

Bayesian Computation in Deep Learning

Wenlong Chen ^{†}, Bolian Li ^{*¶}, Ruqi Zhang ^{§¶} and Yingzhen Li ^{§†}*

1 Introduction

Bayesian computation has achieved profound success in many modeling tasks with statistics tools such as generalized linear models (Dobson and Barnett, 2018; Nelder and Wedderburn, 1972). Yet these traditional tools fail to produce satisfactory predictions for high-dimensional and highly complex data such as images, speech and videos. Deep Learning (LeCun et al., 2015) provides an attractive solution. At the time of late 2023, deep neural networks achieve accurate predictions for image classification (Dehghani et al., 2023), segmentation (Kirillov et al., 2023) and speech recognition tasks (Zhang et al., 2023). Meanwhile they have also demonstrated an astonishing capability for generating photo-realistic and/or artistic images (Rombach et al., 2022), music (Agostinelli et al., 2023) and videos (Liang et al., 2022). Nowadays deep neural networks have become a standard modeling tool for many of the applications in AI and related fields, and the success of deep learning so far are based on training deterministic deep neural networks on big data. So one might ask: is there a place for Bayesian computation in modern deep learning?

The answer is affirmative. Regarding prediction tasks that deep neural networks are primarily used for, Bayesian model averaging is capable for providing further improvements regarding accuracy (Wilson and Izmailov, 2020). More importantly, quantifying uncertainty

*Equal Contribution

†Imperial College London

¶Purdue University

§Equal Contribution

in predictions is crucial for increasing adoptions of neural networks in safety-critical domains such as healthcare (Gal, 2016). Decision-making tasks based on neural network predictions also require reliable uncertainty estimates (Jaynes, 2003; Savage, 1954). All of these motivate the developments of Bayesian neural networks (Blundell et al., 2015; MacKay, 1992; Neal, 1996), which applies Bayesian modeling to addressing the challenge of uncertainty quantification in deep learning. On the other hand, many popular neural network-based generative models for image/speech/video generation are based on deep latent variable models (Kingma and Welling, 2014; Rezende et al., 2014), whose training requires inferring the unobserved latent variables. In this case posterior inference is a required step in generative model training, sharing many advantages and challenges as Bayesian computation in statistical modeling.

There are three main challenges for Bayesian computation in deep learning. First, the random variables to be inferred can be ultra high-dimensional, e.g., Bayesian neural networks require posterior inference over network weights which can be millions (if not billions) of variables. Second, the likelihood function is highly non-linear as a function of the random variable. In deep latent variable models the likelihood function is defined using a neural network, so that multiple distinct values of a latent variable can be mapped to e.g., the same mean of a Gaussian likelihood, which results in a highly multi-modal posterior. Third, the training data sizes for deep learning tasks are massive, e.g., in millions (Deng et al., 2009) to billions (Schuhmann et al., 2022), so evaluating the likelihood functions on the entire dataset becomes challenging. Traditional MCMC tools such as Metropolis-Hastings (Metropolis et al., 1953) and Gibbs sampling (Gelfand, 2000) fail to address these challenges in both fast and memory-efficient ways. These three challenges motivate the developments of *approximate (Bayesian) inference* tools such as variational inference (Beal, 2003; Jordan et al., 1999; Li, 2018; Zhang et al., 2018), stochastic-gradient MCMC (Ma et al., 2015; Welling and Teh, 2011; Zhang et al., 2020b) and their advanced extensions.

This chapter provides an introduction to approximate inference techniques as Bayesian computation methods applied to deep learning models. We organize the chapter by presenting popular computational methods for (1) Bayesian neural networks and (2) deep generative models, explaining their unique challenges in posterior inference as well as the solutions. More specifically, the sections within this chapter are organised as follows.

- Section 2 introduces Bayesian neural networks (BNNs), with a particular focus on the Bayesian computation tools applied to them. These computational tools are categorised into (1) Markov Chain Monte Carlo (MCMC, Section 2.1) and (2) Variational Inference (VI, Section 2.2). As we shall see, a common technique shared across the two themes is the use of stochastic gradient-based optimization (Bottou, 1998), which enables scaling of BNN computations to millions of datapoints.
- Section 3 discusses a variety of deep generative models which are used for density estimation and/or generating new samples from the data distribution. Specifically, MCMC serves as the key computation tool for sampling and training energy-based, score-based and diffusion models (Sections 3.1 & 3.2), while VI is often employed in training deep latent variable models (Section 3.3).

2 Bayesian Neural Networks

Deep neural networks (DNNs), thanks to their expressiveness in modeling complex functions at scale (LeCun et al., 2015), have gained significant attention in a wide range of applications, e.g., recognition tasks in vision and speech domains (Dehghani et al., 2023; Kirillov et al., 2023; Zhang et al., 2023). In such supervised learning settings, a training dataset of N input-output pairs $\mathcal{D} = \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$ with $\mathbf{x} \in \mathbb{R}^{d_{in}}$ is provided. We consider fitting an L -layer, width d_h , feed-forward DNN $f_{\boldsymbol{\theta}}(\mathbf{x}) : \mathbb{R}^{d_{in}} \rightarrow \mathbb{R}^{d_{out}}$ with weight and bias parameters $\boldsymbol{\theta} = \{\mathbf{W}^l, \mathbf{b}^l\}_{l=1}^L$, whose functional form is

$$f_{\boldsymbol{\theta}}(\mathbf{x}) = \mathbf{W}^L g(\mathbf{W}^{L-1} g(\dots g(\mathbf{W}^1 \mathbf{x} + \mathbf{b}^1)) + \mathbf{b}^{L-1}) + \mathbf{b}^L, \quad (2.1)$$

$\mathbf{W}^1 \in \mathbb{R}^{d_h \times d_{in}}, \mathbf{b}^1 \in \mathbb{R}^{d_{in}}, \mathbf{W}^l \in \mathbb{R}^{d_h \times d_h}, \mathbf{b}^l \in \mathbb{R}^{d_h}, l = 2, \dots, L-1, \mathbf{W}^L \in \mathbb{R}^{d_{out} \times d_h}, \mathbf{b}^L \in \mathbb{R}^{d_{out}}$.

Here a non-linear activation $g(\cdot)$ (e.g., ReLU (Nair and Hinton, 2010)) is used to make $f_{\boldsymbol{\theta}}(\mathbf{x})$ a non-linear function w.r.t. \mathbf{x} . For regression problems, we use the neural network to produce the prediction $\hat{\mathbf{y}} = f_{\boldsymbol{\theta}}(\mathbf{x})$. In such case d_{out} equals to the dimensionality of \mathbf{y} , and the likelihood of the model is typically defined as a Gaussian distribution with a

hyper-parameter σ :

$$p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}) = \mathcal{N}(\mathbf{y}; f_{\boldsymbol{\theta}}(\mathbf{x}), \sigma^2 \mathbf{I}). \quad (2.2)$$

On the other hand, for classification problems, categorical likelihood is used, and the neural network returns the logit of the probability vector of the categorical distribution: assuming $\mathbf{y} \in \{1, \dots, C\}$, then $d_{out} = C$ and

$$[p(\mathbf{y} = 1|\mathbf{x}, \boldsymbol{\theta}), \dots, p(\mathbf{y} = C|\mathbf{x}, \boldsymbol{\theta})] = \text{softmax}(f_{\boldsymbol{\theta}}(\mathbf{x})), \quad \text{softmax}(\mathbf{l}) := \frac{\exp(\mathbf{l})}{\sum_{c=1}^C \exp(\mathbf{l}_c)}. \quad (2.3)$$

The power of DNNs lies in their capability in learning expressive feature representations for the input data. To see this, we rewrite the neural network feed-forward pass up to the $L - 1^{\text{th}}$ layer in Eq. (2.1) as $\Phi_{\boldsymbol{\theta}}(\mathbf{x}) = \mathbf{W}^{L-1}g(\mathbf{W}^{L-2}g(\cdot \cdot \cdot g(\mathbf{W}^1\mathbf{x} + \mathbf{b}^1)) + \mathbf{b}^{L-2}) + \mathbf{b}^{L-1}$. Then the neural network $f_{\boldsymbol{\theta}}(\mathbf{x}) = \mathbf{W}^L\Phi_{\boldsymbol{\theta}}(\mathbf{x}) + \mathbf{b}^L$ and the associated likelihood model can be viewed as a (generalized) linear regression model with non-linear features $\Phi_{\boldsymbol{\theta}}(\mathbf{x})$. Different from generalized linear regression models, however, these non-linear features are also learned jointly with the last-layer parameters $\{\mathbf{W}^L, \mathbf{b}^L\}$. While this joint learning procedure returns expressive features and improved predictive accuracy, the weight and bias parameters $\{\mathbf{W}^l, \mathbf{b}^l\}_{l=1}^L$ are often less interpretable in any statistical sense, although sparsity methods may provide explanations akin to feature selection (Ghosh et al., 2019; Lemhadri et al., 2021). See later paragraphs for a further discussion in Bayesian computation context.

A typical estimation procedure for the parameters $\boldsymbol{\theta}$ is maximum likelihood estimation (MLE), which finds the optimal neural network parameters as

$$\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta}} \sum_{n=1}^N \log p(\mathbf{y}_n|\mathbf{x}_n, \boldsymbol{\theta}).$$

However, when the dataset size is substantially smaller than the network parameter size, MLE estimates may overfit. Maximum a Posteriori (MAP) solutions can help alleviate the overfitting issue, however, both MLE and MAP provide point estimates only for the network parameters $\boldsymbol{\theta}$, leaving the underlying uncertainty uncaptured.

Bayesian inference is a principled framework to obtain reliable uncertainty in DNNs (Gal,

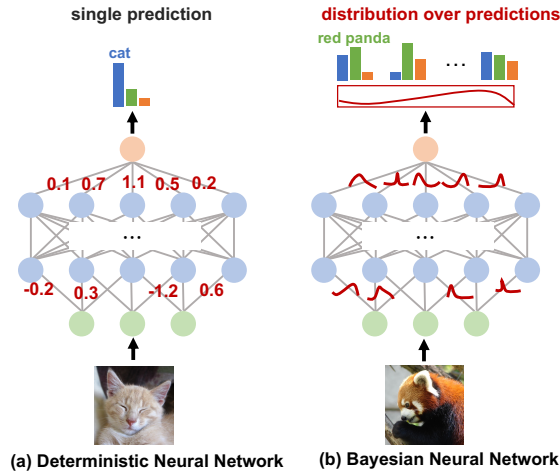


Figure 1: Difference between standard deep neural network and Bayesian neural network. (2016). Instead of using a point estimate of θ , Bayesian inference computes a posterior distribution over θ to explain the observations with uncertainty estimates (See Figure 1):

$$p(\theta|\mathcal{D}) = \frac{p(\theta)p(\mathcal{D}|\theta)}{p(\mathcal{D})} = \frac{p(\theta)p(\mathcal{D}|\theta)}{\int p(\theta)p(\mathcal{D}|\theta)d\theta}. \quad (2.4)$$

Here $p(\theta)$ is the prior distribution representing our belief about the model parameters without observing any data points, and $p(\mathcal{D}|\theta)$ is the likelihood. Often we assume i.i.d. likelihood, which means $p(\mathcal{D}|\theta) := \prod_{n=1}^N p(\mathbf{y}_n|\mathbf{x}_n, \theta)$. In the rest of the chapter we also write $p(\mathbf{y}|f_{\theta}(\mathbf{x})) := p(\mathbf{y}|\mathbf{x}, \theta)$ since the distributional parameters of the likelihood mainly depends on the output of the neural network. The resulting neural networks are called Bayesian neural networks (BNNs) (Hinton and Van Camp, 1993; MacKay, 1992; Neal, 1996). For predictions, the uncertainty in θ translates to the predictive uncertainty for $(\mathbf{x}^*, \mathbf{y}^*)$, via the following *posterior predictive distribution*:

$$p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D}) = \int p(\mathbf{y}^*|\mathbf{x}^*, \theta)p(\theta|\mathcal{D})d\theta. \quad (2.5)$$

With the posterior predictive distribution (Eq. (2.5)) one can compute uncertainty measures from it. In regression problems a typical uncertainty measure is the variance of the prediction $V[\mathbf{y}^*|\mathbf{x}^*] = E_{p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D})}[(\mathbf{y}^* - \boldsymbol{\mu}^*)(\mathbf{y}^* - \boldsymbol{\mu}^*)^\top]$ with $\boldsymbol{\mu}^* = E_{p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D})}[\mathbf{y}^*]$ as the predictive posterior mean. In classification where the likelihood $p(\mathbf{y}|\mathbf{x}, \theta)$ is typically categorical,

one can also use the entropy $H[p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D})] = -\sum_{c=1}^C p(\mathbf{y}^* = c|\mathbf{x}^*, \mathcal{D}) \log p(\mathbf{y}^* = c|\mathbf{x}^*, \mathcal{D})$ to express uncertainty. These two quantities represent *total* uncertainty, which can be further decomposed into *epistemic* and *aleatoric* uncertainty measures, for example (Gal, 2016; Houlsby et al., 2011):

$$H[p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D})] = \underbrace{E_{p(\boldsymbol{\theta}|\mathcal{D})}[H[p(\mathbf{y}^*|\mathbf{x}^*, \boldsymbol{\theta})]]}_{\text{(aleatoric uncertainty)}} + \underbrace{I[\mathbf{y}^*; \boldsymbol{\theta}|\mathbf{x}^*, \mathcal{D}]}_{\text{(epistemic uncertainty)}}. \quad (2.6)$$

The conditional entropy $E_{p(\boldsymbol{\theta}|\mathcal{D})}[H[p(\mathbf{y}^*|\mathbf{x}^*, \boldsymbol{\theta})]] = E_{p(\boldsymbol{\theta}|\mathcal{D})}E_{p(\mathbf{y}^*|\mathbf{x}^*, \boldsymbol{\theta})}[-\log p(\mathbf{y}^*|\mathbf{x}^*, \boldsymbol{\theta})]$ reveals the average “variability” (under posterior $p(\boldsymbol{\theta}|\mathcal{D})$) of each neural network’s predictive distribution $p(\mathbf{y}^*|\mathbf{x}^*, \boldsymbol{\theta})$, and it aims to quantify the intrinsic randomness in data due to e.g., noisy measurements. On the other hand, there are two explanations for epistemic uncertainty, originated from two equivalent definitions of mutual information. The first definition reads

$$I[\mathbf{y}^*; \boldsymbol{\theta}|\mathbf{x}^*, \mathcal{D}] = \mathbb{E}_{p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D})}[\text{KL}[p(\boldsymbol{\theta}|\mathbf{y}^*, \mathbf{x}^*, \mathcal{D})||p(\boldsymbol{\theta}|\mathcal{D})]],$$

which quantifies the expected information gain for $\boldsymbol{\theta}$ (i.e., reduction of uncertainty in $p(\boldsymbol{\theta}|\mathcal{D})$) when adding the prediction \mathbf{y}^* on \mathbf{x}^* to the observations. This concept of “epistemic uncertainty reduction” is particularly useful in data acquisition tasks following Bayesian principles (Houlsby et al., 2011). The second definition of mutual information comes from re-arranging terms in the first definition:

$$I[\mathbf{y}^*; \boldsymbol{\theta}|\mathbf{x}^*, \mathcal{D}] = \mathbb{E}_{p(\boldsymbol{\theta}|\mathcal{D})}[\text{KL}[p(\mathbf{y}^*|\mathbf{x}^*, \boldsymbol{\theta})||p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D})]].$$

It measures the average deviation (under the posterior $p(\boldsymbol{\theta}|\mathcal{D})$) of a single neural network’s prediction $p(\mathbf{y}^*|\mathbf{x}^*, \boldsymbol{\theta})$ from the posterior predictive $p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D})$. Therefore a high value of mutual information means the neural networks parameterized by different $\boldsymbol{\theta}$ from high density regions of $p(\boldsymbol{\theta}|\mathcal{D})$ have large disagreement in their prediction results, and the BNN is said to be epistemically uncertain about its predictions.

Unfortunately, Eq. (2.4) is intractable since the marginal likelihood requires a complex integral $p(\mathcal{D}) = \int p(\boldsymbol{\theta})p(\mathcal{D}|\boldsymbol{\theta})d\boldsymbol{\theta}$, which has no analytic solution due to non-linearities in DNNs. Furthermore, numerical integration is infeasible due to extremely high-dimensional

θ , usually over million or even billions in modern DNNs (Dehghani et al., 2023; Krizhevsky et al., 2012). Therefore, we have to resort to approximate inference techniques (Li, 2018), which uses an approximate posterior distribution $q(\theta) \approx p(\theta|\mathcal{D})$ to replace the exact posterior in Bayesian inference tasks. In such case posterior predictive distribution (Eq. (2.5)) is approximated by (typically also with Monte Carlo estimate):

$$p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D}) \approx \int p(\mathbf{y}^*|\mathbf{x}^*, \theta)q(\theta)d\theta \approx \frac{1}{M} \sum_{m=1}^M p(\mathbf{y}^*|\mathbf{x}^*, \theta_m), \quad \theta_m \sim q(\theta). \quad (2.7)$$

This approximate predictive posterior can then be used subsequently for computing the uncertainty measures such as predictive variance or entropy in Eq. (2.6). We will discuss two classes of approximations: (1) stochastic gradient MCMC (Chen et al., 2014; Welling and Teh, 2011), which constructs $q(\theta) = \frac{1}{M} \sum_{m=1}^M \delta(\theta = \theta_m)$ using samples $\{\theta_m\}_{m=1}^M$ obtained via an MCMC algorithm, and (2) variational inference (Beal, 2003; Jordan et al., 1999) where $q(\theta)$ often has a parametric form and is obtained via an optimization procedure.

Before proceeding to the technical details, an interesting note is that for BNNs, the quality of posterior approximation $q(\theta) \approx p(\theta|\mathcal{D})$ is mostly evaluated via the corresponding approximate posterior predictive (Eq. (2.7)) rather than by direct inspection. First, symmetry exists in neural network parameters θ (i.e., there exists $\theta_1 \neq \theta_2$ such that $f_{\theta_1} = f_{\theta_2}$, e.g., by swapping two hidden units of a hidden layer in the neural network of f_{θ_1} and denoting the corresponding neural network’s weight by θ_2 (Phuong and Lampert, 2020; Rolnick and Kording, 2020)), meaning that the exact posterior $p(\theta|\mathcal{D})$ is highly multi-modal. Second and most distinctively from traditional statistical models, due to the emphasis of prediction accuracy for deep learning models, the network weights θ are not designed to have an associated physical meaning, and in most cases even the exact posterior $p(\theta|\mathcal{D})$ is not interpretable by humans. Instead, the posterior predictive distribution (Eq. (2.7)) directly expresses the uncertainty in prediction, which can then be analyzed and used in further decision-making tasks based on the predictions.

Algorithm 1 SG-MCMC methods for sampling from BNN posteriors

Require: Dataset $\mathcal{D} = \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$, initialized parameter $\boldsymbol{\theta}_1$, batch-size B , SG-MCMC step sizes $\{\alpha_t\}$, total steps T .

- 1: **for** $t = 1, \dots, T$ **do**
 - 2: Sample a mini-batch of datapoints $\Xi = \{(\mathbf{x}, \mathbf{y})\} \sim \mathcal{D}^B$ with $|\Xi| = B$
 - 3: Get the stochastic gradient $\nabla_{\boldsymbol{\theta}_t} \tilde{U}(\boldsymbol{\theta}_t)$ via automatic differentiation
 - 4: Obtain $\boldsymbol{\theta}_{t+1}$ through an SG-MCMC update, using the stochastic gradient $\nabla_{\boldsymbol{\theta}_t} \tilde{U}(\boldsymbol{\theta}_t)$ and the step size α_t
 - 5: **end for**
-

2.1 Stochastic Gradient MCMC

Classical MCMC methods require computation over the entire dataset to obtain the true likelihood or its gradient, making them prohibitively expensive for large datasets. For example, the standard Metropolis-Hastings (MH) algorithm is unsuitable because the accept/reject step necessitates computing the likelihood, which involves summing across the entire dataset. To mitigate this computational burden, minibatched MH algorithms have been developed Korattikara et al. (2014); Maclaurin and Adams (2014); Quiroz et al. (2018); Zhang et al. (2020a). These algorithms utilize a subset of the data in the MH step to reduce costs. However, they often rely on strong assumptions, such as bounded log-likelihood or bounded gradients, which render them impractical for deep neural networks. Despite MCMC methods, notably Hamiltonian Monte Carlo (Duane et al., 1987; Neal et al., 2011), being regarded as the gold standard for Bayesian computations, they are rarely applied in modern BNNs due to their large computational cost. To address this challenge, Stochastic Gradient MCMC (SG-MCMC) methods (Chen et al., 2014; Ma et al., 2015; Welling and Teh, 2011; Zhang et al., 2020b) have been proposed. These methods introduce stochastic gradients into MCMC sampling steps, relying only on a subset of data at each iteration. Consequently, they demand significantly less memory compared to standard MCMC approaches. SG-MCMC enables the use of sampling methods in large-scale Bayesian Neural Networks, such as ResNets (He et al., 2016). In this section, we introduce three popular variants of SG-MCMC algorithms in deep learning. A summary of the procedure of SG-MCMC algorithms is provided in Algorithm 1.

2.1.1 Stochastic Gradient Langevin Dynamics

Stochastic Gradient Langevin Dynamics (SGLD) (Welling and Teh, 2011) stands as the first SG-MCMC algorithm. In contrast to the full energy function $U(\boldsymbol{\theta}) = -\sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \log p(\mathbf{x}|\boldsymbol{\theta}) - \log p(\boldsymbol{\theta})$ for the exact posterior distribution $p(\boldsymbol{\theta}|\mathcal{D}) \propto \exp(-U(\boldsymbol{\theta}))$, the SGLD algorithm considers a mini-batch energy function with a mini-batch of datapoints $\Xi \subseteq \mathcal{D}$:

$$\tilde{U}(\boldsymbol{\theta}) = -\frac{|\mathcal{D}|}{|\Xi|} \sum_{(\mathbf{x}, \mathbf{y}) \in \Xi} \log p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}) - \log p(\boldsymbol{\theta}), \quad (2.8)$$

By resampling the minibatch Ξ at each iteration, the update rule of SGLD follows the *underdamped Langevin dynamics*, replacing the true gradient with the stochastic gradient:

$$\boldsymbol{\theta}_{t+1} \leftarrow \boldsymbol{\theta}_t - \alpha_t \nabla_{\boldsymbol{\theta}_t} \tilde{U}(\boldsymbol{\theta}_t) + \sqrt{2\alpha_t} \cdot \boldsymbol{\epsilon}_t, \quad (2.9)$$

where α_t is the step size at iteration t and $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is the standard Gaussian noise. Compared with the update rule of stochastic gradient descent (SGD), the only difference is the addition of the appropriate Gaussian noise. Welling and Teh (2011) shows that when the step size α_t is sufficiently small, SGLD becomes standard unadjusted Langevin dynamics (Roberts and Tweedie, 1996), as the Gaussian noise term dominates over the stochastic noise. To ensure a low asymptotic estimation error, a decaying step size schedule is typically applied (Bottou, 1998; Ma et al., 2015; Welling and Teh, 2011), satisfying (i) $\sum_{t=1}^{\infty} \alpha_t = \infty$ and (ii) $\sum_{t=1}^{\infty} \alpha_t^2 < \infty$. The first requirement guarantees that the sampler will reach the high probability areas regardless of initialization, and the second requirement ensures a small asymptotic error (Welling and Teh, 2011).

2.1.2 Cyclical Stochastic Gradient MCMC

As mentioned in the last section, decreasing step sizes for SG-MCMC is typically required to ensure asymptotic accuracy of posterior estimation. However, small step sizes will significantly reduce the sampler’s ability to explore the entire domain and to exploit each mode in detail efficiently, resulting in large approximation errors in finite time (Chen et al., 2015;




Figure 2: Comparison between the cyclical stepsize schedule (red) and the traditional decreasing stepsize schedule (blue) for SG-MCMC algorithms. Adapted from Zhang et al. (2020b). Used with the kind permission of Ruqi Zhang.

Vollmer et al., 2016). This issue is exacerbated when estimating Bayesian neural network posteriors, which are complicated and highly multi-modal (Izmailov et al., 2021; Neal, 1996).

To solve this problem, Zhang et al. (2020b) proposes a cyclical step size schedule for SG-MCMC, depicted in Figure 2. This schedule defines the step size at iteration t as

$$\alpha_t = \frac{\alpha_0}{2} \left[\cos \left(\frac{\pi \operatorname{mod}(t-1, \lceil T/M \rceil)}{\lceil T/M \rceil} \right) + 1 \right], \quad (2.10)$$

where α_0 is the initial step size, M is the number of cycles and T is the number of total iterations. The cyclical step size schedule enables the sampler to efficiently explore various modes and accurately characterize local modes at a fine scale.

To efficiently sample from the multi-modal posteriors in Bayesian neural networks, other strategies include SG-MCMC methods combined with parallel tempering (Deng et al., 2020a) and flat histograms (Deng et al., 2020b).

2.1.3 Stochastic Gradient Hamiltonian Monte Carlo

Hamiltonian Monte Carlo (HMC) has achieved the gold standard of performance with small neural networks [Here, refer to the Chapter of HMC]. It incorporates a kinetic energy term, characterized by a set of “momentum” auxiliary variables. The naive stochastic gradient

HMC adopts a mini-batch Hamiltonian function:

$$\tilde{H}(\boldsymbol{\theta}, \mathbf{r}) = \tilde{U}(\boldsymbol{\theta}) + \frac{1}{2} \mathbf{r}^T \mathbf{M}^{-1} \mathbf{r} = -\frac{|\mathcal{D}|}{|\Xi|} \sum_{(\mathbf{x}, \mathbf{y}) \in \Xi} \log p(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta}) - \log p(\boldsymbol{\theta}) + \frac{1}{2} \mathbf{r}^T \mathbf{M}^{-1} \mathbf{r}, \quad (2.11)$$

where \mathbf{r} is the momentum vector and \mathbf{M} is a positive definite mass matrix. Despite its intuitive appeal, Chen et al. (2014) indicates that the naive adaption of HMC algorithm to the stochastic gradient version can diverge from the target distribution. This is because the joint distribution of parameters and momentum (i.e., $\pi(\boldsymbol{\theta}, \mathbf{r}) \propto \exp(-\tilde{H}(\boldsymbol{\theta}, \mathbf{r}))$) are not invariant for this Markov chain due to the introduction of stochastic noise. For corrections, Chen et al. (2014) introduces a “friction” term $C\mathbf{M}^{-1}\mathbf{r}$ where C is the friction weight, to the momentum updating:

$$\frac{d\boldsymbol{\theta}}{dt} = \mathbf{M}^{-1} \mathbf{r} \quad \text{and} \quad \frac{d\mathbf{r}}{dt} = -\nabla_{\boldsymbol{\theta}} U(\boldsymbol{\theta}) - C\mathbf{M}^{-1} \mathbf{r} + \sqrt{2C} \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (2.12)$$

This friction term helps decrease the Hamiltonian $\tilde{H}(\boldsymbol{\theta}, \mathbf{r})$, which counters the effects of random noise. The modified dynamics is commonly denoted as *second-order Langevin dynamics* (Wang and Uhlenbeck, 1945).

2.2 Variational Inference

Variational inference (VI) (Beal, 2003; Jordan et al., 1999) is a class of approximate inference techniques, which approximates the exact posterior using a tractable distribution, typically from a parametric distribution family $\mathcal{Q} := \{q_{\phi}(\boldsymbol{\theta})\}$ such as factorized Gaussians. The best approximate posterior $q_{\phi^*}(\boldsymbol{\theta})$, which is used as $q(\boldsymbol{\theta})$ in Eq. (2.7) and also named as the optimal variational distribution, is selected to be the one that is the closest to the posterior $p(\boldsymbol{\theta} | \mathcal{D})$ within the parametric family as measured by a divergence measure between distributions. The Kullback-Leibler (KL) divergence (Kullback and Leibler, 1951) is a popular choice for VI: here the approximate posterior is obtained by minimizing the KL-divergence with respect to the parameters ϕ of the chosen parametric distributions family:

$$\phi^* = \arg \min_{\phi \in \mathcal{Q}} KL[q_{\phi}(\boldsymbol{\theta}) || p(\boldsymbol{\theta} | \mathcal{D})], \quad KL[q_{\phi}(\boldsymbol{\theta}) || p(\boldsymbol{\theta} | \mathcal{D})] = \int q_{\phi}(\boldsymbol{\theta}) \log \frac{q_{\phi}(\boldsymbol{\theta})}{p(\boldsymbol{\theta} | \mathcal{D})} d\boldsymbol{\theta}. \quad (2.13)$$

Again $p(\boldsymbol{\theta}|\mathcal{D})$ is intractable due to the intractable marginal likelihood $p(\mathcal{D})$, meaning that direct minimization of the KL-divergence is intractable. Fortunately, this KL-divergence minimization task is equivalent to maximizing the following Evidence Lower Bound (ELBO), which is obtained by subtracting the KL-divergence from the log marginal-likelihood which is constant w.r.t. the variational parameters ϕ :

$$\begin{aligned}\mathcal{L}_{ELBO}(\phi) &:= \log p(\mathcal{D}) - KL[q_{\phi}(\boldsymbol{\theta})\|p(\boldsymbol{\theta}|\mathcal{D})] \\ &= E_{q_{\phi}(\boldsymbol{\theta})}[\log p(\mathcal{D}|\boldsymbol{\theta})] - KL[q_{\phi}(\boldsymbol{\theta})\|p(\boldsymbol{\theta})].\end{aligned}\tag{2.14}$$

In Eq. (2.14), the ELBO is decomposed into (i) the expectation of log-likelihood under q_{ϕ} , optimization of which encourages better data fitting, and (ii) a KL regularizer, optimization of which encourages q_{ϕ} to be close to the prior. Combining both terms, the ELBO optimization encourages q_{ϕ} to find small but meaningful deviations from the prior that can explain the observed data. In practice, the expected log-likelihood term (i) in the ELBO can not be analytically evaluated. Therefore an unbiased Monte Carlo estimate is applied to the ELBO with M samples from the variational posterior (Blundell et al., 2015; Graves, 2011):

$$\hat{\mathcal{L}}_{ELBO}(\phi) = \frac{1}{M} \sum_{m=1}^M \log p(\mathcal{D}|\boldsymbol{\theta}_m) - KL[q_{\phi}(\boldsymbol{\theta})\|p(\boldsymbol{\theta})], \quad \boldsymbol{\theta}_m \stackrel{\text{iid}}{\sim} q_{\phi}(\boldsymbol{\theta}).\tag{2.15}$$

The KL regularizer term in ELBO may be evaluated analytically for some choices of prior and variational family (e.g., when they are both Gaussians). When analytic solution is not available, this KL term can also be estimated with Monte Carlo.

VI will return the exact posterior if $p(\boldsymbol{\theta}|\mathcal{D}) \in \mathcal{Q}$, which is not the case in practice. Instead, due to limited computational budget one would prefer simple q distributions that enjoy fast computations and low memory costs. In BNN context, one of the most popular choices is mean-field Gaussian approximation (i.e., factorized Gaussian) (Blundell et al., 2015):

$$q_{\phi}(\boldsymbol{\theta}) = \prod_{l=1}^L q_{\phi}(\mathbf{W}^l) q_{\phi}(\mathbf{b}^l), \quad q_{\phi}(\mathbf{W}^l) = \prod_{ij} \mathcal{N}(\mathbf{W}_{ij}^l; \mathbf{M}_{ij}^l, (\mathbf{S}_{ij}^l)^2), \quad q_{\phi}(\mathbf{b}^l) = \prod_j \mathcal{N}(\mathbf{b}_j^l; \mathbf{m}_j^l, (\mathbf{s}_j^l)^2).\tag{2.16}$$

Here the variational parameters are $\phi = \{\mathbf{M}^l, \mathbf{m}^l, \mathbf{S}^l, \mathbf{s}^l\}_{l=1}^L$ which are optimized by maxi-

mizing the ELBO, typically with Monte Carlo estimates (Eq. (2.15)). Therefore, compared with deterministic DNNs which directly optimizes the weight parameters $\boldsymbol{\theta} = \{\mathbf{W}^l, \mathbf{b}^l\}_{l=1}^L$, mean-field VI for BNNs doubles the amount of parameters to be optimized (hence twice amount of memory consumption), due to the additional variance parameters $\{\mathbf{S}^l, \mathbf{s}^l\}_{l=1}^L$.

The optimization process of ELBO is conducted via gradient ascent, and in order to apply back-propagation with e.g., automatic differentiation techniques (Baydin et al., 2018), the *reparameterization trick* (Kingma and Welling, 2014) is introduced for a range of variational distribution families including Gaussians. This approach assumes the sampling operation $\boldsymbol{\theta} \sim q_\phi(\boldsymbol{\theta})$ is defined as transforming an auxiliary noise variable $\boldsymbol{\epsilon} \sim p_{base}(\boldsymbol{\epsilon})$ through a function $T_\phi(\boldsymbol{\epsilon})$ that is differentiable w.r.t. $\boldsymbol{\phi}$:

$$\boldsymbol{\theta} \sim q_\phi(\boldsymbol{\theta}) \quad \Leftrightarrow \quad \boldsymbol{\theta} = T_\phi(\boldsymbol{\epsilon}), \quad \boldsymbol{\epsilon} \sim p_{base}(\boldsymbol{\epsilon}). \quad (2.17)$$

For mean-field Gaussians, e.g., $q_\phi(\mathbf{W}^l)$ in Eq. (2.16), this transformation is defined as $\mathbf{W}^l = T_\phi(\mathbf{E}^l) := \mathbf{M}^l + \sqrt{\mathbf{S}^l} \odot \mathbf{E}^l$, with $\mathbf{E}_{ij}^l \sim \mathcal{N}(0, 1)$, and $\sqrt{\cdot}$ and \odot denote element-wise square-root and multiplication, respectively. This allows us to rewrite the Monte Carlo estimate of the expected log-likelihood term in Eq. (2.14) using change-of-variable rules, where we collect the mean and variance parameters into $\boldsymbol{\mu} := \{\mathbf{M}^l, \mathbf{m}^l\}_{l=1}^L$ and $\boldsymbol{\sigma} := \{\sqrt{\mathbf{S}^l}, \sqrt{\mathbf{s}^l}\}_{l=1}^L$:

$$E_{q_\phi(\boldsymbol{\theta})}[\log p(\mathcal{D}|\boldsymbol{\theta})] \approx \frac{1}{M} \sum_{m=1}^M \log p(\mathcal{D}|\boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\epsilon}_m), \quad \boldsymbol{\epsilon}_m \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (2.18)$$

The corresponding gradient w.r.t. \mathbf{M}^l for example can then be computed via back-propagation. Recall that $p(\mathcal{D}|\boldsymbol{\theta}) := \prod_{n=1}^N p(\mathbf{y}_n|\mathbf{x}_n, \boldsymbol{\theta})$ with $p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}) := p(\mathbf{y}|f_\theta(\mathbf{x}))$, then

$$\nabla_{\mathbf{M}^l} E_{q_\phi(\boldsymbol{\theta})}[\log p(\mathcal{D}|\boldsymbol{\theta})] \approx \frac{1}{M} \sum_{m=1}^M \sum_{n=1}^N \nabla_{\mathbf{M}^l} f \nabla_f \log p(\mathbf{y}_n|f_{\boldsymbol{\mu}+\boldsymbol{\sigma} \odot \boldsymbol{\epsilon}_m}(\mathbf{x}_n)), \quad \boldsymbol{\epsilon}_m \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}).$$

In practice, using $M = 1$ sample, together with proper initialization of the $\boldsymbol{\phi}$ parameters, is sufficient for the gradient estimation. Similar operations can be applied to computing the gradient of the KL term in the ELBO when Monte Carlo estimate is also required.

Lastly, evaluating $\log p(\mathcal{D}|\boldsymbol{\theta})$ requires processing the entire dataset \mathcal{D} which can be costly

Algorithm 2 SGD Training for a BNN with mean-field Gaussian VI

Require: Dataset $\mathcal{D} = \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$, initialised variational parameters $\phi = \{\boldsymbol{\mu}, \boldsymbol{\sigma}\}$, batch-size B , SGD step sizes $\{\alpha_t\}$, SGD total steps T .

- 1: **for** $t = 1, \dots, T$ **do**
 - 2: Compute the KL term in the ELBO: $\tilde{\mathcal{L}}_{ELBO}(\phi) \leftarrow -KL[q_\phi(\boldsymbol{\theta})\|p(\boldsymbol{\theta})]$
 - 3: Sample BNN weights with the reparameterization trick: $\boldsymbol{\theta} \leftarrow \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\epsilon}$, $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 4: Sample a mini-batch of datapoints $\Xi = \{(\mathbf{x}, \mathbf{y})\} \sim \mathcal{D}^B$ with $|\Xi| = B$
 - 5: **for** $(\mathbf{x}, \mathbf{y}) \in \Xi$ **do**
 - 6: Add-in data likelihood: $\tilde{\mathcal{L}}_{ELBO}(\phi) \leftarrow \tilde{\mathcal{L}}_{ELBO}(\phi) + \frac{|\mathcal{D}|}{|\Xi|} \log p(\mathbf{y}|f_\theta(\mathbf{x}))$
 - 7: **end for**
 - 8: Compute SGD updates via automatic differentiation: $\phi \leftarrow \phi + \alpha_t \nabla_\theta \tilde{\mathcal{L}}_{ELBO}(\phi)$
 - 9: **end for**
-

in big data setting. Similar to training deterministic DNNs and running SG-MCMC for BNNs, stochastic optimization methods such as stochastic gradient descent (SGD) (Bottou, 1998) are also applicable to variational inference. With $\Xi \subseteq \mathcal{D}$ a mini-batch sampled from the entire dataset, the ELBO can be estimated as (with $M = 1$ Monte Carlo sample and the reparameterization trick):

$$\tilde{\mathcal{L}}_{ELBO}(\phi) = \frac{|\mathcal{D}|}{|\Xi|} \sum_{(\mathbf{x}, \mathbf{y}) \in \Xi} \log p(\mathbf{y}|f_{\boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\epsilon}}(\mathbf{x})) - KL[q_\phi(\boldsymbol{\theta})\|p(\boldsymbol{\theta})], \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (2.19)$$

Automatic differentiation can be directly applied to $\tilde{\mathcal{L}}_{ELBO}(\phi)$, and the optimization of this ELBO estimate has the same time complexity as running stochastic gradient descent training of deterministic DNNs. All the above techniques combined enable BNNs with mean-field Gaussian VI to scale up to modern DNN architectures such as ResNets (He et al., 2016), and a summary of the training procedure is provided in Algorithm 2.

2.2.1 Alternative Divergences

Standard VI based on minimizing $KL[q_\phi(\boldsymbol{\theta})\|p(\boldsymbol{\theta}|\mathcal{D})]$, while enjoying many computational advantages, tends to underestimate posterior uncertainty (Turner and Sahani, 2011). The Gaussian approximation from standard VI often captures only one mode of the exact posterior, which has been referred to as the mode-seeking phenomenon (Hernández-Lobato et al., 2015; Li and Turner, 2016). This issue can be mitigated by extending the VI framework to

other divergences. For example, α -divergence is a subset of f -divergence family (Csiszár, 1963) with a hyper-parameter α determining the degree of the posterior mass coverage of the variational distribution (Rényi, 1961; Van Erven and Harremoës, 2014):

$$D_\alpha[p(\boldsymbol{\theta}|\mathcal{D})\|q_\phi(\boldsymbol{\theta})] = \frac{1}{(\alpha - 1)} \log \int p(\boldsymbol{\theta}|\mathcal{D})^\alpha q_\phi(\boldsymbol{\theta})^{1-\alpha} d\boldsymbol{\theta}. \quad (2.20)$$

When α is set to approach 1 and 0, we recover the two special cases: $KL[q_\phi(\boldsymbol{\theta})\|p(\boldsymbol{\theta}|\mathcal{D})]$ (exclusive KL as used in standard VI) and $KL[p(\boldsymbol{\theta}|\mathcal{D})\|q_\phi(\boldsymbol{\theta})]$ (inclusive KL) respectively. There exist other equivalent α -divergence definitions, e.g., Amari’s (Amari, 1985), which also include the two KL divergences as special cases.

Direct minimization of the α -divergence in Eq. (2.20) is intractable again due to the intractable $p(\mathcal{D})$. This is mitigated by again an equivalent maximization task on another lower bound of the log marginal likelihood, which can also be estimated with Monte Carlo, potentially with the reparameterization trick as well (Li and Turner, 2016):

$$\begin{aligned} \mathcal{L}_\alpha(\phi) &:= \log p(\mathcal{D}) - D_\alpha[p(\boldsymbol{\theta}|\mathcal{D})\|q_\phi(\boldsymbol{\theta})] \\ &= \frac{1}{1 - \alpha} \log E_{q_\phi(\boldsymbol{\theta})} \left[\left(\frac{p(\mathcal{D}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{q_\phi(\boldsymbol{\theta})} \right)^{1-\alpha} \right] \\ &\approx \frac{1}{1 - \alpha} \log \frac{1}{M} \sum_{m=1}^M \left(\frac{p(\mathcal{D}|\boldsymbol{\theta}_m)p(\boldsymbol{\theta}_m)}{q_\phi(\boldsymbol{\theta}_m)} \right)^{1-\alpha}, \quad \boldsymbol{\theta}_m \stackrel{\text{iid}}{\sim} q_\phi(\boldsymbol{\theta}) \end{aligned} \quad (2.21)$$

Figure 3 illustrates the coverage of mean-field approximations fitted with α -divergences, where the target distribution is a correlated Gaussian. Indeed by decreasing α , one can interpolate between mode-seeking and mass-covering behaviors for the approximate posterior.




Figure 3: Factorized Gaussians fitted by minimizing α -divergences with different α 's for a correlated 2D Gaussian target:

$$\mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 2.0 & 1.5 \\ 1.5 & 1.6 \end{pmatrix} \right).$$

2.2.2 Distribution Family

One drawback of using mean-field Gaussian approximations is the negligence of covariance structures in the exact posterior. In this regard a Gaussian distribution with a full covariance matrix is favorable. However, storing and decomposing a covariance matrix require $O(d^2)$ memory and $O(d^3)$ time, respectively, which quickly become computationally intractable even if the number of BNN weights d is around tens of thousands. Recall that modern neural networks for image recognition has millions, if not billions, of parameters (Dehghani et al., 2023; Krizhevsky et al., 2012). Therefore Gaussians with freely parameterized full covariance matrices are not viable solutions, and further approximations are needed.

Recent attention on developing correlated Gaussian approximate posteriors focuses on so called “low-rank + diagonal” structure for the covariance matrices. In other words, instead of searching over the entire space of positive semi-definite matrices for the covariance matrix, the Gaussian variational distribution $q_\phi(\boldsymbol{\theta})$ has the following structure (Tomczak et al., 2020):

$$q_\phi(\boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\phi}; \boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad \boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}^\top + \mathbf{D}, \quad (2.22)$$

where $\mathbf{L} \in \mathbb{R}^{d \times r}$ is a low-rank matrix (with $r \ll d$) and \mathbf{D} is a diagonal matrix with positive diagonal entries. Here the variational parameters are $\boldsymbol{\phi} = \{\boldsymbol{\mu}, \mathbf{L}, \mathbf{D}\}$ which can be freely optimized using ELBO. In practice correlated Gaussian approximations are only applied layer-wise, i.e., the q_ϕ distribution still factorizes over layers, but within each layer a low-rank Gaussian approximate posterior is fitted.

Another approach to introduce structured Gaussian approximations is to use matrix normal distributions. Notice that the popular choice of isotropic Gaussian prior $p(\mathbf{W}^l) = \mathcal{N}(\text{vec}(\mathbf{W}^l); \mathbf{0}, \sigma^2 \mathbf{I}_{d_o d_i})$ for the l^{th} layer weight $\mathbf{W}^l \in \mathbb{R}^{d_o \times d_i}$ has an equivalent matrix normal form $\mathcal{MN}(\mathbf{W}^L; \mathbf{0}, \sigma_r^2 \mathbf{I}_{d_o}, \sigma_c^2 \mathbf{I}_{d_i})$ with $\sigma_r > 0$ and $\sigma_c > 0$ the row and column standard deviations satisfying $\sigma = \sigma_r \sigma_c$. This inspires the design of matrix normal approximate posteriors $q_\phi(\mathbf{W}^l) = \mathcal{MN}(\mathbf{W}^l; \mathbf{M}^l, \Sigma_r, \Sigma_c)$ with $\Sigma_r \in \mathbb{R}^{d_o \times d_o}$ and $\Sigma_c \in \mathbb{R}^{d_i \times d_i}$, reducing the costs to $\mathcal{O}(d_i^2 + d_o^2)$ memory and $\mathcal{O}(d_i^3 + d_o^3)$ time instead of $\mathcal{O}(d_i^2 d_o^2)$ and $\mathcal{O}(d_i^3 d_o^3)$ respectively for using full covariance Gaussian approximations. Further low-rank approximation techniques

can be applied to parameterize Σ_r and Σ_c to further reduce the complexity figures (Louizos and Welling, 2017; Ritter et al., 2021).

Non-Gaussian posterior approximations have also been explored. Hierarchical conjugate priors (Kessler et al., 2021) and sparsity-inducing priors (Bai et al., 2020; Ghosh et al., 2019) can be used to induce desirable structural properties of the weights, and in such cases posterior approximations typically share similar conjugacy or sparsity structure assumptions. Another line of work considers DNNs as flexible transformations to obtain expressive posterior approximations (Li and Turner, 2018; Louizos and Welling, 2017; Mescheder et al., 2017; Tran et al., 2017). In detail, with a learnable neural network $F_\phi(\boldsymbol{\omega})$ parameterized by ϕ and operated on a potentially lower-dimensional variable $\boldsymbol{\omega}$, the approximate posterior $q_\phi := (F_\phi)_\#q_0$ is defined as the push-forward distribution of a based distribution $q_0(\boldsymbol{\omega})$ that is often a standard Gaussian. When $F_\phi(\boldsymbol{\omega})$ represents an invertible function (e.g., normalizing flows (Louizos and Welling, 2017; Papamakarios et al., 2021; Rezende and Mohamed, 2015)), the resulting approximate posterior density can be obtained via change-of-variable rule: $q_\phi(\boldsymbol{\theta}) = q_0(F_\phi^{-1}(\boldsymbol{\theta}))\left|\frac{dF_\phi^{-1}(\boldsymbol{\theta})}{d\boldsymbol{\theta}}\right|$, which can then be plugged into the ELBO (Eq. (2.14)), and techniques such as Monte Carlo estimate and the reparameterization trick are applicable. On the other hand, for non-invertible $F_\phi(\boldsymbol{\omega})$ transforms, the resulting approximate posterior $q_\phi(\boldsymbol{\theta})$, while still supporting posterior predictive inference via Monte Carlo estimation and the reparameterization trick, no longer has a tractable density. Therefore further approximations to the ELBO objective and its gradients have been proposed (Li and Turner, 2018; Mescheder et al., 2017). When $\boldsymbol{\omega}$ is lower-dimensional, these neural network transformed posterior may have mismatched support as compared with the prior. A solution to this issue is to add in e.g., Gaussian noise after neural network transform, resulting in an approximate posterior $q_\phi(\boldsymbol{\theta}) = \int \mathcal{N}(\boldsymbol{\theta}|F_\phi(\boldsymbol{\omega}), \sigma^2\mathbf{I})q_0(\boldsymbol{\omega})d\boldsymbol{\omega}$ which is then fitted using a lower-bound approximation to the ELBO (Salimans et al., 2015; Yin and Zhou, 2018).

3 Deep Generative Models

Deep generative models (DGMs) are a class of deep learning models that aim to learn a $\boldsymbol{\theta}$ -parameterized distribution $p_\theta(\mathbf{x})$ from the training set $\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^N$ as an approximation for the true data distribution $p(\mathbf{x})$. After learning, these models can generate new data samples

from $p_{\theta}(\mathbf{x})$. DGMs have gained significant attention due to their ability to generate diverse and high-quality data, such as images (Kawar et al., 2023; Lugmayr et al., 2022; Saharia et al., 2022), text (Gong et al., 2022; Gu et al., 2022; Yang et al., 2023), audio (Chung et al., 2015; Huang et al., 2022; Kong et al., 2020) and molecules (Collins and Bettens, 2015; Li et al., 2014; Xu et al., 2021).

Roughly speaking, DGMs can be categorized into two main types. The first type directly parameterizes the data distribution $p_{\theta}(\mathbf{x})$. Examples include energy-based models (LeCun et al., 2006) and score-based models (Song et al., 2021). The second type introduces latent variables \mathbf{z} to aid in the learning process and acquire latent representations. These models represent the data distribution as $p_{\theta}(\mathbf{x}) = \int p_{\theta}(\mathbf{x}|\mathbf{z})p(\mathbf{z})d\mathbf{z}$. Examples include variational autoencoders (Kingma and Welling, 2014).

3.1 Energy-based Models

Energy-based models (EBMs) (LeCun et al., 2006) are probabilistic models with minimum restrictions on the functional form. They define an *energy function* $E_{\theta}(\mathbf{x})$ in the form of unnormalized log-probabilities. Practically, the energy function can be directly parameterized by any deep neural networks, which makes it flexible towards a variety of model architectures. The density given by EBMs is defined as:

$$p_{\theta}(\mathbf{x}) := \frac{1}{Z_{\theta}} \exp(-E_{\theta}(\mathbf{x})), \quad (3.1)$$

where $Z_{\theta} = \int_{\mathbf{x}} \exp(-E_{\theta}(\mathbf{x}))d\mathbf{x}$ is the normalizing constant. This normalizing constant is a function of the model parameter θ , and thus cannot be ignored when learning θ .

The ultimate goal of EBMs is to approximate the real data distribution $p(\mathbf{x})$ with the θ -parameterized density $p_{\theta}(\mathbf{x})$. One standard way to achieve such an approximation is by the maximum likelihood estimation (MLE), in which the expected log-likelihood is maximized:

$$\theta^* := \max_{\theta} E_{\mathbf{x} \sim p(\mathbf{x})}[\log p_{\theta}(\mathbf{x})]. \quad (3.2)$$




Figure 4: The training procedure of energy-based models. Generated data is obtained by sampling from the neural network-parameterized distribution $\mathbf{X}_{\text{fake}} \sim \exp(-E_{\theta}(\mathbf{x}))$. The model parameter θ is updated by minimizing the contrastive divergence between the real and generated data.

The MLE approximation is equivalent to minimizing the KL-divergence $KL[p(\mathbf{x})||p_{\theta}(\mathbf{x})]$. Unfortunately, gradient-based optimization cannot be directly applied to the MLE objective, since the the gradients of the logarithm of $p_{\theta}(\mathbf{x})$ contains an intractable term $\nabla_{\theta} \log Z_{\theta}$:

$$\nabla_{\theta} \log p_{\theta}(\mathbf{x}) = -\nabla_{\theta} E_{\theta}(\mathbf{x}) - \nabla_{\theta} \log Z_{\theta}. \quad (3.3)$$

A popular solution is based on Monte Carlo estimation with a random sample drawn from the EBM distribution, resulting in an approximate gradient to Eq. (3.3):

$$\nabla_{\theta} \log Z_{\theta} = -E_{\mathbf{x} \sim p_{\theta}(\mathbf{x})}[\nabla_{\theta} E_{\theta}(\mathbf{x})] \approx -\nabla_{\theta} E_{\theta}(\tilde{\mathbf{x}}), \quad \tilde{\mathbf{x}} \sim p_{\theta}(\mathbf{x}), \quad (3.4)$$

$$\Rightarrow \nabla_{\theta} \log p_{\theta}(\mathbf{x}) \approx \nabla_{\theta} E_{\theta}(\tilde{\mathbf{x}}) - \nabla_{\theta} E_{\theta}(\mathbf{x}), \quad \tilde{\mathbf{x}} \sim p_{\theta}(\mathbf{x}). \quad (3.5)$$

The training procedure of EBMs is illustrated in Fig. 4. As we shall see, sampling from the EBM is typically done via MCMC, often with finite number of transition steps. When using the empirical data distribution as the initial distribution, the corresponding approach is named *contrastive divergence* (Hinton, 2002), which is motivated as approximately minimizing the gap between two KL-divergences.

3.1.1 Langevin Sampling

During both training and generation, we need to obtain samples $\tilde{\mathbf{x}}$ from the EBM distribution $p_{\theta}(\mathbf{x})$. However, $p_{\theta}(\mathbf{x})$ lacks a closed form due to the intractable normalizing constant.

Algorithm 3 Training Energy-based methods

Require: Dataset $\mathcal{D} = \{\mathbf{x}_n\}_{n=1}^N$, sampling step sizes $\{\alpha_k\}$, total optimization steps T , total sampling steps K .

- 1: **for** $t = 1, \dots, T$ **do**
- 2: \triangleright Generate a sample from the current EBM via Langevin Sampling
- 3: Initialize \mathbf{x}_0 (see discussions in the main text)
- 4: **for** $t = 1, \dots, K$ **do**
- 5: $\mathbf{x}_{k+1} \leftarrow \mathbf{x}_k - \alpha_k \nabla_{\mathbf{x}_k} E_{\theta}(\mathbf{x}_k) + \sqrt{2\alpha_k} \cdot \boldsymbol{\epsilon}_k, \quad \boldsymbol{\epsilon}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 6: **end for**
- 7: Set the negative sample to be $\mathbf{x}^- = \mathbf{x}_K$
- 8: \triangleright Maximize the log-likelihood of the EBM
- 9: Sample a datapoint from the dataset: $\mathbf{x}^+ \sim \mathcal{D}$
- 10: Obtain the gradient of the log-likelihood: $\nabla_{\mathbf{x}} \log p_{\theta}(\mathbf{x}) = \nabla_{\theta} E_{\theta}(\mathbf{x}^-) - \nabla_{\theta} E_{\theta}(\mathbf{x}^+)$
- 11: Update θ by MLE using gradient-based optimizers
- 12: **end for**

Therefore, MCMC methods, particularly gradient-based MCMC, are commonly applied to acquire samples from $p_{\theta}(\mathbf{x})$. For example, Langevin dynamics (Grenander and Miller, 1994; Parisi, 1981; Welling and Teh, 2011) utilizes the gradient of $\log p_{\theta}(\mathbf{x})$:

$$\nabla_{\mathbf{x}} \log p_{\theta}(\mathbf{x}) = -\nabla_{\mathbf{x}} E_{\theta}(\mathbf{x}) - \nabla_{\mathbf{x}} \log Z_{\theta} = -\nabla_{\mathbf{x}} E_{\theta}(\mathbf{x}) \quad (3.6)$$

The second equation holds because the normalizing constant Z_{θ} is independent of x . Consequently, the update rule for x using Langevin dynamics is:

$$\mathbf{x}_{k+1} \leftarrow \mathbf{x}_k - \alpha_k \nabla_{\mathbf{x}_k} E_{\theta}(\mathbf{x}_k) + \sqrt{2\alpha_k} \cdot \boldsymbol{\epsilon}_k, \quad (3.7)$$

where $\boldsymbol{\epsilon}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is a standard Gaussian noise and α_k is the stepsize. With an annealed step size schedule, \mathbf{x}_k is guaranteed to converge to the EBM distribution $p_{\theta}(\mathbf{x})$ as $k \rightarrow \infty$. During the training process, we can use the empirical data distribution as the initial distribution (Hinton, 2002), and run a short Markov chain to generate samples $\tilde{\mathbf{x}}$ for computing the gradient in Eq. (3.5). The initialization of the MCMC can also be either random noise (Nijkamp et al., 2019) or persistent states obtained from previous iterations (Tieleman, 2008). A summary of the full training procedure for EBMs with MCMC sampling assistance is provided in Algorithm 3.

3.1.2 Score Matching

In Eq. (3.7), generating new data only needs the gradient of the energy $\nabla_{\mathbf{x}_k} E_{\theta}(\mathbf{x}_k)$ rather than the energy itself $E_{\theta}(\mathbf{x}_k)$. Therefore, another way to learn EBMs is by directly parameterizing the gradient of $\log p_{\theta}(\mathbf{x})$, referred to as the *score function*: $s_{\theta}(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_{\theta}(\mathbf{x}) \approx \nabla_{\mathbf{x}} \log p(\mathbf{x})$. Then the sampling rule during generation becomes:

$$\mathbf{x}_{k+1} \leftarrow \mathbf{x}_k + \alpha_k s_{\theta}(\mathbf{x}_k) + \sqrt{2\alpha_k} \cdot \epsilon_k. \quad (3.8)$$

Generative models that learn the data distribution through the score function are called *score-based models* (Song and Ermon, 2019). One benefit of learning the score function is that it is free from the normalizing constant. From Eq. (3.6), we have:

$$s_{\theta}(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_{\theta}(\mathbf{x}) = -\nabla_{\mathbf{x}} E_{\theta}(\mathbf{x}). \quad (3.9)$$

The training of score-based models involves *score matching* (Hyvärinen, 2005), in which the Fisher divergence between the learned score and true score is minimized:

$$D_F(p(x)||p_{\theta}(x)) = \frac{1}{2} E_{\mathbf{x} \sim p(\mathbf{x})} [\|s_{\theta}(\mathbf{x}) - \nabla_{\mathbf{x}} \log p(\mathbf{x})\|^2] \propto \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} \left[\text{tr}(\nabla_{\mathbf{x}} s_{\theta}(\mathbf{x})) + \frac{1}{2} \|s_{\theta}(\mathbf{x})\|^2 \right]. \quad (3.10)$$

However, computing the trace of the Jacobian $\text{tr}(\nabla_{\mathbf{x}} s_{\theta}(\mathbf{x}))$ is prohibitively costly for deep neural networks. To address this issue, *denoising score matching* is proposed (Vincent, 2011), which introduces a small Gaussian perturbation to the original data $q(\tilde{\mathbf{x}}|\mathbf{x}) = N(\tilde{\mathbf{x}}|\mathbf{x}, \sigma^2 I)$, and returns an accurate estimate of the score function when $\sigma \rightarrow 0$. Consequently,

$$\begin{aligned} D_F(q(\tilde{\mathbf{x}}|\mathbf{x})||p_{\theta}(\tilde{\mathbf{x}})) &= \frac{1}{2} E_{p(\mathbf{x})} E_{q(\tilde{\mathbf{x}}|\mathbf{x})} [\|s_{\theta}(\tilde{\mathbf{x}}) - \nabla_{\tilde{\mathbf{x}}} \log q(\tilde{\mathbf{x}}|\mathbf{x})\|^2] \\ &= \frac{1}{2} E_{p(\mathbf{x})} E_{q(\tilde{\mathbf{x}}|\mathbf{x})} \left[\left\| s_{\theta}(\tilde{\mathbf{x}}) + \frac{\tilde{\mathbf{x}} - \mathbf{x}}{\sigma^2} \right\|^2 \right] \\ &= \frac{1}{2} E_{p(\mathbf{x})} E_{\epsilon \sim \mathcal{N}(0, I)} \left[\left\| s_{\theta}(\mathbf{x} + \sigma\epsilon) + \frac{\epsilon}{\sigma} \right\|^2 \right] \end{aligned} \quad (3.11)$$

In practice with $\sigma > 0$, the learned score approximates the score of the perturbed data distribution, not the original data distribution. Even if an accurate data score function

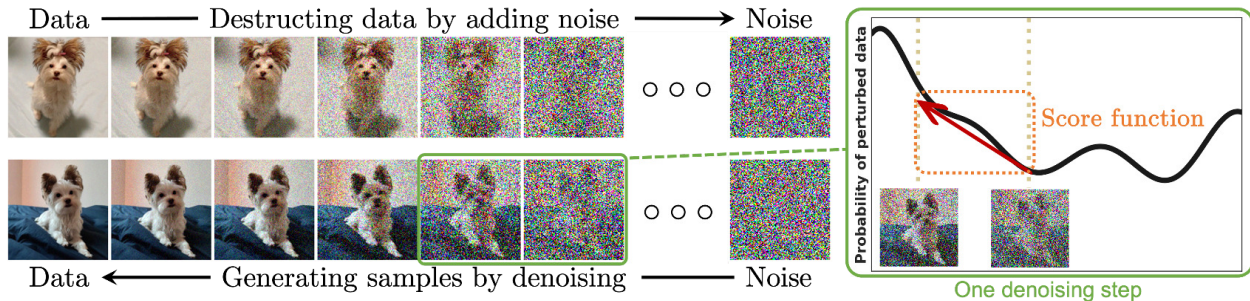


Figure 5: In diffusion models, the forward process progressively introduces noise, whereas the reverse process aims to denoise the perturbed data. The denoising step typically involves the estimation of the score function. Adapted from Yang et al. (2022). Used with kind permission of Ling Yang.

estimate is obtained with $\sigma \approx 0$, sampling from the data distribution with Langevin dynamics (Eq. (3.8)) remains challenging in e.g., fast mixing. To resolve these issues, Song and Ermon (2019) proposes learning the score function across different noise scales via *Noise Conditional Score Networks*, which parameterize the score function as $s_\theta(\mathbf{x}, \sigma)$ with the noise scale σ as an additional input. During generation, *Annealed Langevin dynamics* (Song and Ermon, 2019) is used which simulates Eq. (3.8) using $s_\theta(\mathbf{x}, \sigma)$ with different σ , starting from the most noise-perturbed distribution (with $\sigma \gg 0$) and then progressively reducing the noise scales until $\sigma \approx 0$. Consequently, the samples become close to the original data distribution.

3.2 Diffusion Models

Diffusion models are probabilistic generative methods that generate new data by reversing the process of injecting noise into data (Ho et al., 2020; Sohl-Dickstein et al., 2015). The intuition of diffusion models is illustrated in Fig. 5. Diffusion models are closely related to score-based models, essentially sharing the same training objective outlined in Eq. (3.11), differing only by a constant. Subsequently, a continuous-time generalization of diffusion models through Stochastic Differential Equations (SDE) has been developed (Song et al., 2021). The SDE perspective sheds light on the connection between diffusion models and MCMC methods, which we will discuss in detail below.

3.2.1 Stochastic Differential Equations

The diffusion process that injects noise to data is governed by the following stochastic differential equation:

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, t)dt + g(t)d\mathbf{w}, \quad (3.12)$$

where $\mathbf{f}(\mathbf{x}, t)$ and $g(t)$ are the diffusion and drift functions respectively, and \mathbf{w} is a standard Wiener process (Brownian motion). Anderson (1982) shows that any diffusion process can be reversed by solving the following reverse-time SDE:

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g(t)^2 \nabla_{\mathbf{x}} \log q_t(\mathbf{x})]dt + g(t)d\bar{\mathbf{w}}, \quad (3.13)$$

where $\bar{\mathbf{w}}$ is a standard Wiener process when time flows backwards, and dt denotes an infinitesimal negative time step.

The reverse-time SDE can be simulated using various general-purpose numerical SDE solvers. To address the approximation errors arising from these numerical methods, gradient-based MCMC methods, specifically Langevin dynamics, have been employed. The reason is that we have the score function $s_{\theta}(\mathbf{x}, t) \approx \nabla \log p_t(\mathbf{x})$ which can be used with gradient-based MCMC to sample from $p_t(\mathbf{x})$. This forms a hybrid sampler, denoted as *Predictor-Corrector* (PC) sampler (Song et al., 2021). Here, at each time step t , the numerical SDE solver first provides an estimate of the sample at the next time step, \mathbf{x}_{t+1} , serving as a “predictor”. Subsequently, the gradient-based MCMC method adjusts the marginal distribution of the estimated sample by running Monte Carlo sampling, using \mathbf{x}_{t+1} as the initial value, serving as a “corrector”.

Recent developments of diffusion model’s SDE design have been inspired by the continuous-time MCMC literature. Notably, continuous-time diffusion models have been found to align with overdamped Langevin dynamics with high friction coefficients (Dockhorn et al., 2022). To enhance the convergence towards equilibrium, a critically-damped Langevin diffusion has been introduced, leveraging concepts from Hamiltonian dynamics (Dockhorn et al., 2022).




Figure 6: Graphical model of a deep latent variable model.

3.3 Deep Latent Variable Models

Apart from generating realistic data and accurate density estimation where EBMs/SBMs and diffusion models are well suited, another important application of deep generative models is representation learning, possibly with dimensionality reduction. In such scenario we are interested in obtaining a (possibly lower dimensional) representation \mathbf{z} for a given datapoint \mathbf{x} sampled from the underlying data distribution. A starting example is Probabilistic PCA (Tipping and Bishop, 1999) which assumes the following model:

$$\mathbf{z} \sim p_{\theta}(\mathbf{z}) := \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad \mathbf{x} \sim p_{\theta}(\mathbf{x}|\mathbf{z}) := \mathcal{N}(\mathbf{W}\mathbf{z} + \mathbf{b}, \sigma^2\mathbf{I}). \quad (3.14)$$

The model parameters $\theta = \{\mathbf{W}, \mathbf{b}, \sigma\}$ is then fitted via maximum likelihood estimation (MLE) given a dataset $\mathcal{D} = \{\mathbf{x}_n\}_{n=1}^N$:

$$\theta^* = \arg \max_{\theta} \sum_{n=1}^N \log p_{\theta}(\mathbf{x}_n), \quad p_{\theta}(\mathbf{x}_n) = \int p_{\theta}(\mathbf{x}|\mathbf{z})p_{\theta}(\mathbf{z})d\mathbf{z}. \quad (3.15)$$

Since both the prior and the likelihood are Gaussians, the log marginal likelihood $\log p_{\theta}(\mathbf{x})$ is tractable, and there exist analytic solutions for the model parameters θ .

(Probabilistic) PCA returns meaningful representations when the underlying data generation process follows a linear transform of some latent features. Therefore it is less suited for typical deep learning applications with inputs as e.g., images, where useful features must be obtained in a non-linear way. Instead, Deep Latent Variable Models (DLVMs) (Kingma and Welling, 2014; Rezende et al., 2014) are typically used to extract non-linear features from such high-dimensional and complex datapoints, since they are generative models that assume a data-generating process as transforming latent variables through a neural network.

A typical DLVM assumes the observation \mathbf{x} is generated from the conditional distribution $p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})$ parameterized by a DNN with parameters $\boldsymbol{\theta}$. Compared with Probabilistic PCA, the conditional distribution in Eq. (3.14) is now replaced by $p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{x}; f_{\boldsymbol{\theta}}(\mathbf{z}), \sigma^2\mathbf{I})$ with a neural network $f_{\boldsymbol{\theta}}(\mathbf{z})$. It can then be viewed as a non-linear extension of Probabilistic PCA, since this formulation recovers Probabilistic PCA if $f_{\boldsymbol{\theta}}(\mathbf{z})$ is defined as a linear function.

The model parameters $\boldsymbol{\theta}$ in a DLVM is also estimated via maximum likelihood estimation, however, here the MLE objective (Eq. (3.15)) becomes intractable, due to the complex integral w.r.t. the latent variables \mathbf{z}_n when a non-linear function $f_{\boldsymbol{\theta}}(\mathbf{z})$ is in use. In practice, we again resort to approximation, where popular approaches are often VI-based (Kingma and Welling, 2014; Rezende et al., 2014). Different from VI in BNNs, amortized inference (Gershman and Goodman, 2014) plays a key role in VI for DLVMs to further reduce computational costs. As we shall see, instead of constructing individual approximations to the posterior $p_{\boldsymbol{\theta}}(\mathbf{z}_n|\mathbf{x}_n)$ for each \mathbf{x}_n , another DNN is employed to directly map \mathbf{x}_i to the variational parameters of its corresponding variational distribution.

3.3.1 Variational Autoencoders

Variational Autoencoders (VAEs) (Kingma and Welling, 2014; Rezende et al., 2014) are a class of DLVMs paired with *inference networks* (also called encoder networks) for amortized VI. Typically the Gaussian variational distribution family is used: $q_{\phi}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_{\phi}(\mathbf{x}), \boldsymbol{\sigma}_{\phi}^2(\mathbf{x}))$, where the posterior mean and variance for an observation \mathbf{x} are obtained by an inference network $g_{\phi}(\mathbf{x}) := [\boldsymbol{\mu}_{\phi}(\mathbf{x}), \log \boldsymbol{\sigma}_{\phi}(\mathbf{x})]$ with variational parameters ϕ . The model parameters $\boldsymbol{\theta}$ and the variational parameters ϕ can be jointly trained by maximizing the ELBO:

$$\boldsymbol{\theta}^*, \phi^* = \arg \max_{\boldsymbol{\theta}, \phi} \sum_{n=1}^N \mathcal{L}_{ELBO}(\mathbf{x}_n, \boldsymbol{\theta}, q_{\phi}), \quad (3.16)$$

$$\log p_{\boldsymbol{\theta}}(\mathbf{x}) \geq \mathcal{L}_{ELBO}(\mathbf{x}, \boldsymbol{\theta}, q_{\phi}) = E_{q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})] - KL[q_{\phi}(\mathbf{z}|\mathbf{x})||p_{\boldsymbol{\theta}}(\mathbf{z})].$$

Again the intractable expectation can be approximated with Monte Carlo estimation (often with only 1 sample), and the reparameterization trick (Kingma and Welling, 2014; Rezende

Algorithm 4 SGD Training for a variational auto-encoder with Gaussian q distribution

Require: Dataset $\mathcal{D} = \{\mathbf{x}_n\}_{n=1}^N$, initialized parameters $\boldsymbol{\theta}, \boldsymbol{\phi}$, batch-size B , SGD step sizes $\{\alpha_t\}$, SGD total steps T .

```
1: for  $t = 1, \dots, T$  do
2:   Sample a mini-batch of datapoints  $\Xi = \{\mathbf{x}\} \sim \mathcal{D}^B$  with  $|\Xi| = B$ 
3:   for  $\mathbf{x} \in \Xi$  do
4:     Encode the input data:  $g_\phi(\mathbf{x}) \leftarrow [\boldsymbol{\mu}_\phi(\mathbf{x}), \log \boldsymbol{\sigma}_\phi(\mathbf{x})]$ 
5:     Compute the KL term in the ELBO:  $\mathcal{L}_{ELBO}(\mathbf{x}, \boldsymbol{\theta}, q_\phi) \leftarrow -KL[q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})]$ 
6:     Sample the latent variable:  $\mathbf{z} \leftarrow \boldsymbol{\mu}_\phi(\mathbf{x}) + \boldsymbol{\sigma}_\phi(\mathbf{x}) \odot \boldsymbol{\epsilon}, \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
7:     Decode the latent variable to  $f_\theta(\mathbf{z})$ 
8:     Add-in data likelihood:  $\mathcal{L}_{ELBO}(\mathbf{x}, \boldsymbol{\theta}, q_\phi) \leftarrow \mathcal{L}_{ELBO}(\mathbf{x}, \boldsymbol{\theta}, q_\phi) + \log p(\mathbf{x}|\mathbf{z})$ 
9:   end for
10:  Compute SGD updates:  $(\boldsymbol{\theta}, \boldsymbol{\phi}) \leftarrow (\boldsymbol{\theta}, \boldsymbol{\phi}) + \alpha_t \nabla_{(\boldsymbol{\theta}, \boldsymbol{\phi})} \frac{1}{|\Xi|} \sum_{\mathbf{x} \in \Xi} \mathcal{L}_{ELBO}(\mathbf{x}, \boldsymbol{\theta}, q_\phi)$ 
11: end for
```

et al., 2014) is also applied to enable back-propagation:

$$\mathcal{L}_{ELBO}(\mathbf{x}, \boldsymbol{\theta}, q_\phi) \approx \log p_\theta(\mathbf{x}|\boldsymbol{\mu}_\phi(\mathbf{x}) + \boldsymbol{\sigma}_\phi(\mathbf{x}) \odot \boldsymbol{\epsilon}) - KL[q_\phi||p_\theta], \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (3.17)$$

This ELBO optimization can be viewed as training a stochastic autoencoder (a stochastic encoder $q_\phi(\mathbf{z}|\mathbf{x})$ plus a stochastic decoder $p_\theta(\mathbf{x}|\mathbf{z})$), hence the name “autoencoder”. The full algorithm for training VAEs is provided as Algorithm 4 with stochastic gradient descent.

The class of VAE models can be extended by considering VI with alternative divergences. For example, we can again control the coverage behavior of $q_\phi(\mathbf{z}|\mathbf{x})$ (from mode-seeking to mass covering) by using α -divergences with different α hyper-parameters (Hernández-Lobato et al., 2015; Li and Turner, 2016). The standard VAE can be viewed as a special case within this framework where the exclusive KL-divergence ($KL[q||p]$) is chosen within the α -divergence family. Another special case is the use of the inclusive KL-divergence ($KL[p||q]$), the resulting model becomes an Importance-Weighted Autoencoder (IWAE) (Burda et al., 2015), which was originally proposed to tighten the lower bound of $\log p_\theta(\mathbf{x})$ for optimization:

$$\log p_\theta(\mathbf{x}) \geq \mathcal{L}_{IWAE}(\mathbf{x}, \boldsymbol{\theta}, q_\phi) := E_{\prod_{m=1}^M q_\phi(\mathbf{z}^{(m)}|\mathbf{x})} \left[\log \frac{1}{M} \sum_{m=1}^M \left[\frac{p_\theta(\mathbf{x}|\mathbf{z}^{(m)})p_\theta(\mathbf{z}^{(m)})}{q_\phi(\mathbf{z}^{(m)}|\mathbf{x})} \right] \right]. \quad (3.18)$$

Again Monte Carlo estimation and the reparameterization trick apply to IWAE. The lower bound can be made tighter by using more MC samples, and as $M \rightarrow \infty$, \mathcal{L}_{IWAE} converges

to $\log p_{\theta}(\mathbf{x})$. While benefiting learning of model parameters θ , in such case the fitting of $q_{\phi}(\mathbf{z}|\mathbf{x})$ may be sub-optimal, as any reasonable importance sampling proposal can return the exact marginal likelihood with infinite amount of samples (Rainforth et al., 2018).

3.3.2 Combining VI and MCMC in DLVM

The ELBO (Eq. (3.16)) is a biased approximation to the MLE objective (Eq. (3.15)) unless the variational posterior $q_{\phi}(\mathbf{z}_i|\mathbf{x}_i)$ matches the exact posterior $p_{\theta}(\mathbf{z}_i|\mathbf{x}_i)$. This motivates the use of MCMC to generate more accurate approximate posterior samples whose distribution asymptotically converges to the exact posterior (Doucet et al., 2023; Hoffman, 2017). Starting from an initial distribution $q_0(\mathbf{z}|\mathbf{x})$, one runs T steps of MCMC transitions to obtain an improved approximate posterior $q_T(\mathbf{z}|\mathbf{x})$, and optimizes the model parameters θ by maximizing $\mathcal{L}_{ELBO}(\mathbf{x}, \theta, q_T)$ with the MCMC sample distribution $q_T(\mathbf{z}|\mathbf{x})$ as the approximate posterior. In DLVM context, gradient-based samplers such as Hamiltonian Monte-Carlo (HMC) (Duane et al., 1987; Neal et al., 2011) are preferred, since they adapt to the decoder more efficiently than e.g., Metropolis-Hastings (Metropolis et al., 1953) with random walk proposals (Metropolis et al., 1953; Sherlock et al., 2010). Still MCMC is computationally more expensive than VI and in practice may take prohibitively many transitions to converge, especially when the posterior is high-dimensional.

To speed-up MCMC, a popular approach is to draw its initial samples from a distribution q_0 that is already close to the exact posterior, with the hope of reducing the number of MCMC transition steps to obtain high-quality posterior samples (Geffner and Domke, 2021; Hoffman, 2017). A natural choice for such initial distributions is the variational posterior: $q_0^{(\phi)}(\mathbf{z}|\mathbf{x}) = q_{\phi}(\mathbf{z}|\mathbf{x})$, where one can fit the variational parameters ϕ by maximizing $\mathcal{L}_{ELBO}(\mathbf{x}, \theta, q_{\phi})$. Another approach involves distilling the initial distribution using the feedback from the final samples of a MCMC: one can improve $q_0^{(\phi)}(\mathbf{z}|\mathbf{x})$ by minimizing $D(q_0^{(\phi)}(\mathbf{z}|\mathbf{x}), q_T(\mathbf{z}|\mathbf{x}))$ for a chosen distance/divergence measure between distributions. Note that $q_T(\mathbf{z}|\mathbf{x})$ implicitly depends on ϕ but in practical implementations we apply a “stop gradient” operation to it and treat it as a constant w.r.t. ϕ . Also the density of $q_T(\mathbf{z}|\mathbf{x})$ is unknown as it is implicitly defined based on MCMC samples. These suggest the use of the inclusive KL-

divergence $KL[q_T(\mathbf{z}|\mathbf{x})||q_0^{(\phi)}(\mathbf{z}|\mathbf{x})] = E_{q_T}[-\log q_0^{(\phi)}(\mathbf{z}|\mathbf{x})] + C$ for such divergence measure, as minimizing it w.r.t. to ϕ doesn't require the density of $q_T(\mathbf{z}|\mathbf{x})$ (Li et al., 2017). Ruiz and Titsias (2019) further proposed a new divergence tailored to this task.

In addition to improving the initial distributions, variational inference can also be applied to tuning hyper-parameters of MCMC samplers (Campbell et al., 2021; Caterini et al., 2018; Gong et al., 2019; Salimans et al., 2015), such as the step size and momentum variance in HMC, which are known to be critical for fast convergence of the samplers (Neal et al., 2011). A line of research, as illustrated in Figure 7, treats the MCMC hyper-parameters ψ as additional variational parameters, meaning the improved approximate posterior $q_T^{(\psi)}(\mathbf{z}_T|\mathbf{x})$ can be further optimized w.r.t. ψ based on certain divergences between the exact posterior and the implicit distribution. Since $q_T^{(\psi)}(\mathbf{z}_T|\mathbf{x})$ is now implicit, its density is intractable, therefore the ELBO $\mathcal{L}_{ELBO}(\mathbf{x}, \boldsymbol{\theta}, q_T^{(\psi)})$ cannot be computed. In particular, the ELBO $\mathcal{L}_{ELBO}(\mathbf{x}, \boldsymbol{\theta}, q_T^{(\psi)})$ in Eq. (3.16) can also be reformulated as follows:

$$\mathcal{L}_{ELBO}(\mathbf{x}, \boldsymbol{\theta}, q_T^{(\psi)}) = E_{q_T^{(\psi)}(\mathbf{z}_T|\mathbf{x})}[\log p_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}_T)] + E_{q_T^{(\psi)}(\mathbf{z}|\mathbf{x})}[-\log q_T^{(\psi)}(\mathbf{z}_T|\mathbf{x})]. \quad (3.19)$$

Maximizing the differential entropy $H[q_T^{(\psi)}] = E_{q_T^{(\psi)}(\mathbf{z}|\mathbf{x})}[-\log q_T^{(\psi)}(\mathbf{z}_T|\mathbf{x})]$ prevents $q_T^{(\psi)}$ from collapsing to a point mass located at a mode of the true posterior $p_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x})$. To circumvent the intractability of $q_T^{(\psi)}(\mathbf{z}|\mathbf{x})$ density, Salimans et al. (2015) proposed a variational inference approach in the expanded space of all intermediate states along the Markov chain (i.e. $\mathbf{z}_{0:T}$), and derived a lower bound of \mathcal{L}_{ELBO} as the optimization objective without computing $H[q_T^{(\psi)}]$. Noticing that gradient-based optimization is in use, Li and Turner (2018) and Gong et al. (2019) proposed to directly approximate the gradient of $H[q_T^{(\psi)}]$ with respect to ψ in the optimization procedure. Campbell et al. (2021) proposed to simply drop $H[q_T^{(\psi)}]$ from the ELBO (Eq. (3.19)) and prevent $q_T^{(\psi)}$ from collapsing by using a mass-covering initial distribution $q_0^{(\phi)}$. Alternatively, one can also bypass the intractability of the density of $q_T^{(\psi)}$ by using divergences which only require samples from $q_T^{(\psi)}$ as objective, for example Kernelized Stein Discrepancy (Chwialkowski et al., 2016; Liu et al., 2016).

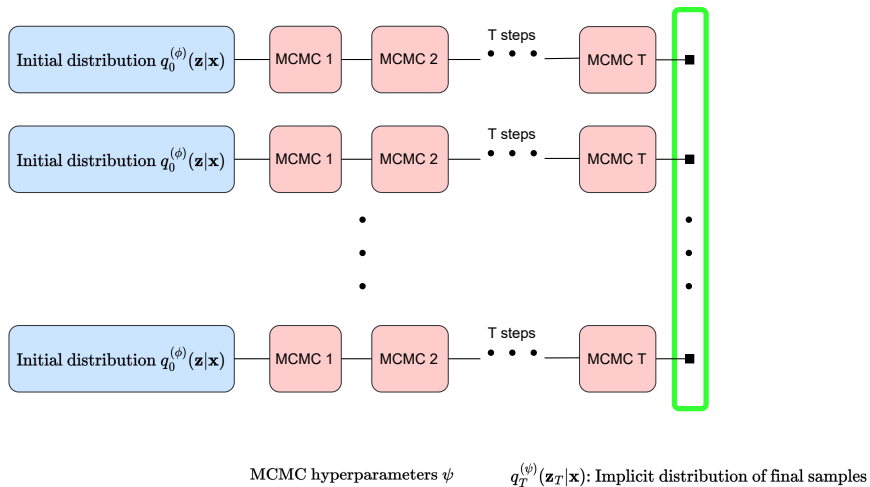


Figure 7: Implicit variational distribution $q_T^{(\psi)}(\mathbf{z}_T|\mathbf{x})$ defined via a T-step MCMC. The hyper-parameters of MCMC ψ are treated as learnable variational parameters.

4 Discussion

We have discussed scalable MCMC and approximate inference techniques for Bayesian computation in deep learning, with applications to quantifying uncertainty in DNNs and training deep generative models. Advances in fast posterior inference have contributed significantly to the rapid development of Bayesian neural networks as well as the profound success of deep generative models. Bayesian computation has been, and will continue to be, instrumental for embedding probabilistic principles into deep learning models.

The emphases on the need of posterior inference exhibit differences in the discussed two categories of modeling applications. For BNNs, posterior inference is the main computation tool for predictions and uncertainty estimates, although approximate inference for BNNs may require fitting parametric approximate posteriors, thereby introducing the concept of “BNN training” similar to training deterministic DNNs (Blundell et al., 2015; Gal, 2016; Wilson and Izmailov, 2020). For deep generative models, posterior inference serves as part of the training routine, and it may not be needed for data generation tasks in test time. Interestingly, when considering DLVMs for unsupervised or semi-supervised learning, posterior inference becomes handy again in obtaining useful representations of the observed data, which will then be used for further downstream tasks in test time (Higgins et al., 2016; Kingma et al.,

2014; Tschannen et al., 2018). An interesting recipe in this case is to use fast inference (e.g., mean-field VI) in DLVM training, and high-fidelity approaches (e.g., MCMC with rejection steps) in test time to obtain more accurate and robust inference results (Kuzina et al., 2022).

We conclude this chapter by discussing a few important research challenges to be addressed for more accurate and scalable Bayesian computation in deep learning context.

- For many popular DNNs the network weights are non-identifiable due to weight symmetries (Chen et al., 1993; Phuong and Lampert, 2020). This leads to many symmetric modes in the exact posterior of a BNN and increased difficulties of MCMC sampling (Izmailov et al., 2021; Papamarkou et al., 2022). But ultimately one cares more about the *posterior predictive distribution* (Eq. 2.5) of the neural network outputs, where averaging over symmetric modes in the exact posterior provides no additional advantage in this regard. So future research efforts should focus on accurate computation of the predictive posterior, which may provide exciting opportunities for fast and memory-efficient weight-space posterior approximations (Ma et al., 2019; Ritter et al., 2021; Sun et al., 2019). Optimization aspects are also under-explored for both SG-MCMC inference and parametric approximate posterior fitting of BNN posteriors. Apart from incorporating popular adaptive stochastic gradient methods (Chen et al., 2016; Li et al., 2016), the study of neural network loss landscape (Li et al., 2018; Maddox et al., 2019) may also inspire novel algorithms and analyses of posterior inference methods.
- For training deep generative models with posterior inference as a sub-routine, the bias in the training procedure (as compared with e.g., MLE), resulted from errors in inference, is much less well understood. For example, contrastive divergence (Hinton, 2002) as an approximation to MLE for energy-based models has been shown to perform adversarial optimization (Yair and Michaeli, 2021), but it is still unclear theoretically about the impact of MCMC convergence to the quality of learned energy-based model. For variational auto-encoders, although research has discussed the impact of amortization gap (Cremer et al., 2018; Marino et al., 2018) in variational inference and the aforementioned research in combining MCMC and VI, the combined impact of amortization and restricted variational family on latent variable model training is still an open question.

A note on author contributions

Wenlong Chen and Yingzhen Li wrote the sections regarding variational inference training for both Bayesian neural networks and deep latent variable models. Bolian Li and Ruqi Zhang wrote the sections on MCMC methods for Bayesian neural networks, energy-based models and diffusion models. Ruqi Zhang and Yingzhen Li edited the chapter in its final form for consistent presentations.

References

- Agostinelli, A., Denk, T. I., Borsos, Z., Engel, J., Verzett, M., Caillon, A., Huang, Q., Jansen, A., Roberts, A., Tagliasacchi, M., et al. (2023). Musiclm: Generating music from text. *arXiv preprint arXiv:2301.11325*.
- Amari, S.-i. (1985). Differential-geometrical methods in statistic. *Springer, New York*.
- Anderson, B. D. (1982). Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326.
- Bai, J., Song, Q., and Cheng, G. (2020). Efficient variational inference for sparse deep learning with theoretical guarantee. *Advances in Neural Information Processing Systems*, 33:466–476.
- Baydin, A. G., Pearlmutter, B. A., Radul, A. A., and Siskind, J. M. (2018). Automatic differentiation in machine learning: a survey. *The Journal of Machine Learning Research*.
- Beal, M. J. (2003). *Variational algorithms for approximate Bayesian inference*. PhD thesis, University of London.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. (2015). Weight uncertainty in neural networks. In *The International Conference on Machine Learning*.
- Bottou, L. (1998). Online learning and stochastic approximations. *On-line learning in neural networks*, 17(9):25.

- Burda, Y., Roger, G., and Salakhutdinov, R. (2015). Importance weighted autoencoders. In *International Conference on Learning Representations*.
- Campbell, A., Chen, W., Hernández-Lobato, J. M., and Zhang, Y. (2021). A gradient based strategy for hamiltonian monte carlo hyperparameter optimization. In *International Conference on Machine Learning*.
- Caterini, A. L., Doucet, A., and Sejdinovic, D. (2018). Hamiltonian variational auto-encoder. *Advances in Neural Information Processing Systems*, 31.
- Chen, A. M., Lu, H.-m., and Hecht-Nielsen, R. (1993). On the geometry of feedforward neural network error surfaces. *Neural computation*, 5(6):910–927.
- Chen, C., Carlson, D., Gan, Z., Li, C., and Carin, L. (2016). Bridging the gap between stochastic gradient mcmc and stochastic optimization. In *Artificial Intelligence and Statistics*, pages 1051–1060. PMLR.
- Chen, C., Ding, N., and Carin, L. (2015). On the convergence of stochastic gradient mcmc algorithms with high-order integrators. *Advances in neural information processing systems*, 28.
- Chen, T., Fox, E., and Guestrin, C. (2014). Stochastic gradient hamiltonian monte carlo. In *International conference on machine learning*, pages 1683–1691. PMLR.
- Chung, J., Kastner, K., Dinh, L., Goel, K., Courville, A. C., and Bengio, Y. (2015). A recurrent latent variable model for sequential data. *Advances in neural information processing systems*, 28.
- Chwialkowski, K., Strathmann, H., and Gretton, A. (2016). A kernel test of goodness of fit. In *International conference on machine learning*, pages 2606–2615. PMLR.
- Collins, M. A. and Bettens, R. P. (2015). Energy-based molecular fragmentation methods. *Chemical reviews*, 115(12):5607–5642.
- Cremer, C., Li, X., and Duvenaud, D. (2018). Inference suboptimality in variational autoencoders. In *International Conference on Machine Learning*, pages 1078–1086. PMLR.

- Csiszár, I. (1963). Eine informationstheoretische ungleichung und ihre anwendung auf den beweis der ergodizitat von markoffschen ketten. *Magyar. Tud. Akad. Mat. Kutató Int. Közl*, 8:85–108.
- Dehghani, M., Djolonga, J., Mustafa, B., Padlewski, P., Heek, J., Gilmer, J., Steiner, A. P., Caron, M., Geirhos, R., Alabdulmohsin, I., et al. (2023). Scaling vision transformers to 22 billion parameters. In *International Conference on Machine Learning*, pages 7480–7512. PMLR.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee.
- Deng, W., Feng, Q., Gao, L., Liang, F., and Lin, G. (2020a). Non-convex learning via replica exchange stochastic gradient mcmc. In *International Conference on Machine Learning*, pages 2474–2483. PMLR.
- Deng, W., Lin, G., and Liang, F. (2020b). A contour stochastic gradient langevin dynamics algorithm for simulations of multi-modal distributions. *Advances in neural information processing systems*, 33:15725–15736.
- Dobson, A. J. and Barnett, A. G. (2018). *An introduction to generalized linear models*. CRC press.
- Dockhorn, T., Vahdat, A., and Kreis, K. (2022). Score-based generative modeling with critically-damped langevin diffusion. *International Conference on Learning Representations*.
- Doucet, A., Moulines, E., and Thin, A. (2023). Differentiable samplers for deep latent variable models. *Philosophical Transactions of the Royal Society A*, 381(2247):20220147.
- Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D. (1987). Hybrid monte carlo. *Physics letters B*, 195(2):216–222.
- Gal, Y. (2016). Uncertainty in deep learning. *PhD dissertation, University of Cambridge*.
- Geffner, T. and Domke, J. (2021). Mcmc variational inference via uncorrected hamiltonian annealing. *Advances in Neural Information Processing Systems*, 34:639–651.

- Gelfand, A. E. (2000). Gibbs sampling. *Journal of the American statistical Association*, 95(452):1300–1304.
- Gershman, S. J. and Goodman, N. D. (2014). Amortized inference in probabilistic reasoning. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 36.
- Ghosh, S., Yao, J., and Doshi-Velez, F. (2019). Model selection in bayesian neural networks via horseshoe priors. *Journal of Machine Learning Research*, 20(182):1–46.
- Gong, S., Li, M., Feng, J., Wu, Z., and Kong, L. (2022). Diffuseq: Sequence to sequence text generation with diffusion models. In *The Eleventh International Conference on Learning Representations*.
- Gong, W., Li, Y., and Hernández-Lobato, J. M. (2019). Meta-learning for stochastic gradient mcmc. In *International Conference on Learning Representations*.
- Graves, A. (2011). Practical variational inference for neural networks. *Advances in neural information processing systems*, 24.
- Grenander, U. and Miller, M. I. (1994). Representations of knowledge in complex systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 56(4):549–581.
- Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., and Guo, B. (2022). Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10696–10706.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 770–778.
- Hernández-Lobato, J. M., Li, Y., Rowland, M., Hernández-Lobato, D., Bui, T., and Turner, R. E. (2015). Black-box α -divergence minimization. *The International Conference on Machine Learning*.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. (2016). Beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International conference on learning representations*.

- Hinton, G. E. (2002). Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800.
- Hinton, G. E. and Van Camp, D. (1993). Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pages 5–13. ACM.
- Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851.
- Hoffman, M. D. (2017). Learning deep latent Gaussian models with Markov Chain Monte Carlo. *Proceedings of the International Conference on Machine Learning (ICML)*.
- Houlsby, N., Huszár, F., Ghahramani, Z., and Lengyel, M. (2011). Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*.
- Huang, R., Lam, M., Wang, J., Su, D., Yu, D., Ren, Y., and Zhao, Z. (2022). Fastdiff: A fast conditional diffusion model for high-quality speech synthesis. In *IJCAI International Joint Conference on Artificial Intelligence*, pages 4157–4163. IJCAI: International Joint Conferences on Artificial Intelligence Organization.
- Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4).
- Izmailov, P., Vikram, S., Hoffman, M. D., and Wilson, A. G. G. (2021). What are bayesian neural network posteriors really like? In *International conference on machine learning*, pages 4629–4640. PMLR.
- Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge university press.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. (1999). An introduction to variational methods for graphical models. *Machine learning*, 37:183–233.
- Kawar, B., Zada, S., Lang, O., Tov, O., Chang, H., Dekel, T., Mosseri, I., and Irani, M. (2023). Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6007–6017.

- Kessler, S., Nguyen, V., Zohren, S., and Roberts, S. J. (2021). Hierarchical indian buffet neural networks for bayesian continual learning. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, volume 161 of *Proceedings of Machine Learning Research*, pages 749–759. PMLR.
- Kingma, D. P., Mohamed, S., Jimenez Rezende, D., and Welling, M. (2014). Semi-supervised learning with deep generative models. *Advances in neural information processing systems*, 27.
- Kingma, D. P. and Welling, M. (2014). Auto-encoding variational bayes. *stat*, 1050:1.
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., et al. (2023). Segment anything. *arXiv preprint arXiv:2304.02643*.
- Kong, Z., Ping, W., Huang, J., Zhao, K., and Catanzaro, B. (2020). Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*.
- Korattikara, A., Chen, Y., and Welling, M. (2014). Austerity in mcmc land: Cutting the metropolis-hastings budget. In *International conference on machine learning*, pages 181–189. PMLR.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- Kullback, S. and Leibler, R. A. (1951). On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86.
- Kuzina, A., Welling, M., and Tomczak, J. (2022). Alleviating adversarial attacks on variational autoencoders with mcmc. *Advances in Neural Information Processing Systems*, 35:8811–8823.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature*, 521(7553):436–444.
- LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., and Huang, F. (2006). A tutorial on energy-based learning. *Predicting structured data*, 1(0).

- Lemhadri, I., Ruan, F., Abraham, L., and Tibshirani, R. (2021). Lassonet: A neural network with feature sparsity. *Journal of Machine Learning Research*, 22(127):1–29.
- Li, C., Chen, C., Carlson, D., and Carin, L. (2016). Preconditioned stochastic gradient langevin dynamics for deep neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30.
- Li, H., Xu, Z., Taylor, G., Studer, C., and Goldstein, T. (2018). Visualizing the loss landscape of neural nets. *Advances in neural information processing systems*, 31.
- Li, S., Li, W., and Ma, J. (2014). Generalized energy-based fragmentation approach and its applications to macromolecules and molecular aggregates. *Accounts of chemical research*, 47(9):2712–2720.
- Li, Y. (2018). Approximate inference: New visions. *PhD dissertation, University of Cambridge*.
- Li, Y. and Turner, R. E. (2016). Rényi divergence variational inference. In *Advances in Neural Information Processing Systems*.
- Li, Y. and Turner, R. E. (2018). Gradient estimators for implicit models. In *International Conference on Learning Representations*.
- Li, Y., Turner, R. E., and Liu, Q. (2017). Approximate inference with amortised mcmc. *arXiv:1702.08343*.
- Liang, J., Wu, C., Hu, X., Gan, Z., Wang, J., Wang, L., Liu, Z., Fang, Y., and Duan, N. (2022). NUWA-infinity: Autoregressive over autoregressive generation for infinite visual synthesis. *Advances in Neural Information Processing Systems*, 35:15420–15432.
- Liu, Q., Lee, J. D., and Jordan, M. I. (2016). A kernelized stein discrepancy for goodness-of-fit tests. In *International Conference on Machine Learning (ICML)*.
- Louizos, C. and Welling, M. (2017). Multiplicative normalizing flows for variational bayesian neural networks. In *International Conference on Machine Learning*, pages 2218–2227. PMLR.

- Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., and Van Gool, L. (2022). Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471.
- Ma, C., Li, Y., and Hernández-Lobato, J. M. (2019). Variational implicit processes. In *International Conference on Machine Learning*, pages 4222–4233. PMLR.
- Ma, Y.-A., Chen, T., and Fox, E. (2015). A complete recipe for stochastic gradient mcmc. *Advances in neural information processing systems*, 28.
- MacKay, D. J. (1992). A practical Bayesian framework for backpropagation networks. *Neural Computation*, 4(3):448–472.
- Maclaurin, D. and Adams, R. P. (2014). Firefly monte carlo: Exact mcmc with subsets of data. *The Conference on Uncertainty in Artificial Intelligence*.
- Maddox, W. J., Izmailov, P., Garipov, T., Vetrov, D. P., and Wilson, A. G. (2019). A simple baseline for bayesian uncertainty in deep learning. *Advances in neural information processing systems*, 32.
- Marino, J., Yue, Y., and Mandt, S. (2018). Iterative amortized inference. In *International Conference on Machine Learning*, pages 3403–3412. PMLR.
- Mescheder, L., Nowozin, S., and Geiger, A. (2017). Adversarial variational bayes: Unifying variational autoencoders and generative adversarial networks. In *International conference on machine learning*, pages 2391–2400. PMLR.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092.
- Nair, V. and Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. In *The International Conference on Machine Learning*.
- Neal, R. M. (1996). Bayesian learning for neural networks. *Lecture Notes in Statistics*.
- Neal, R. M. et al. (2011). Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2.

- Nelder, J. A. and Wedderburn, R. W. (1972). Generalized linear models. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 135(3):370–384.
- Nijkamp, E., Hill, M., Zhu, S.-C., and Wu, Y. N. (2019). Learning non-convergent non-persistent short-run mcmc toward energy-based model. *Advances in Neural Information Processing Systems*, 32.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. (2021). Normalizing flows for probabilistic modeling and inference. *The Journal of Machine Learning Research*, 22(1):2617–2680.
- Papamarkou, T., Hinkle, J., Young, M. T., and Womble, D. (2022). Challenges in markov chain monte carlo for bayesian neural networks. *Statistical Science*, 37(3):425–442.
- Parisi, G. (1981). Correlation functions and computer simulations. *Nuclear Physics B*, 180(3):378–384.
- Phuong, M. and Lampert, C. H. (2020). Functional vs. parametric equivalence of relu networks. In *International conference on learning representations*.
- Quiroz, M., Kohn, R., Villani, M., and Tran, M.-N. (2018). Speeding up mcmc by efficient data subsampling. *Journal of the American Statistical Association*.
- Rainforth, T., Kosiorrek, A., Le, T. A., Maddison, C., Igl, M., Wood, F., and Teh, Y. W. (2018). Tighter variational bounds are not necessarily better. In *International Conference on Machine Learning*, pages 4277–4285. PMLR.
- Rényi, A. (1961). On measures of entropy and information. *Fourth Berkeley symposium on mathematical statistics and probability*, 1.
- Rezende, D. and Mohamed, S. (2015). Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR.
- Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31st International Conference on Machine Learning*, pages 1278–1286.

- Ritter, H., Kukla, M., Zhang, C., and Li, Y. (2021). Sparse uncertainty representation in deep learning with inducing weights. *Advances in Neural Information Processing Systems*, 34:6515–6528.
- Roberts, G. O. and Tweedie, R. L. (1996). Exponential convergence of langevin distributions and their discrete approximations. *Bernoulli*, pages 341–363.
- Rolnick, D. and Kording, K. (2020). Reverse-engineering deep relu networks. In *International conference on machine learning*, pages 8178–8187. PMLR.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.
- Ruiz, F. and Titsias, M. (2019). A contrastive divergence for combining variational inference and MCMC. *Proceedings of the International Conference on Machine Learning (ICML)*.
- Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., and Norouzi, M. (2022). Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10.
- Salimans, T., Kingma, D., and Welling, M. (2015). Markov chain monte carlo and variational inference: Bridging the gap. In *International conference on machine learning*, pages 1218–1226. PMLR.
- Savage, L. J. (1954). *The Foundations of Statistics*. Wiley Publications in Statistics.
- Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al. (2022). Laion-5B: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294.
- Sherlock, C., Fearnhead, P., and Roberts, G. O. (2010). The random walk metropolis: Linking theory and practice through a case study. *Statistical Science*, 25(2):172–190.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. (2015). Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR.

- Song, Y. and Ermon, S. (2019). Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. (2021). Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*.
- Sun, S., Zhang, G., Shi, J., and Grosse, R. (2019). Functional variational bayesian neural networks. In *International Conference on Learning Representations*.
- Tieleman, T. (2008). Training restricted boltzmann machines using approximations to the likelihood gradient. In *Proceedings of the 25th international conference on Machine learning*, pages 1064–1071.
- Tipping, M. E. and Bishop, C. M. (1999). Probabilistic principal component analysis. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 61(3):611–622.
- Tomczak, M., Swaroop, S., and Turner, R. (2020). Efficient low rank gaussian variational inference for neural networks. *Advances in Neural Information Processing Systems*, 33:4610–4622.
- Tran, D., Ranganath, R., and Blei, D. (2017). Hierarchical implicit models and likelihood-free variational inference. In *Advances in Neural Information Processing Systems*, pages 5529–5539.
- Tschannen, M., Bachem, O., and Lucic, M. (2018). Recent advances in autoencoder-based representation learning. *arXiv preprint arXiv:1812.05069*.
- Turner, R. E. and Sahani, M. (2011). Two problems with variational expectation maximisation for time-series models. In Barber, D., Cemgil, T., and Chiappa, S., editors, *Bayesian Time series models*, chapter 5, pages 109–130. Cambridge University Press.
- Van Erven, T. and Harremoës, P. (2014). Rényi divergence and Kullback-Leibler divergence. *Information Theory, IEEE Transactions on*, 60(7):3797–3820.
- Vincent, P. (2011). A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674.

- Vollmer, S. J., Zygalkis, K. C., and Teh, Y. W. (2016). Exploration of the (non-) asymptotic bias and variance of stochastic gradient langevin dynamics. *The Journal of Machine Learning Research*, 17(1):5504–5548.
- Wang, M. C. and Uhlenbeck, G. E. (1945). On the theory of the brownian motion ii. *Reviews of modern physics*, 17(2-3):323.
- Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688. Citeseer.
- Wilson, A. G. and Izmailov, P. (2020). Bayesian deep learning and a probabilistic perspective of generalization. *Advances in neural information processing systems*, 33:4697–4708.
- Xu, M., Yu, L., Song, Y., Shi, C., Ermon, S., and Tang, J. (2021). Geodiff: A geometric diffusion model for molecular conformation generation. In *International Conference on Learning Representations*.
- Yair, O. and Michaeli, T. (2021). Contrastive divergence learning is a time reversal adversarial game. In *International Conference on Learning Representations*.
- Yang, D., Yu, J., Wang, H., Wang, W., Weng, C., Zou, Y., and Yu, D. (2023). Diffsound: Discrete diffusion model for text-to-sound generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Yang, L., Zhang, Z., Song, Y., Hong, S., Xu, R., Zhao, Y., Zhang, W., Cui, B., and Yang, M.-H. (2022). Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*.
- Yin, M. and Zhou, M. (2018). Semi-implicit variational inference. In *International Conference on Machine Learning*, pages 5660–5669. PMLR.
- Zhang, C., Bütepage, J., Kjellström, H., and Mandt, S. (2018). Advances in variational inference. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):2008–2026.
- Zhang, R., Cooper, A. F., and De Sa, C. M. (2020a). Asymptotically optimal exact minibatch metropolis-hastings. *Advances in Neural Information Processing Systems*, 33:19500–19510.

- Zhang, R., Li, C., Zhang, J., Chen, C., and Wilson, A. G. (2020b). Cyclical stochastic gradient mcmc for bayesian deep learning. In *International Conference on Learning Representations*.
- Zhang, Y., Han, W., Qin, J., Wang, Y., Bapna, A., Chen, Z., Chen, N., Li, B., Axelrod, V., Wang, G., et al. (2023). Google USM: Scaling automatic speech recognition beyond 100 languages. *arXiv preprint arXiv:2303.01037*.