Advances in Neural Information Processing Systems*, volume 32, pages 4287–4299. Curran Associates, Inc., 2019.
- [78] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [79] F. Pittorino, A. Ferraro, G. Perugini, C. Feinauer, C. Baldassi, and R. Zecchina. Deep Networks on Toroids: Removing Symmetries Reveals the Structure of Flat Regions in the Landscape Geometry. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 17759–17781. PMLR, 17–23 Jul 2022.
- [80] G. Pleiss and J. P. Cunningham. The Limitations of Large Width in Neural Networks: A Deep Gaussian Process Perspective. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.
- [81] Y. Polyanskiy and Y. Wu. Strong Data-Processing Inequalities for Channels and Bayesian Networks. In E. Xarlen, M. Madiman, and E. M. Werner, editors, *Convexity and Concentration*, pages 211–249, New York, NY, 2017. Springer New York.
- [82] A. Radford and K. Narasimhan. Improving Language Understanding by Generative Pre-Training. 2018.
- [83] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005. ISBN 026218253X.
- [84] H. Ritter, A. Botev, and D. Barber. A Scalable Laplace Approximation for Neural Networks. In *Proceedings of the 6th International Conference on Learning Representations*. OpenReview.net, 2018.
- [85] H. Robbins and S. Monro. A Stochastic Approximation Method. 22(3):400–407, 1951.

- [86] G. Roeder, Y. Wu, and D. Duvenaud. Sticking the Landing: Simple, Lower-Variance Gradient Estimators for Variational Inference. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6925–6934, 2017.
- [87] S. Rossi, P. Michiardi, and M. Filippone. Good Initializations of Variational Bayes for Deep Models. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5487–5497. PMLR, 2019.
- [88] S. Rossi, S. Marmin, and M. Filippone. Walsh-Hadamard Variational Inference for Bayesian Deep Learning. In *Advances in Neural Information Processing Systems 33*, 2020.
- [89] T. G. J. Rudner, Z. Chen, Y. W. Teh, and Y. Gal. Tractable Function-Space Variational Inference in Bayesian Neural Networks. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [90] F. R. Ruiz, M. Titsias, and D. Blei. The Generalized Reparameterization Gradient. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [91] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning Representations by Back-propagating Errors. 323(6088):533–536, 1986-10.
- [92] T. Salimans and D. A. Knowles. Fixed-Form Variational Posterior Approximation through Stochastic Linear Regression. 8(4):837–882, 2013.
- [93] H. Salimbeni and M. Deisenroth. Doubly Stochastic Variational Inference for Deep Gaussian Processes. In *Advances in Neural Information Processing Systems 30*, pages 4588–4599. Curran Associates, Inc., 2017.
- [94] B. Simsek, F. Ged, A. Jacot, F. Spadaro, C. Hongler, W. Gerstner, and J. Brea. Geometry of the Loss Landscape in Overparameterized Neural Networks: Symmetries and Invariances. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 9722–9732. PMLR, 18–24 Jul 2021.
- [95] E. Snelson and Z. Ghahramani. Sparse Gaussian Processes using Pseudo-Inputs. In *Advances in Neural Information Processing Systems 18*, pages 1257–1264. MIT Press, 2006.
- [96] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [97] J. Song, C. Meng, and S. Ermon. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*, 2021.
- [98] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*, 2021.
- [99] S. Sun, G. Zhang, J. Shi, and R. B. Grosse. Functional Variational Bayesian Neural Networks. In *Proceedings of the 7th International Conference on Learning Representations*. OpenReview.net, 2019.
- [100] A. Takatsu. Wasserstein geometry of Gaussian measures. *Osaka Journal of Mathematics*, 48(4):1005 – 1026, 2011.
- [101] Tishby, Levin, and Solla. Consistent inference of probabilities in layered networks: predictions and generalizations. In *International 1989 Joint Conference on Neural Networks*, pages 403–409 vol.2, 1989.
- [102] B.-H. Tran, S. Rossi, D. Milios, P. Michiardi, E. V. Bonilla, and M. Filippone. Model Selection for Bayesian Autoencoders. In *Advances in Neural Information Processing Systems*, volume 34, pages 19730–19742. Curran Associates, Inc., 2021.
- [103] B.-H. Tran, S. Rossi, D. Milios, and M. Filippone. All You Need is a Good Functional Prior for Bayesian Deep Learning. In *The Journal of Machine Learning Research*, 2022.
- [104] C. Villani. *Optimal Transport: Old and New*, volume 338. Springer Science & Business Media, 2008.



- [105] M. Welling and Y. W. Teh. Bayesian Learning via Stochastic Gradient Langevin Dynamics. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pages 681–688, Madison, WI, USA, 2011. Omnipress. ISBN 9781450306195.
- [106] F. Wenzel, K. Roth, B. S. Veeling, J. Świątkowski, L. Tran, S. Mandt, J. Snoek, T. Salimans, R. Jenatton, and S. Nowozin. How Good is the Bayes Posterior in Deep Neural Networks Really? In *Proceeding of the 37th International Conference on Machine Learning*, 2020.
- [107] J. G. Wiese, L. Wimmer, T. Papamarkou, B. Bischl, S. Günnemann, and D. Rügamer. Towards Efficient MCMC Sampling in Bayesian Neural Networks by Exploiting Symmetry, 2023.
- [108] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. *CoRR*, abs/1708.07747, 2017.
- [109] G. Yehudai and O. Shamir. On the Power and Limitations of Random Features for Understanding Neural Networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [110] C. Zeno, I. Golan, A. Pakman, and D. Soudry. Why Cold Posteriors? On the Suboptimal Generalization of Optimal Bayes Estimates. In *Third Symposium on Advances in Approximate Bayesian Inference*, 2021.
- [111] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.
- [112] R. Zhang, C. Li, J. Zhang, C. Chen, and A. G. Wilson. Cyclical Stochastic Gradient MCMC for Bayesian Deep Learning. In *Proceedings of the 8th International Conference on Learning Representations*. OpenReview.net, 2020.

## A Future work and open questions

Besides the specific open questions discussed above, we anticipate some possible future work. The framework we proposed to analyze the loss barriers in BNNs is general and can be applied to approximations other than VI. Future work could undertake to extend this analysis to the Laplace approximation [68]. However, this raises methodological challenges since Algorithm 1 is not directly applicable, due to the use of a dense covariance matrix (the inverse of the Hessian). Additionally, it would be worth extending these results to sample-based inference, like stochastic gradient Hamiltonian Monte Carlo (SGHMC) [19], or particle-based inference, like Stein variational inference [63]. This would call for a careful analysis, starting from the solution of Eq. (6), which would require approximations [21].

## B Additional details on the alignment method

In the main paper, to align the distributions with respect to permutation matrices we argue to use the Wasserstein distance rather than the KL divergence. Indeed, by considering the KL divergence  $\text{KL}[P_{\#}q_1 \parallel q_0]$  between Gaussians we have

$$\text{KL}[P_{\#}q_1 \parallel q_0] = \log \det \text{diag}(s_0) - \log \det \text{diag}(Ps_1) + \text{Tr}(\text{diag}(Ps_1s_0^{-1})) + \quad (20)$$

$$(m_0 - Pm_1)^\top \text{diag}(s_0^{-1})(m_0 - Pm_1) \quad (21)$$

It’s easy to verify that the first three terms do not depend on  $P$ , leading to just a distance between means and disregarding any covariance information. In the figure below, we visualize the difference between doing LAP with the KL cost and LAP with the Wasserstein cost.

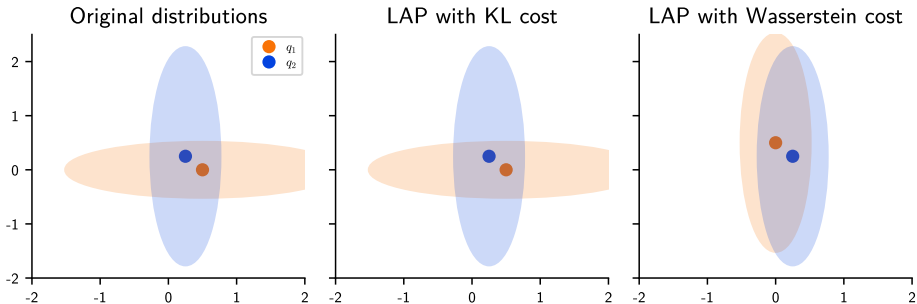


Figure A13: **Alignment using different objectives.** Given two distributions symmetrical w.r.t. the  $y = x$  plane, using the KL cost LAP results in the identity permutation (which fails to recover the symmetry), while the Wasserstein cost better aligns the two distributions

## C Experimental setup

If not otherwise stated, we start by training two replicas of BNNs with variational inference and we compute the marginalized barrier as  $\mathcal{B}(q_0, q_1) = \max_{\tau} \mathcal{L}(q_{\tau}) - ((1 - \tau)\mathcal{L}(q_0) + \tau\mathcal{L}(q_1))$  where  $\mathcal{L}(\cdot)$  is the predictive likelihood and  $\tau \in [0, 1]$ , from which we take 25 evenly distributed points. In particular, we seek to understand what happens to the VI solutions with and without alignment applied to one of the two distribution ( $q_1$  in our case). We experiment with MLPs with three layers and ResNet20 [41] with various widths on MNIST [60], Fashion-MNIST [108] and CIFAR10 [56]. All models are trained without data augmentation, because the random augmentations introduce stochasticity which lacks a proper Bayesian interpretation in the formulation of the likelihood function [106]. Finally, we set the prior to be Gaussian  $\mathcal{N}(0, \alpha^2 I)$ , with the flexibility of choosing the variance. All VI models are trained using the classic SGD optimizer with momentum [85, 91] using the reparameterization trick [54] with one sample during training and 128 samples during testing. We use the categorical distribution and the Gaussian distribution as classification and regression likelihood, respectively. Tables 2 and 3 show details on the MLPs and convolutional neural networks (CNNs) base architectures used in our experimental campaign, while Table 1 reports the hyperparameters used in the experiments. Note that differently from Entezari et al. [30] and Ainsworth

Table 1: Hyperparameters used for the experiments

<b>Dataset</b>	CIFAR10	MNIST
<b>Model</b>	ResNet20	MLP
Data Aug.	False	False
Batch size	500	500
Temperature	1.0	1.0
Test samples	128	128
Train samples	1	1
VI std. init	0.01	0.01
Base features	16	512
Prior var	0.01	0.0025
Learning rate	0.000001	0.000001
Train epochs	1000	1000

et al. [2], we don't use data augmentation. A possible protocol for handling data augmentation in BNNs is presented by Osawa et al. [77] and involves carefully tuning the likelihood temperature to correctly counting the number of data points.

Table 2: MLP

<b>Layer</b>	<b>Dimensions</b>
Linear-ReLU	$512 \times D_{in}$
Linear-ReLU	$512 \times 512$
Linear-ReLU	$512 \times 512$
Linear-Softmax	$D_{out} \times 512$

Table 3: ResNet20

<b>Layer</b>	<b>Dimensions</b>
Conv2D	$16 \times 3 \times 3 \times D_{in}$
Residual Block	$\begin{matrix} 3 \times 3, 16 \\ 3 \times 3, 16 \end{matrix} \times 3$
Residual Block	$\begin{matrix} 3 \times 3, 32 \\ 3 \times 3, 32 \end{matrix} \times 3$
Residual Block	$\begin{matrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{matrix} \times 3$
AvgPool	$8 \times 8$
Linear-Softmax	$D_{out} \times 64$

### C.1 Computing platform

The experiments have been performed using JAX [15] and run on two AWS p4d.24xlarge instances with 8 NVIDIA A100 GPUs. Experiments were conducted using in the eu-west-1 region, which has a carbon efficiency of 0.62 kgCO<sub>2</sub>eq/kWh. A cumulative of 6500 hours of computation was performed on GPUs and it includes interactive sessions as well as small experiments with very low GPU usage, providing a pessimistic estimation of the true utilization. Total emissions are estimated to be 1007.5 kgCO<sub>2</sub>eq of which 100 percents were directly offset by AWS.

## D Additional results

We present timings obtained by profiling the time needed to solve the SOLAP with the Wasserstein cost, as well as the time for the deterministic case [2]. In Fig. A14 we show the results for MLP and ResNet20 architectures, varying the model width. It is evident that, in the majority of cases, the algorithm completes within a minute. Moreover, as anticipated, in case of VI solving our distribution alignment problem for wide neural networks is more computationally demanding compared to merely matching weights from SGD solutions.

Finally, we also test our setup on the CIFAR100 dataset [56]. Surprisingly, we were not able to replicate the same level of performance as in the other cases. In Fig. A15, we see that, despite converging well, we fall short to find zero-barrier solutions. Similarly to the comments of Ainsworth et al. [2], we also stress that the failure to align distributions does not rule out the existence of a proper permutation map that the algorithm couldn't find. Nonetheless, this raises a number of questions: the Bayesian posterior is the product of two ingredients, the prior and the likelihood, conditioned to observing a dataset.

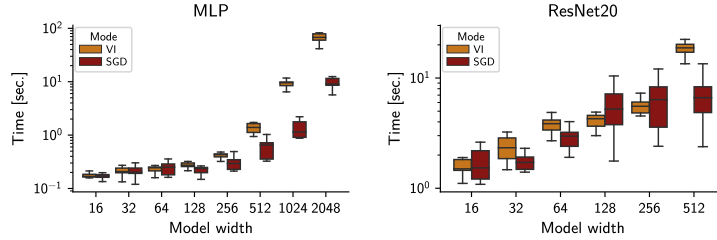


Figure A14: **Timings.** Profile of the algorithms to align the VI solutions and to match weights from SGD solutions.

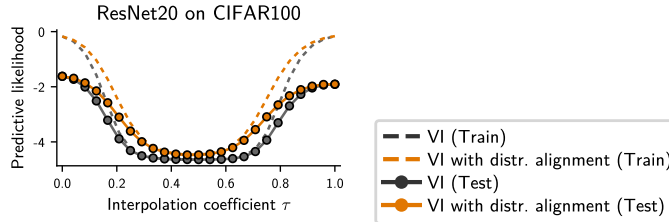


Figure A15: **Alignment failure.** The method proposed fails to recover zero-barrier solutions for CIFAR100.

Finally, as an additional check, we analyze the log-posterior by projecting the density into two dimensional slices, following the setup in [47, 36]. We study the two dimensional subspace of the parameter space supported by the hyperplane  $H$  of the form

$$H = \{\theta \in \mathbb{R}^d \mid \theta = a\theta_a + b\theta_b + (1 - a - b)\theta_c\},$$

where  $a, b \in \mathbb{R}$  and  $\theta_a, \theta_b$  and  $\theta_c$  are the means of  $q_0, q_1$  and  $P_{\#}q_1$ . With this configuration, all three solutions lie on this hyper-plane. In Fig. 6, we present the visualization of ResNet20 trained on CIFAR10. We see that the distributions  $q_0$  and  $P_{\#}q_1$  are connected by higher density regions than the ones between  $q_0$  and  $q_1$ . Also, as expected the symmetries arise from the form of the likelihood, and the prior has a comparable strength with respect to the three posteriors. Later, we also study in more details the effect of the prior’s variance in finding low-barrier solutions.

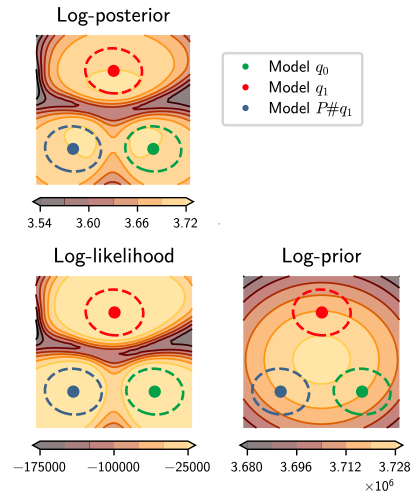


Figure A16: **Posterior density visualization.** All three solutions are local approximation of posterior, but  $q_0$  and  $P_{\#}q_1$  are connected by lower density regions.

## E A primer on variational inference for Bayesian neural networks

VI is a classic tool to tackle intractable Bayesian inference [50, 13]. VI casts the inference problem into an optimization-based procedure to compute a tractable approximation of the true posterior. Assume a generic parametric model  $f$  parameterized by some unknown parameters  $\theta$  (i.e.  $f(\cdot, \theta)$ ) and a collection of data  $\mathbf{y} \in \mathbb{R}^N$  corresponding to some input points  $\mathbf{X} = \{\mathbf{x}_i \mid \mathbf{x}_i \in \mathbb{R}^{D_{in}}\}_{i=1, \dots, N}$ . In our setting, we have a probabilistic model  $p(\mathbf{y} \mid f(\mathbf{X}; \theta))$  with parameters  $\theta$ , a prior distributions on them  $p(\theta)$  and a set of observations  $\{\mathbf{X}, \mathbf{y}\}$ . In a nutshell, the general recipe of VI consists of (i) introducing a set  $\mathcal{Q}$  of distributions; (ii) defining a tractable objective that “measure” the distance between any arbitrary distribution  $q(\theta) \in \mathcal{Q}$  and the true posterior  $p(\theta \mid \mathbf{y})$ ; and finally (iii) providing a programmatic way to find the distribution  $\tilde{q}(\theta)$  that minimizes such distance. In practice,  $q(\theta)$  has some free parameters  $\nu$  (also known as *variational parameters*), which are optimized such that

the approximating distribution  $q(\boldsymbol{\theta}; \boldsymbol{\nu})$  is as closer as possible to the true posterior  $p(\boldsymbol{\theta} | \mathbf{y})$ . We can derive the variational objective starting from the definition of the KL,

$$\begin{aligned} \text{KL} [q(\boldsymbol{\theta}; \boldsymbol{\nu}) \parallel p(\boldsymbol{\theta} | \mathbf{y})] &= \mathbb{E}_{q(\boldsymbol{\theta}; \boldsymbol{\nu})} [\log q(\boldsymbol{\theta}; \boldsymbol{\nu}) - \log p(\boldsymbol{\theta} | \mathbf{y})] = \\ &= \mathbb{E}_{q(\boldsymbol{\theta}; \boldsymbol{\nu})} [\log q(\boldsymbol{\theta}; \boldsymbol{\nu}) - \log p(\mathbf{y} | \boldsymbol{\theta}) - \log p(\boldsymbol{\theta})] + \log p(\mathbf{y}) \end{aligned} \quad (22)$$

Rearranging we have that

$$\log p(\mathbf{y}) - \text{KL} [q(\boldsymbol{\theta}; \boldsymbol{\nu}) \parallel p(\boldsymbol{\theta} | \mathbf{y})] = \mathbb{E}_{q(\boldsymbol{\theta}; \boldsymbol{\nu})} [\log q(\boldsymbol{\theta}; \boldsymbol{\nu}) - \log p(\mathbf{y} | \boldsymbol{\theta}) - \log p(\boldsymbol{\theta})] \quad (23)$$

The r.h.s. of the equation defines our variational objective, also known as ELBO, that can be arranged as follows,

$$\mathcal{L}_{\text{ELBO}}(\boldsymbol{\nu}) = \underbrace{\mathbb{E}_{q(\boldsymbol{\theta}; \boldsymbol{\nu})} \log p(\mathbf{y} | \boldsymbol{\theta})}_{\text{Model fitting term}} - \underbrace{\text{KL} [q(\boldsymbol{\theta}; \boldsymbol{\nu}) \parallel p(\boldsymbol{\theta})]}_{\text{Regularization term}}. \quad (24)$$

This formulation highlights the property of this objective, which is made of two components: the first one is the expected log-likelihood under the approximate posterior  $q$  and measures how the model fits the data. The second term, on the other hand, has the regularization effect of penalizing posteriors that are far from the prior as measured by the KL. Before diving into the challenges of optimization of the ELBO, we shall spend a brief moment discussing the form of the approximating distribution  $q$ . One of the simplest and easier choice is the mean field approximation [42], where each variable  $\theta_i$  is taken to be independent with respect to the remaining  $\boldsymbol{\theta}_{-i}$ . Effectively, this imposes a factorization of the posterior,

$$q(\boldsymbol{\theta}; \boldsymbol{\nu}) = \prod_{i=1}^K q(\theta_i; \boldsymbol{\nu}_i) \quad (25)$$

where  $\boldsymbol{\nu}_i$  is the set of variational parameters for the parameter  $\theta_i$ . On top of this approximation,  $q(\theta_i)$  is often chosen to be Gaussian,

$$q(\theta_i) = \mathcal{N}(\mu_i, \sigma_i^2) \quad (26)$$

Now, the collection of all means and variances  $\{\mu_i, \sigma_i^2\}_{i=1}^K$  defines the set of variational parameters to optimize.

For BNNS the analytic evaluation of the ELBO (and its gradients) is always untractable due the non-linear nature of the expectation of the log-likelihood under the variational distribution. Nonetheless, this can be easily estimated via Monte Carlo integration [72], by sampling  $N_{\text{MC}}$  times from  $q_{\boldsymbol{\nu}}$ ,

$$\mathbb{E}_{q(\boldsymbol{\theta}; \boldsymbol{\nu})} \log p(\mathbf{y} | \boldsymbol{\theta}) \approx \frac{1}{N_{\text{MC}}} \sum_{j=1}^{N_{\text{MC}}} \log p(\mathbf{y} | \tilde{\boldsymbol{\theta}}_j), \quad \text{with } \tilde{\boldsymbol{\theta}}_j \sim q(\boldsymbol{\theta}; \boldsymbol{\nu}) \quad (27)$$

In practice, this is as simple as re-sampling the weights and the biases for all the layers  $N_{\text{MC}}$  times and computing the output for each new sample.

We now have a tractable objective that needs to be optimized with respect to the variational parameters  $\boldsymbol{\nu}$ . Very often the KL term is known, making its differentiation trivial. On the other hand the expectation of the likelihood is not available, making the computation of its gradients more challenging. This problem can be solved using the so-called *reparameterization trick* [92, 54]. The reparameterization trick aims at constructing  $\boldsymbol{\theta}$  as an invertible function  $\mathcal{T}$  of the variational parameters  $\boldsymbol{\nu}$  and of another random variable  $\boldsymbol{\varepsilon}$ , so that  $\boldsymbol{\theta} = \mathcal{T}(\boldsymbol{\varepsilon}; \boldsymbol{\nu})$ . Generally, a  $\mathcal{T}$  that suits this constraint might not exists; Ruiz et al. [90] discuss how to build “weakly” dependent transformation  $\mathcal{T}$  for distributions like Gamma, Beta and Log-normal. For discrete distributions, instead, one could use a continuous relaxation, like the Concrete [69].  $\boldsymbol{\varepsilon}$  is chosen such that its marginal  $p(\boldsymbol{\varepsilon})$  does not depend on the variational parameters. With this parameterization,  $\mathcal{T}$  separates the deterministic components of  $q$  from the stochastic ones, making the computation of its gradient straightforward. For a Gaussian distribution with mean  $\mu$  and variance  $\sigma^2$ ,  $\mathcal{T}$  corresponds to as simple scale-location transformation of an isotropic Gaussian noise,

$$\theta \sim \mathcal{N}(\mu, \sigma^2) \iff \theta = \mu + \sigma \varepsilon \quad \text{with } \varepsilon \sim \mathcal{N}(0, 1). \quad (28)$$

This simple transformation ensures that  $p(\boldsymbol{\varepsilon}) = \mathcal{N}(0, 1)$  does not depends on the variational parameters  $\boldsymbol{\nu} = \{\mu, \sigma^2\}$ . The gradients of the ELBO can be therefore computed as

$$\nabla_{\boldsymbol{\nu}} \mathcal{L}_{\text{ELBO}} = \mathbb{E}_{p(\boldsymbol{\varepsilon})} [\nabla_{\boldsymbol{\theta}} \log p(\mathbf{y} | \boldsymbol{\theta}) |_{\boldsymbol{\theta}=\mathcal{T}(\boldsymbol{\varepsilon}; \boldsymbol{\nu})} \nabla_{\boldsymbol{\nu}} \mathcal{T}(\boldsymbol{\varepsilon}; \boldsymbol{\nu})] - \nabla_{\boldsymbol{\nu}} \text{KL} [q(\boldsymbol{\theta}; \boldsymbol{\nu}) \parallel p(\boldsymbol{\theta})]. \quad (29)$$

The gradient  $\nabla_{\theta} \log p(\mathbf{y} | \theta)$  depends on the model and it can be derived with automatic differentiation tools [1, 78], while  $\nabla_{\nu} \mathcal{T}(\varepsilon; \nu)$  doesn't have any stochastic components and therefore can be known deterministically. Note that the reparameterization trick can be also used when the KL is not analytically available. In that case, we would end up with,

$$\nabla_{\nu} \mathcal{L}_{\text{ELBO}} = \mathbb{E}_{p(\varepsilon)} [\nabla_{\theta} \log p(\mathbf{y} | \theta) + \log q(\theta; \nu) - \log p(\theta)]_{\theta=\mathcal{T}(\varepsilon; \nu)} \nabla_{\nu} \mathcal{T}(\varepsilon; \nu) \quad (30)$$

Roeder et al. [86] argue that when we believe that  $q(\theta; \nu) \approx p(\mathbf{y} | \theta)$ , Eq. (30) should be preferred over Eq. (29) even if computing analytically the KL is possible. Note that this case is very unlikely for BNN posteriors, and that the additional randomness introduced by the Monte Carlo estimation of the KL could be harmful.

In case of large datasets and complex models, the formulation summarized in Eq. (29) can be computationally challenging, due to the evaluation of the likelihood and its gradients  $N_{\text{MC}}$  times. Assuming factorization of the likelihood,

$$p(\mathbf{y} | \theta) = p(\mathbf{y} | f(\mathbf{X}; \theta)) = \prod_{i=1}^N p(y_i | f(\mathbf{x}_i; \theta)) \quad (31)$$

this quantity can be approximated using mini-batching [39, 43]. Recalling  $\mathbf{y}$  as the set of labels of our dataset with  $N$  examples, by taking  $\mathcal{B} \subset \mathbf{y}$  as a random subset of  $\mathbf{y}$ , the likelihood term can be estimated in an unbiased way as

$$\log p_{\theta}(\mathbf{y} | \theta) \approx \frac{N}{M} \sum_{y_i \sim \mathcal{B}} \log p(y_i | \theta). \quad (32)$$

where  $M$  is the number of points in the minibatch. At the cost of increase “randomness”, we can use Eq. (29) to compute the gradients of the ELBO with the minibatch formulation in Eq. (32). Stochastic optimization, e.g. any version of SGD, will converge to a local optimum provided with a decreasing learning rate and sufficient gradient updates [85].