GENERATIVE MODELING WITH BAYESIAN SAMPLE INFERENCE

Anonymous authorsPaper under double-blind review

ABSTRACT

We derive a novel generative model from iterative Gaussian posterior inference. By treating the generated sample as an unknown variable, we can formulate the sampling process in the language of Bayesian probability. Our model uses a sequence of prediction and posterior update steps to iteratively narrow down the unknown sample starting from a broad initial belief. In addition to a rigorous theoretical analysis, we establish a connection between our model and diffusion models and show that it includes Bayesian Flow Networks (BFNs) as a special case. In our experiments, we demonstrate that our model improves sample quality on ImageNet32 over both BFNs and the closely related Variational Diffusion Models, while achieving equivalent log-likelihoods on ImageNet32 and CIFAR10.

1 Introduction

The field of deep learning has produced a multitude of generative models over the years (Harshvardhan et al., 2020). Variational autoencoders, for example, learn the data distribution by compressing data into a lower-dimensional representation (Kingma & Welling, 2013). Normalizing flows learn to map between a prior and the data distribution via invertible transformations, enabling exact likelihood computation (Rezende & Mohamed, 2015). Generative adversarial networks generate samples by pitting two models against each other such that one proposes artificial data samples while the other tries to distinguish real and generated (Goodfellow et al., 2014). Recently, diffusion models (DMs) have become a cornerstone of generative modeling (Sohl-Dickstein et al., 2015; Ho et al., 2020). They define a multi-step forward process that gradually adds noise to the data, turning it into pure noise. Then, a model is trained to reverse this process, enabling the generation of new data samples by starting from noise and iteratively denoising.

In this work, we take a Bayesian viewpoint of sample generation to propose a new generative model. Imagine that a sample \boldsymbol{x} from the data distribution $p(\boldsymbol{x})$ is fixed but unknown to us; however, we can receive noisy measurements $\boldsymbol{y}_i \sim \mathcal{N}(\boldsymbol{x}, \alpha_i^{-1})$ of it. Then, we can infer the unknown \boldsymbol{x} by combining the information in these measurements. To be more precise, we start with a broad belief $p(\boldsymbol{x}) = \mathcal{N}(\boldsymbol{x} \mid \boldsymbol{\mu}_0, \lambda_0^{-1})$ about \boldsymbol{x} in the form of a Normal distribution with low precision λ , i.e. high variance, that encompasses the entire data distribution. Then, we can take a first noisy measurement \boldsymbol{y}_1 and form a posterior belief $p(\boldsymbol{x} \mid \boldsymbol{y}_1)$ about the sample, which will be a little more precise and a little more correct. Iterating this process allows us to refine our estimate $p(\boldsymbol{x} \mid \boldsymbol{y}_1, \dots, \boldsymbol{y}_k)$ to any desired level of precision.

We transform this inference process into a generative model by introducing a prediction model f_{θ} that estimates x from our current Gaussian belief about it. Since the true x is unknown at generation time, we substitute it with an estimate $\hat{x} = f_{\theta}(\mu_i, \lambda_i)$ and sample $y_{i+1} \sim \mathcal{N}(\hat{x}, \alpha_{i+1}^{-1})$ instead. Maximizing an evidence lower bound (ELBO) for the likelihood that this simple process assigns to the training data, trains f_{θ} to reconstruct true x from uncertain belief states (μ_i, λ_i) about them. Consequently, the noisy measurements y_i of predicted samples \hat{x} become indistinguishable from those of real samples x, and our generative process converges toward producing new samples from the data distribution.

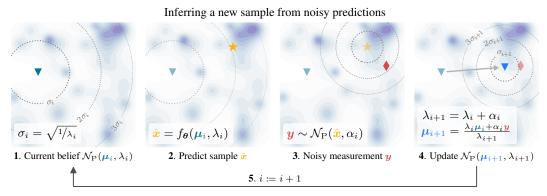


Figure 1. We view generation as the problem of inferring the identity of an unknown sample x from noisy observations. 1. To begin, our belief about x is so broad as to cover the complete data distribution. 2. We use a model f_{θ} to guess which x likely corresponds to the information we have collected so far. 3. Now, we pretend that \hat{x} is the true x and take a noisy measurement y. 4. We form the posterior belief about x to incorporate the information contained in y. 5. Repeat until we have identified a new sample with sufficient precision λ_i .

Our key **contributions** can be summarized as follows.

- We present a new generative model based on iterative posterior inference from noisy predictions.
- We derive an ELBO to enable effective likelihood optimization and show how we can reduce the variance of the training loss with importance sampling.
- Further, we compare our model in detail to Variational Diffusion Models (VDMs) (Kingma et al., 2023) and Bayesian Flow Networks (BFNs) (Graves et al., 2023).
- We show that the simple generative process described above includes BFN as a special case, providing a novel and simplified perspective on them, and analyze the relationship to DMs.
- Finally, we describe our model design and demonstrate empirically that our model surpasses both VDM and BFN in terms of sample quality on ImageNet32 while achieving equivalent log-likelihoods.

Notation We parametrize Normal distributions either with a variance σ^2 as $\mathcal{N}(\boldsymbol{\mu}, \sigma^2 \boldsymbol{I})$ or with a precision $\lambda = 1/\sigma^2$ as $\mathcal{N}_{P}(\boldsymbol{\mu}, \lambda \boldsymbol{I})$. Since all Normal distributions in this work are isotropic, we shorten these to $\mathcal{N}(\boldsymbol{\mu}, \sigma^2)$ and $\mathcal{N}_{P}(\boldsymbol{\mu}, \lambda)$. [n] is the set of integers $1, \ldots, n$ and \mathbb{R}_+ refers to the non-negative reals.

2 Sample Discovery through Iterative Measurement

Consider a sample $x \in \mathbb{R}^n$ that is unknown to us, but we can access noisy measurements $y_i \sim \mathcal{N}_P(x,\alpha_i)$ of it. Then we can infer x from the sequence of measurements y_i through Bayesian inference. We start with a broad initial belief $p(x) \sim \mathcal{N}_P(\mu_0,\lambda_0)$ and update it with information contained in y_i per the following lemma.

Lemma 2.1 (Posterior Update). Let $\mathbf{x}, \boldsymbol{\mu} \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}_+$ such that \mathbf{x} is latent and $p(\mathbf{x}) = \mathcal{N}_P(\mathbf{x} \mid \boldsymbol{\mu}, \lambda)$ is a prior on \mathbf{x} ; and $\mathbf{y} \sim \mathcal{N}_P(\mathbf{x}, \alpha)$ where $\alpha \in \mathbb{R}_+$. Then the posterior is $p(\mathbf{x} \mid \mathbf{y}) = \mathcal{N}_P(\mathbf{x} \mid \boldsymbol{\mu}', \lambda')$ with

$$\lambda' = \lambda + \alpha \quad and \quad \mu' = 1/\lambda' \left[\lambda \mu + \alpha y \right].$$
 (1)

Proof. See (Murphy, 2012, Section 4.4.1). □

We can now iterate over the noisy measurements and update our belief until $p(x \mid y_1, \dots, y_k) \sim \mathcal{N}_P(\mu_k, \lambda_k)$ identifies x with sufficient precision. Sufficiency depends on the application but could be defined, for example in the case of images, such that most of the probability mass for each dimension of an image x is contained within a single color intensity bin of width 1/256 for 8-bit color. Note that, at each step, all information contained in y_1, \dots, y_k is captured in the current μ_k .

109 110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129 130

131

132 133

134 135

136 137

138

139

140141

142

143 144

145

146

147 148

149

150 151

152 153 154

155 156

157 158 159

160

161

3 Sample Generation with Posterior Inference

We turn the procedure in Section 2 into a generative model, which we call *Bayesian Sample Inference* (BSI), as follows. We begin with an initial belief (μ_0, λ_0) about the sample x which we will generate in the end, with μ_0 sampled from a suitable prior distribution $p(\mu_0)$ and λ_0 fixed. Obviously, x is unknown a priori, so we cannot measure it, but we can estimate it from the information we have gathered so far.

Let $f_{\boldsymbol{\theta}}: \mathbb{R}^n \times \mathbb{R}_+ \to \mathbb{R}^n$ be a learned model with parameters θ that estimates which unknown sample x we have observed so far from our current belief (μ_i, λ_i) . We estimate x as $\hat{x}_{i-1} = f_{\theta}(\mu_{i-1}, \lambda_{i-1})$ and sample a noisy measurement $y_i \sim \mathcal{N}_{P}(\hat{x}_{i-1}, \alpha_i)$ of \hat{x}_{i-1} in place of x with precision α_i . Then, we can update our belief with y_i and Lemma 2.1 to the posterior (μ_i, λ_i) . Now, we alternate between these two steps, i.e. predicting and taking a noisy measurement followed by updating our current belief, until the posterior precision λ_i is sufficient. Finally, we return $\hat{x}^* = f_{\theta}(\mu_k, \lambda_k)$ as our generated sample. See Algorithm 1 for a formal description and Fig. 1 for a visual explanation.

Algorithm 1: Sampling with posterior infer**input** Initial precision λ_0 , precision schedule α_i for $i \in [k]$ output Sample \hat{x}^* 1: Initialize belief (μ_0, λ_0) with $\mu_0 \sim p(\mu_0)$ 2: **for** i = 1, 2, ..., k **do** $\hat{\boldsymbol{x}}_{i-1} = f_{\boldsymbol{\theta}}(\boldsymbol{\mu}_{i-1}, \lambda_{i-1})$ $oldsymbol{y}_i \sim \mathcal{N}_{\mathrm{P}}(\hat{oldsymbol{x}}_{i-1}, lpha_i)$ 4: Update belief 5: $p(\boldsymbol{x} \mid \boldsymbol{y}_1, \dots, \boldsymbol{y}_i) = \mathcal{N}_{P}(\boldsymbol{\mu}_i, \lambda_i)$: $\lambda_i = \lambda_{i-1} + \alpha_i$ $\boldsymbol{\mu}_i = \frac{1}{\lambda_i} [\lambda_{i-1} \boldsymbol{\mu}_{i-1} + \alpha_i \boldsymbol{y}_i]$ 6: 7: 8: end for 9: Return $\hat{\boldsymbol{x}}^* = f_{\boldsymbol{\theta}}(\boldsymbol{\mu}_k, \lambda_k)$

Since the posterior precision λ_i does not depend on the generated sample \hat{x}_i , we can choose the number of measurement rounds k and precision schedule α_i a priori such that λ_k will always be sufficiently large.

We have collected the proofs of all formal statements in this section in Section D.

3.1 EVIDENCE LOWER BOUND

By interpreting BSI as a hierarchical latent variable model, we derive an ELBO (Kingma & Welling, 2013), i.e. a lower bound on $\log p(x)$ assigned to a data point by our model. The ELBO will then serve as a natural training target for f_{θ} to ensure that true data samples have high likelihood under our model.

We form our hierarchy out of the sequence of belief means $\{\mu_i\}$, giving us

$$p(\boldsymbol{x}) = \underset{p(\boldsymbol{\mu}_0) \cdot p(\boldsymbol{\mu}_1 | \boldsymbol{\mu}_0) \cdots p(\boldsymbol{\mu}_k | \boldsymbol{\mu}_{k-1}))}{\mathbb{E}} [p(\boldsymbol{x} | \boldsymbol{\mu}_k)]. \tag{2}$$

The precisions $\{\lambda_i\}$ are not included as latent variables, because they do not depend on x. With this hierarchy, we can derive the following ELBO.

Theorem 3.1. Let $x \in \mathbb{R}^n$ and $\alpha_R, \alpha_i \in \mathbb{R}_+, i \in [k]$. Then the log-likelihood of x is lower-bounded as

$$\log p(\boldsymbol{x}) \ge -\mathcal{L}_{R} - \mathcal{L}_{M}^{k} \tag{3}$$

by a reconstruction term \mathcal{L}_{R} and a measurement term $\mathcal{L}_{\mathrm{M}}^k$,

$$\mathcal{L}_{R} = \underset{\mathbf{q}(\boldsymbol{\mu}_{k} \mid \boldsymbol{x}, \lambda_{k})}{\mathbb{E}} \left[-\log \mathcal{N}_{P}(\boldsymbol{x} \mid \hat{\boldsymbol{x}}_{k}, \alpha_{R}) \right] \quad and \quad \mathcal{L}_{M}^{k} = \frac{k}{2} \underset{\mathbf{q}(\boldsymbol{\mu}_{i} \mid \boldsymbol{x}, \lambda_{i})}{\mathbb{E}} \left[\alpha_{i+1} \| \boldsymbol{x} - \hat{\boldsymbol{x}}_{i} \|_{2}^{2} \right]$$
(4)

where

$$q(\boldsymbol{\mu}_i \mid \boldsymbol{x}, \lambda_i) = \mathbb{E}_{p(\boldsymbol{\mu}_0)} [p(\boldsymbol{\mu}_i \mid \boldsymbol{\mu}_0, \boldsymbol{x}, \lambda_i)], \quad \hat{\boldsymbol{x}}_i = f_{\boldsymbol{\theta}}(\boldsymbol{\mu}_i, \lambda_i) \quad and \quad \lambda_i = \lambda_0 + \sum_{j=1}^i \alpha_j. \quad (5)$$

The measurement term $\mathcal{L}_{\mathrm{M}}^{k}$ corresponds to the noisy measurement and update loop in Algorithm 1 and \mathcal{L}_{R} to the final computation of the sample \hat{x}^{*} . $q(\mu_{i} \mid x, \lambda_{i})$ is the distribution of our belief (μ_{i}, λ_{i}) about the unknown sample x after i steps if we would have observed the true x instead of

 $\hat{x}_1, \dots, \hat{x}_i$. $p(\mu_i \mid \mu_0, x, \lambda_i)$ is the marginal distribution of possible posterior beliefs (μ_i, λ_i) with posterior precision λ_i reachable from an initial belief (μ_0, λ_0) . Equivalently, $p(\mu_i \mid \mu_0, x, \lambda_i)$ is the distribution of beliefs (μ_i, λ_i) after updating our initial belief (μ_0, λ_0) with a single measurement of x with Lemma 2.1 – marginalized over all possible noisy measurements y at precision $\alpha = \lambda_i - \lambda_0$.

On closer examination, we see that \mathcal{L}_R , measuring how accurately we can reconstruct x at the end, only depends on the total precision λ_k that we accumulated in the first phase of the algorithm. However, \mathcal{L}_M^k depends both on the number of rounds k and the precision schedule α_i . We can derive an ELBO that is independent of k and α_i by considering the limit as $k \to \infty$ and refining the precision schedule $\{\alpha_i\}_{i=1}^k$ into smaller and smaller steps while keeping the total precision $\alpha_M = \sum_{i=1}^k \alpha_i$ constant.

Theorem 3.2. Let $\alpha_R, \alpha_M \in \mathbb{R}_+$. For any sequence of precision schedules $\alpha_{k,i}$ for $k \in \mathbb{N}, i \in [k]$ such that $\sum_{i=1}^k \alpha_{k,i} = \alpha_M$ and the sequence of functions $[k] \to \mathbb{R}_+ : i \mapsto \alpha_{k,i}$ converges uniformly to 0, we can take the limit of Theorem 3.1 as $k \to \infty$ to get

$$\mathcal{L}_{R} = \underset{q(\boldsymbol{\mu}_{\lambda_{M}}|\boldsymbol{x},\lambda_{M})}{\mathbb{E}} \left[-\log \mathcal{N}_{P}(\boldsymbol{x} \mid \hat{\boldsymbol{x}}_{\lambda_{M}}, \alpha_{R}) \right] \quad and \quad \mathcal{L}_{M}^{\infty} = \frac{\alpha_{M}}{2} \underset{q(\boldsymbol{\mu}_{\lambda}|\boldsymbol{x},\lambda)}{\mathbb{E}} \left[\|\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda}\|_{2}^{2} \right] \quad (6)$$

where
$$q(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{x}, \lambda) = \mathbb{E}_{p(\boldsymbol{\mu}_{0})} [p(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{\mu}_{0}, \boldsymbol{x}, \lambda)], \lambda_{M} = \lambda_{0} + \alpha_{M} \text{ and } \hat{\boldsymbol{x}}_{\lambda} = f_{\boldsymbol{\theta}}(\boldsymbol{\mu}_{\lambda}, \lambda).$$

As long as our model is more accurate in reconstructing x from more precise measurements, a reasonable assumption, Theorem 3.2 is a tighter bound on the log-likelihood than Theorem 3.1. To see this, we rewrite $\mathcal{L}_{\mathrm{M}}^{\infty}$ in terms of the expected squared error at belief precision λ

$$h(\lambda) = \mathbb{E}_{\mathbf{q}(\boldsymbol{\mu}_{\lambda}|\boldsymbol{x},\lambda)} \|\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda}\|_{2}^{2}$$
(7)

as

$$\mathcal{L}_{\mathbf{M}}^{\infty} = \frac{\alpha_{\mathbf{M}}}{2} \underset{\lambda \sim \mathcal{U}(\lambda_{0}, \lambda_{\mathbf{M}})}{\mathbb{E}} [h(\lambda)] \tag{8}$$

for which we have the following result.

Lemma 3.3. If h is strictly decreasing, $\mathcal{L}_{\mathrm{M}}^{\infty} < \mathcal{L}_{\mathrm{M}}^{k}$ for any k and any precision schedule $\{\alpha_{i}\}$.

3.2 PRIOR DISTRIBUTION

Let's consider possible priors of the form $p(\mu_0) = \mathcal{N}_P(\mathbf{0}, \gamma_0)$ for our initial belief. Then we have the following result for the encoding distribution $q(\mu_\lambda \mid x, \lambda)$ in Theorems 3.1 and 3.2.

Lemma 3.4. Let $\lambda_0, \gamma_0 \in \mathbb{R}_+$, $p(\boldsymbol{\mu}_0) = \mathcal{N}_P(\mathbf{0}, \gamma_0)$ and $\lambda \geq \lambda_0$. Then

$$q(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{x}, \lambda) = \mathcal{N}_{P}\left(\frac{\lambda - \lambda_{0}}{\lambda} \boldsymbol{x}, \frac{\lambda^{2}}{\lambda - \lambda_{0} + \lambda_{0}^{2}/\gamma_{0}}\right). \tag{9}$$

How should we choose γ_0 ? We start the sampling process with initial precision, i.e. confidence, λ_0 . If λ_0 was larger than γ_0 , we would be unreasonably confident in our initial belief, since we know that μ_0 has more uncertainty than λ_0 . From this, we deduce that the reasonable range for γ_0 is $[\lambda_0, \infty]$. At the same time, we want to avoid unwarranted assumptions by the prior, so we choose $\gamma_0 = \lambda_0$ for our model, which also gives us a particularly simple form for the encoding distribution.

Corollary 3.5. Let $\lambda_0 \in \mathbb{R}_+$, $p(\mu_0) \sim \mathcal{N}_P(\mathbf{0}, \lambda_0)$ and $\lambda \geq \lambda_0$. Then

$$q(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{x}, \lambda) = \mathcal{N}_{P}\left(\frac{\lambda - \lambda_{0}}{\lambda} \boldsymbol{x}, \lambda\right). \tag{10}$$

3.3 VARIANCE REDUCTION

The squared distance $\|\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda}\|_2^2$ in $\mathcal{L}_{\mathrm{M}}^{\infty}$ will necessarily vary significantly across the range of λ with large values for small λ where $q(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{x}, \lambda) \approx p(\boldsymbol{\mu}_0)$ and small values for large λ when $\boldsymbol{\mu}_{\lambda} \approx \boldsymbol{x}$. We can reduce the variance of Monte Carlo (MC) estimates of $\mathcal{L}_{\mathrm{M}}^{\infty}$ for ELBO evaluation or gradient computation in training with importance sampling with a suitable proposal distribution $p(\lambda)$.

Corollary 3.6. Let $p(\lambda)$ be a probability distribution with support $[\lambda_0, \lambda_M]$. Then we have

$$\mathcal{L}_{\mathbf{M}}^{\infty} = \frac{1}{2} \underset{\mathbf{q}(\boldsymbol{\mu}_{\lambda}|\boldsymbol{x},\lambda)}{\mathbb{E}} \left[\frac{1}{\mathbf{p}(\lambda)} \|\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda}\|_{2}^{2} \right]. \tag{11}$$

We can further rewrite $\mathcal{L}_{\mathrm{M}}^{\infty}$ as

216

217

218

219 220 221

222

223 224

225

226

227

228

229

230

231

232

233

234

235

236

237

238

239

240

241

242

243

244

245

246 247

248 249

250

251

253

254

255 256

257 258

259

260

261

262

263

264

265

266 267

268

$$\mathcal{L}_{\mathbf{M}}^{\infty} = \frac{1}{2} \underset{\lambda \sim \mathbf{p}(\lambda)}{\mathbb{E}} \left[\frac{h(\lambda)}{\mathbf{p}(\lambda)} \right]$$
 (12)

with h as defined in Eq. (7). To minimize the variance of MC estimates of $\mathcal{L}_{\mathrm{M}}^{\infty}$, we want to bring $h(\lambda)/p(\lambda)$ as close to a constant as possible. If it were actually constant, the variance of the MC estimate would be zero.

Let's begin by examining h more closely. If we approximate f_{θ} as $f_{\theta}(\mu, \lambda) = \mu$ and assume that x is normalized to zero mean and unit variance, we get the closed form

$$\mathbb{E}[h(\lambda)] \propto \frac{\lambda_0^2}{\lambda^2} + \frac{1}{\lambda}.$$
 (13)

While $f_{\theta}(\mu, \lambda) = \mu$ might seem a crude approximation at first, it is not too far off for large λ where the model just needs to predict a small correction to its input.

Eq. (13) suggests that we should choose $p(\lambda) \propto$ $\lambda_0^2/\lambda^2 + 1/\lambda$ to minimize the variance of MC estimates. While evaluating $p(\lambda)$ is simple enough, we would need to invert its cumulative distribution function (CDF) numerically to sample from it. Instead, we recognize that $1/\lambda$ dominates λ_0^2/λ^2 except for the smallest λ and choose $p(\lambda) \propto 1/\lambda$, i.e. a standard Log-Uniform(λ_0, λ_M) distribution.

3.4 Training & Sampling

We train our model with the ELBO from Theorem 3.2 by optimizing $2\mathcal{L}_{\mathrm{M}}^{\infty}/n$. We do not

Algorithm 2: Estimating the BSI training loss

input Data sample x**output** Monte Carlo estimate of $\mathcal{L}_{\mathrm{M}}^{\infty}$

1: Sample $t \sim \mathcal{U}(0,1), \varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

2:
$$\lambda = \exp\left((\log \lambda_{\mathrm{M}} - \log \lambda_{0}) \cdot t + \log(\lambda_{0})\right)$$

3:
$$\mu_{\lambda} = (\lambda - \lambda_0)/\lambda x + \sqrt{1/\lambda} \varepsilon$$

4: Return

$$(\log \lambda_{\mathrm{M}} - \log \lambda_{0}) \lambda \cdot \|\boldsymbol{x} - f_{\boldsymbol{\theta}}(\boldsymbol{\mu}_{\lambda}, \lambda)\|_{2}^{2}$$

Algorithm 3: Sampling with BSI

input Initial precision λ_0 ,

precision schedule α_i for $i \in [k]$

output Sample \hat{x}^*

1: Sample
$$\boldsymbol{\varepsilon}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I}), i = 0, \dots, k$$

2:
$$\mu_0 = \sqrt{1/\lambda_0} \, \varepsilon_0$$

2:
$$\mu_0 = \sqrt{1/\lambda_0} \, \varepsilon_0$$

3: **for** $i = 1, 2, ..., k$ **do**

4:
$$\hat{\boldsymbol{x}}_{i-1} = f_{\boldsymbol{\theta}}(\boldsymbol{\mu}_{i-1}, \lambda_{i-1})$$

5:
$$\lambda_i = \lambda_{i-1} + \alpha_i$$

6:
$$\mu_i = \lambda_{i-1}^{-1} + \alpha_i$$

$$\lambda_i^{-1} \left(\lambda_{i-1} \mu_{i-1} + \alpha_i (\hat{x}_{i-1} + \sqrt{1/\alpha_i} \varepsilon_i) \right)$$
7: **end for**

8: Return
$$\hat{\boldsymbol{x}}^* = f_{\boldsymbol{\theta}}(\boldsymbol{\mu}_k, \lambda_k)$$

optimize \mathcal{L}_R directly as its magnitude is negligible for sufficiently large α_M and it is structurally similar to $\mathcal{L}_{\mathrm{M}}^{\infty}$, i.e. both amount to a squared distance. Algorithm 2 shows the resulting algorithm with our belief prior $p(\mu_0)$ and proposal distribution $p(\lambda)$. Similarly, Algorithm 3 implements the abstract Algorithm 1 with our belief prior.

DISCUSSION

We are aware of two generative models that are closely related to BSI, BFN (Graves et al., 2023) and VDM (Kingma et al., 2023). BFNs are generative models motivated from an information theory perspective with a sender and a receiver communicating about the sample. We show in Section A.1 that BFNs are a special case of our framework in Section 3 if we translate them to the probabilistic perspective and choose $\gamma_0 = \infty$ instead of $\gamma_0 = \lambda_0$ for the initial belief prior, meaning that sampling always starts from the deterministic $\mu_0 = 0$. VDM are a type of DM that have shown excellent performance in likelihood-based modeling. They are similar to BSI insofar as they specify the distribution of latent variables directly rather than defining a Markovian noising process as classical DMs do.

All three models admit an ELBO of the form

$$-\log p(\boldsymbol{x}) \leq \mathcal{L}_{R} + \frac{\bar{\omega} - \omega}{2} \underset{q(\boldsymbol{\psi}_{\omega}|\boldsymbol{x},\omega)}{\mathbb{E}} \left[\|\boldsymbol{x} - \hat{\boldsymbol{x}}_{\omega}\|_{2}^{2} \right]$$
(14)

Table 1. Central structures of VDM, BFN and BSI. To improve comparability, we parametrize VDM in terms of the signal-to-noise ratio (SNR) ν . BFN and BSI are parametrized with the precision α as introduced in Section 3. $\varepsilon_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is sampling noise.

Model	ELBO Encoder $\mathbf{q}(oldsymbol{\psi} \mid oldsymbol{x}, \omega)$	Latent Prior	Update Step for Sampling
VDM	$q(\boldsymbol{z} \mid \boldsymbol{x}, \nu) = \mathcal{N}_{P}\Big(\sqrt{\frac{\nu}{1+\nu}} \boldsymbol{x}, 1+\nu\Big)$	$z_T \sim \mathcal{N}_{\mathrm{P}}(0, 1)$	$z_i = \frac{\sqrt{\nu_{i+1}(1+\nu_{i+1})}z_{i+1} + (\nu_i - \nu_{i+1}) \left(\hat{x}_i + \sqrt{\frac{1}{\nu_i - \nu_{i+1}}}\boldsymbol{\varepsilon}_i\right)}}{\sqrt{\nu_i(1+\nu_i)}}$
BFN	$\mathrm{q}(oldsymbol{\mu} \mid oldsymbol{x}, \lambda) = \mathcal{N}_{\mathrm{P}}({}^{(\lambda-\lambda_0)}\!/_{\!\lambda} oldsymbol{x}, {}^{\lambda^2}\!/_{\!(\lambda-\lambda_0)})$	$\mu_0 = 0$	$\lambda_{i-1}\mu_{i-1} + lpha_i \left(\hat{m{x}}_{i-1} + \sqrt{rac{1}{lpha_i}}m{arepsilon}_i ight)$
BSI	$\mathrm{q}(oldsymbol{\mu} \mid oldsymbol{x}, \lambda) = \mathcal{N}_{\mathrm{P}}({}^{(\lambda - \lambda_0)} / _{\!\! \lambda} oldsymbol{x}, \lambda)$	$\mu_0 \sim \mathcal{N}_{\mathrm{P}}(0, \lambda_0)$	$\mu_i = \frac{1}{\lambda_{i-1} + \alpha_i}$

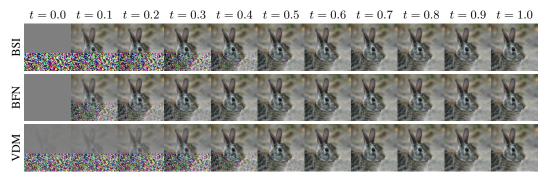


Figure 2. ELBO encoders q, i.e. training inputs, of BSI, BFN and VDM. t parametrizes the precision levels by the respective model's precision schedule with t=0 being pure noise, ideally, and t=1 almost equaling the data. Top half shows the mean of q and bottom half a sample. Mean ${\bf 0}$ is gray because all models rescale the data to [-1,1]. BFNs apply little noise overall and reach a deterministic state at t=0. For VDM, significant information about the sample is preserved in the structure of the mean at the highest noise level. In contrast, BSI converges to its latent prior distribution.

for a set of latent variables ψ at precision levels ω between $\underline{\omega}$ and $\bar{\omega}$. For BSI and BFN, the precision level ω is the belief precision λ between λ_0 and $\lambda_{\rm M}$ and $\psi_{\omega}=\mu_{\lambda}$. For VDM, the latent variables ψ are called z and they parametrize ω as the SNR ν between e^{-5} and $e^{13.3}$.

Despite this shared ELBO form, the models vary significantly. Table 1 lists the encoding distribution $q(\psi \mid x, \omega)$ for each model, their prior, from which they begin the sampling process, and the update step that the models iterate during sampling. First, we see that VDM starts sampling from a standard Normal vector and BFN from the deterministic 0. Only BSI allows sampling from an initial precision λ_0 less than 1, which has been shown to improve sample diversity in consistency models (Song & Dhariwal, 2024). Second, the update step shared between BSI and BFN is significantly simpler than the VDM update with respect to the precision parameter and does not require working in log-space for numerical stability as recommended for VDM (Kingma et al., 2023).

For the encoding distribution $q(\psi \mid x, \omega)$, which provides the training inputs when the models optimize their ELBO, we turn to Fig. 2. First, we note that BFN adds little noise overall due to their noise variance $(\lambda - \lambda_0)/\lambda^2$ going to 0 for both small and large λ . Next, we notice the encoding distribution $q(\psi \mid x, \omega)$ with the most noise at t=0. While it agrees exactly with the latent prior used for sampling for BSI and BFN, for VDM it becomes approximately $\mathcal{N}_P(0.08x, 1)$, which differs significantly from the standard Normal prior for sampling. In fact, the image motif is still clearly discernible in the distribution mean for VDM at its maximum noise level. The amount of signal remaining in the mean for BSI at high noise levels is counteracted by much higher noise variance, e.g. 15.85 at t=0.1 for BSI compared to 0.96 for VDM.

Diffusion Models If we currently hold the belief (μ', λ') , the distribution over beliefs $(\mu, \lambda' - \alpha)$ that are α less precise is

$$p(\boldsymbol{\mu} \mid \boldsymbol{\mu}', \boldsymbol{x}) = \mathcal{N}\left(\xi^{-1} \left[\frac{\lambda \lambda'}{\alpha} \boldsymbol{\mu}' - \lambda_0 \boldsymbol{x}\right], \xi\right)$$
(15)

for a certain precision ξ . This shows that BSI can be written as a DM with a non-Markovian forward or "noising" process. See Section A.2 for a detailed derivation of this connection. There we also

exploit that BFNs are a special case of BSI to derive the forward process for BFN and show that it is Markov, in contrast to the BSI process.

5 MODEL DESIGN

In this section, we introduce a design for the prediction model in BSI. We begin by deriving a preconditioning structure for f_{θ} , i.e. a type of model structure similar to noise prediction in DMs. Then, we describe how we bring λ into a suitable range as an input for deep learning. Finally, we give our choice of the hyperparameters λ_0 , $\alpha_{\rm M}$ and $\alpha_{\rm R}$ and report the model architectures we used as the backbone of f_{θ} .

5.1 Preconditioning

It has long been known in the context of DMs that training models to predict x directly from a noisy input can hinder learning and limit sample quality (Karras et al., 2022; Ho et al., 2020). For probabilistic modeling, it is especially important that the model prediction stays close to the true sample if the input is already at a low noise level to achieve high ELBOs. This can be seen, for example, in Corollary 3.6 where prediction errors for high-precision input beliefs with large λ have a higher weight. Instead of predicting x, DMs commonly either predict a variation of the noise in the model input (Ho et al., 2020; Song et al., 2021a) or an adaptive mixture of the noise and the true sample (Salimans & Ho, 2021). In the end, these approaches amount to adding a skip connection to the model with specific weights.

For BSI, we derive such a preconditioning structure with the adaptive-mixture approach from Karras et al. (2022). Let f'_{θ} be our neural network. Then we define the preconditioned f_{θ} as

$$f_{\theta}(\mu, \lambda) = c_{\text{skip}}\mu + c_{\text{out}}f'_{\theta}(c_{\text{in}}\mu, \lambda)$$
(16)

and find the parameters through the conditions proposed by Karras et al. (2022). $c_{\rm in}$ and $c_{\rm out}$ are chosen such that the input to f_{θ}' and its training target have unit variance. $c_{\rm skip}$ is then chosen to minimize $c_{\rm out}$, which minimizes the influence of prediction errors and ensures that f_{θ} retains most of the signal already contained in μ at large precisions λ .

From these conditions, we derive

$$c_{\rm skip} = (\lambda - \lambda_0)/\kappa, \quad c_{\rm out} = \sqrt{1/\kappa}, \quad c_{\rm in} = \sqrt{\lambda/\kappa}$$
 (17)

where $\kappa = 1 + (\lambda - \lambda_0)^2 / \lambda$ in Section C. λ is the precision of our current belief about x and the input to f_{θ} .

5.2 Precision Encoding

The magnitude of λ makes it impractical as a feature for neural networks. However, the CDF F of $p(\lambda)$ is a natural way to scale λ from $[\lambda_0, \lambda_M]$ to [0,1] as in DMs and flow matching (FM) (Lipman et al., 2023). In practice, we use $f_{\theta}(\mu, t)$ instead of $f_{\theta}(\mu, \lambda)$ where

$$t = F(\lambda) = \frac{\log \lambda - \log \lambda_0}{\log(\lambda_{\rm M}) - \log \lambda_0}.$$
 (18)

Compared to linear re-scaling, our method makes it easier for f_{θ} to distinguish belief precisions in the high-noise regime.

5.3 Hyperparameters

Apart from f_{θ} , BSI has three hyperparameters, λ_0 , $\alpha_{\rm M}$ and $\alpha_{\rm R}$. λ_0 should be small enough that the initial belief covers the whole data distribution. We have found experimentally that $\lambda_0=10^{-2}$ works well for images rescaled to [-1,1]. This agrees with the finding of Song & Dhariwal (2024) that large initial noise scales improve sample diversity in consistency models.

 $\alpha_{\rm M}$ should be large enough that a noisy measurement at precision $\alpha_{\rm M}$ identifies an ${\boldsymbol x}$, e.g. for images almost all probability mass of ${\mathcal N}_{\rm P}({\boldsymbol x},\alpha_{\rm M})$ should be contained within a single 8-bit color intensity bin. We choose $\alpha_{\rm M}=10^6$, which Graves et al. (2023) also picked for BFN. While ${\mathcal L}_{\rm M}^\infty$ dwarfs ${\mathcal L}_{\rm R}$, $\alpha_{\rm R}=2\alpha_{\rm M}$ gives a slight edge in likelihood, empirically, as also observed by Graves et al. (2023).

5.4 ARCHITECTURE

After the preconditioning and mapping λ to a $t \in [0,1]$, there are two more steps to turn the inputs μ and t of f'_{θ} into effective features for a neural network. Regarding t, we convert it into a 32-dimensional precision embedding with a sinusoidal position encoding (Vaswani et al., 2017).

The Fourier features proposed by Kingma et al. (2023) are an essential component to reach high likelihoods, because they help the model distinguish fine details at high likelihoods, i.e. for inputs that are already close to the data distribution. They are basically a sinusoidal embedding of every dimension of μ . In particular, we extend μ to the vector

$$(\boldsymbol{\mu} \quad \sin(2^i \pi \boldsymbol{\mu}) \quad \cos(2^i \pi \boldsymbol{\mu})) \in \mathbb{R}^{(3+2(n_{\max}-n_{\min})) \cdot n} \qquad i \in n_{\min}, \dots, n_{\max}$$
 (19)

before passing it into the neural network. We choose $n_{\min} = 6$ and $n_{\max} = 8$, in effect increasing the dimensionality of the input to the neural network from n to 7n.

For the neural network itself, we use two architectures, U-Nets (Ronneberger et al., 2015) and Vision Transformers (ViTs) (Dosovitskiy et al., 2020). We use the U-Net configuration proposed by Kingma et al. (2023) which adapts the widely used configuration from (Ho et al., 2020) for likelihood estimation. Most notably, the (Kingma et al., 2023) version has no downsampling between layers of the U-Net, which lets them increase the number of U-Net levels to 32.

ViTs are a more recent architecture inspired by the success of transformers (Vaswani et al., 2017). They represent images as a set of patches with a 2D position embedding and process them with global attention, in contrast to convolutional architectures like the U-Net where communication happens primarily locally. We opt for the Diffusion Transformer (DiT) architecture (Peebles & Xie, 2023) which has been shown to improve sample quality over variants of the (Ho et al., 2020) U-Net model.

6 EXPERIMENTS

We evaluate BSI on the ImageNet (Deng et al., 2009) and CIFAR10 (Krizhevsky, 2009) datasets in terms of log-likelihood and sample quality. While BSI as a method is general and not specific to images, we chose image datasets, because they are established benchmarks in the probabilistic modeling literature. In our experiments, we compare against BFN (Graves et al., 2023) and VDM (Kingma et al., 2023). BFNs are a special case of our framework (see Section 3) and provides an important reference point for the effect of the non-deterministic hyper-prior $p(\mu_0)$ in BSI. VDMs are a representative of the diffusion family of models specifically designed for probabilistic modeling that is structurally similar to BSI as we explained in Section 4.

In Section B, we describe how we compute the ELBO, which we derived in Section 3.1 for continuous x, on discretized images with 8-bit color channels. Section E lists hyperparameters and training details and Section F shows some generated samples.

6.1 IMAGENET32

For this evaluation, we train a DiT-L-2 (Peebles & Xie, 2023) in the BFN, VDM and BSI model, respectively, on the official 32×32 version of ImageNet (Chrabaszcz et al., 2017). We train each model from three seeds for 2 M steps and evaluate the log-likelihood of the test set in bits per dimension (BPD) and the sample quality in terms of Fréchet inception distance (FID) against the test set. For the log-likelihood, we evaluate each model's ELBO with 5 samples from the respective equivalent of $\mathcal{L}_{\rm M}^{\rm c}$ and 2 samples from the respective equivalent of $\mathcal{L}_{\rm R}$. For the

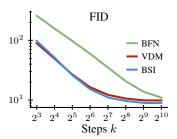
Table 2. Log-likelihood in BPD and sample quality (FID) against the test set on ImageNet32. We compute standard deviations over 3 seeds. Lower is better.

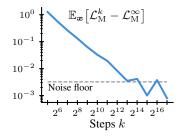
Model	BPD ↓	FID↓
VDM	3.452 ± 0.006	9.9 ± 0.5
BFN	3.448 ± 0.005	11.0 ± 0.1
BSI	3.448 ± 0.006	8.9 ± 0.1

sample quality, we draw 50 000 *unconditional* samples from each model with 1024 steps and then compute the FID between the generated samples and the test set.

Table 2 shows that BSI achieves equivalent likelihoods to VDM and BFN while generating higher-quality samples in terms of FID. This aligns with the result for consistency models by Song & Dhariwal (2024) that a larger variance of the initial state – initial belief μ_0 for BSI – improves the

sample diversity. Ordering the models by improving FID, we have BFN first with an initial variance of 0 ($\mu_0 = 0$), then VDM with an initial variance of 1 and finally BSI with an initial variance of $\lambda_0^{-1} = 100$. Furthermore, Fig. 3 shows that BSI generates significantly better samples than the closely related BFN with fewer steps.





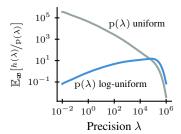


Figure 3. BSI's sample quality converges quickly and to a lower FID with increasing number of steps.

Figure 4. $\mathcal{L}_{\mathrm{M}}^{k}$ converges to $\mathcal{L}_{\mathrm{M}}^{\infty}$ from above as predicted in Lemma 3.3.

Figure 5. Our proposal distribution shrinks the range of $h(\lambda)/p(\lambda)$, reducing ELBO variance.

ELBO Convergence Fig. 4 shows how the finite step ELBO from Theorem 3.1 converges towards its infinite step counterpart as $k \to \infty$ on the test set of ImageNet32. For this plot, we sampled 100 precisions λ per image for the Monte Carlo estimates of $\mathcal{L}_{\mathrm{M}}^k$ and $\mathcal{L}_{\mathrm{M}}^\infty$. The convergence trend continues right to the noise floor where the noise overshadows the signal, marked in the plot by the standard deviation of the Monte Carlo estimator for the difference between the two terms.

6.2 CIFAR10

We train the same U-Net architecture as VDM (Kingma et al., 2023) and BFN (Graves et al., 2023) on CIFAR10. Table 3 shows that BSI achieves equivalent log-likelihoods in terms of BPD. Due to the significant number of training steps (10 M), we followed (Kingma et al., 2023; Graves et al., 2023) and trained only a single model on this dataset.

Variance Reduction Fig. 5 verifies the effect of importance sampling with a log-uniform distribution that we propose in Section 3.3. It reduces the range of the $h(\lambda)/P(\lambda)$ term in Eq. (12) by about 4 orders of magnitude on CIFAR10 and therefore the variance of a Monte Carlo estimate of the ELBO.

Table 3. Test set log-likelihood on CI-FAR10 of the same U-Net in different models.

Model	Training Steps	BPD
VDM BSI	10 M	2.65 2.64
BFN BSI	5 M	2.66 2.65

7 CONCLUSION

We have introduced our generative model BSI through a novel perspective on generative modeling that frames sample generation as iterative Bayesian inference. We have derived an ELBO for both finite steps and the infinite step limit and an importance sampling distribution to minimize the training loss variance. In addition, we have thoroughly discussed how BSI relates to BFN and DMs and shown that BSI includes BFN as a special case. Our experiments have demonstrated that BSI generates better samples than both VDM and BFN while achieving equivalent log-likelihoods on established image datasets. Overall, BSI contributes a Bayesian perspective to the landscape of probabilistic generative modeling that is theoretically simple and empirically effective.

SOFTWARE

For our results, we rely on excellent software packages, notably numpy (Harris et al., 2020), pytorch (Paszke et al., 2019), einops (Rogozhnikov, 2022), matplotlib (Hunter, 2007), h5py (Collette, 2013), hydra (Yadan, 2019) and jupyter (Granger & Pérez, 2021).

REFERENCES

486

487

496

497

501

511

518

532

- Ayadi, S., Hetzel, L., Sommer, J., Theis, F. J., and Günnemann, S. Unified Guidance for Geometry-488 Conditioned Molecular Generation. In Neural Information Processing Systems, November 2024. 489
- 490 Chen, Z., Yuan, H., Li, Y., Kou, Y., Zhang, J., and Gu, Q. Fast Sampling via Discrete Non-Markov 491 Diffusion Models with Predetermined Transition Time. In The Thirty-eighth Annual Conference 492 on Neural Information Processing Systems, November 2024. 493
- Chrabaszcz, P., Loshchilov, I., and Hutter, F. A Downsampled Variant of ImageNet as an Alternative 494 to the CIFAR datasets, August 2017. 495
 - Collette, A. Python and HDF5. O'Reilly, 2013.
- 498 Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A Large-scale Hierarchical 499 Image Database. In Computer Vision and Pattern Recognition Conference, 2009. 500
- Dhariwal, P. and Nichol, A. Diffusion Models Beat GANs on Image Synthesis. In Neural Information Processing Systems, 2021. 502
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., 504 Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An Image is Worth 16x16 505 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning* 506 Representations, October 2020. 507
- 508 Duda, J., Tahboub, K., Gadgil, N. J., and Delp, E. J. The use of asymmetric numeral systems as an 509 accurate replacement for Huffman coding. In 2015 Picture Coding Symposium (PCS), pp. 65–69, May 2015. doi: 10.1109/PCS.2015.7170048. 510
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., 512 and Bengio, Y. Generative Adversarial Networks. In Neural Information Processing Systems. 513 arXiv, 2014. 514
- 515 Granger, B. E. and Pérez, F. Jupyter: Thinking and Storytelling With Code and Data. Computing in 516 Science & Engineering, 23(2):7–14, March 2021. ISSN 1558-366X. doi: 10.1109/MCSE.2021. 517 3059263.
- Graves, A., Srivastava, R. K., Atkinson, T., and Gomez, F. Bayesian Flow Networks, November 519 2023. 520
- 521 Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, 522 E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., 523 Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, 524 T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. Array programming with NumPy. Nature, 585(7825):357-362, September 2020. doi: 10.1038/s41586-020-2649-2. 525
- 526 Harshvardhan, GM., Gourisaria, M. K., Pandey, M., and Rautaray, S. S. A comprehensive survey 527 and analysis of generative models in machine learning. Computer Science Review, 38:100285, 528 November 2020. ISSN 1574-0137. doi: 10.1016/j.cosrev.2020.100285. 529
- 530 Ho, J., Jain, A., and Abbeel, P. Denoising Diffusion Probabilistic Models. In Neural Information 531 Processing Systems, 2020. doi: 10.48550/arXiv.2006.11239.
- Hunter, J. D. Matplotlib: A 2D graphics environment. Computing in Science & Engineering, 9(3): 533 90-95, 2007. doi: 10.1109/MCSE.2007.55. 534
- 535 Karras, T., Aittala, M., Aila, T., and Laine, S. Elucidating the Design Space of Diffusion-Based 536 Generative Models, October 2022. 537
- 538 Kingma, D. P. and Welling, M. Auto-Encoding Variational Bayes, 2013.
 - Kingma, D. P., Salimans, T., Poole, B., and Ho, J. Variational Diffusion Models, April 2023.

547

548

549

550

551 552

553

554

556

561

562

567

568

569

570 571

572

573

574

575

576

577578

- Kollovieh, M., Ansari, A. F., Bohlke-Schneider, M., Zschiegner, J., Wang, H., and Wang, Y. Predict,
 Refine, Synthesize: Self-Guiding Diffusion Models for Probabilistic Time Series Forecasting. In
 Neural Information Processing Systems. arXiv, 2023. doi: 10.48550/arXiv.2307.11494.
- Kollovieh, M., Gosch, L., Lienen, M., Scholten, Y., Schwinn, L., and Günnemann, S. Assessing
 Robustness via Score-Based Adversarial Image Generation. *Transactions on Machine Learning Research*, August 2024a. ISSN 2835-8856.
 - Kollovieh, M., Lienen, M., Lüdke, D., Schwinn, L., and Günnemann, S. Flow Matching with Gaussian Process Priors for Probabilistic Time Series Forecasting. In *International Conference on Learning Representations*, October 2024b.
 - Krizhevsky, A. Learning Multiple Layers of Features from Tiny Images, 2009.
 - Lewis, S., Hempel, T., Jiménez-Luna, J., Gastegger, M., Xie, Y., Foong, A. Y. K., Satorras, V. G., Abdin, O., Veeling, B. S., Zaporozhets, I., Chen, Y., Yang, S., Schneuing, A., Nigam, J., Barbero, F., Stimper, V., Campbell, A., Yim, J., Lienen, M., Shi, Y., Zheng, S., Schulz, H., Munir, U., Clementi, C., and Noé, F. Scalable emulation of protein equilibrium ensembles with generative deep learning, December 2024.
- Lienen, M., Lüdke, D., Hansen-Palmus, J., and Günnemann, S. From Zero to Turbulence: Generative Modeling for 3D Flow Simulation. In *International Conference on Learning Representations*, 2024.
 - Lipman, Y., Chen, R. T. Q., Ben-Hamu, H., Nickel, M., and Le, M. Flow Matching for Generative Modeling, February 2023.
- Lüdke, D., Biloš, M., Shchur, O., Lienen, M., and Günnemann, S. Add and Thin: Diffusion for Temporal Point Processes. In *Neural Information Processing Systems*. arXiv, 2023. doi: 10.48550/arXiv.2311.01139.
 - Lüdke, D., Raventós, E. R., Kollovieh, M., and Günnemann, S. Unlocking Point Processes through Point Set Diffusion, October 2024.
 - Murphy, K. P. Machine Learning: A Probabilistic Perspective. MIT Press, 2012.
 - Nichol, A. and Dhariwal, P. Improved Denoising Diffusion Probabilistic Models. In *International Conference on Machine Learning*, 2021. doi: 10.48550/arXiv.2102.09672.
 - Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Neural Information Processing Systems*, 2019.
 - Peebles, W. and Xie, S. Scalable Diffusion Models with Transformers. In *International Conference on Computer Vision*. arXiv, 2023. doi: 10.48550/arXiv.2212.09748.
- Rezende, D. J. and Mohamed, S. Variational Inference with Normalizing Flows. In *International Conference on Machine Learning*, 2015.
- Rogozhnikov, A. Einops: Clear and Reliable Tensor Manipulations with Einstein-like Notation. In
 International Conference on Learning Representations, 2022.
- Ronneberger, O., Fischer, P., and Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In Navab, N., Hornegger, J., Wells, W. M., and Frangi, A. F. (eds.), *Medical Image Computing and Computer-Assisted Intervention*, Lecture Notes in Computer Science, pp. 234–241, Cham, 2015. Springer International Publishing. ISBN 978-3-319-24574-4. doi: 10.1007/978-3-319-24574-4_28.
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E. L., Ghasemipour, K., Gontijo Lopes,
 R., Karagol Ayan, B., Salimans, T., Ho, J., Fleet, D. J., and Norouzi, M. Photorealistic Text-to Image Diffusion Models with Deep Language Understanding. In *Neural Information Processing Systems*, volume 35, pp. 36479–36494, 2022.

- Salimans, T. and Ho, J. Progressive Distillation for Fast Sampling of Diffusion Models. In *International Conference on Learning Representations*, October 2021.
 - Saydemir, A., Lienen, M., and Günnemann, S. Unfolding Time: Generative Modeling for Turbulent Flows in 4D. In *AI for Science: Scaling in AI for Scientific Discovery Workshop, ICML*, 2024.
 - Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *International Conference on Machine Learning*, 2015.
 - Song, J., Meng, C., and Ermon, S. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*, January 2021a.
 - Song, Y. and Dhariwal, P. Improved Techniques for Training Consistency Models. In *International Conference on Learning Representations*, 2024.
 - Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*, 2021b.
 - Theis, L., van den Oord, A., and Bethge, M. A note on the evaluation of generative models. In *International Conference on Learning Representations*. arXiv, 2016. doi: 10.48550/arXiv.1511. 01844.
 - Townsend, J., Bird, T., Kunze, J., and Barber, D. HiLLoC: Lossless image compression with hierarchical latent variable models. In *International Conference on Learning Representations*, September 2019.
 - Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention Is All You Need. In *Neural Information Processing Systems*, 2017.
 - Xue, K., Zhou, Y., Nie, S., Min, X., Zhang, X., Zhou, J., and Li, C. Unifying Bayesian Flow Networks and Diffusion Models through Stochastic Differential Equations. In *Forty-First International Conference on Machine Learning*, June 2024.
 - Yadan, O. Hydra A framework for elegantly configuring complex applications. Github, 2019.

A How BSI relates to ...

A.1 BAYESIAN FLOW NETWORKS

BFNs are a recent class of generative models for continuous and discrete data motivated from an information-theoretic perspective (Graves et al., 2023). In it, a sender communicates a latent sample to a receiver while trying to minimize the transported data volume. The sender compresses the data with entropy coding, so that minimizing the data volume is equivalent to the receiver maximizing the log-likelihood of the latent sample based on the information that it has received from the sender so far. Finally, a sample can be generated when the receiver also assumes the role of the sender and repeatedly refines its belief.

Our generative approach in Section 3 includes BFN for continuous data as a special case. To see this, we begin by choosing our belief prior $p(\mu_0)$ as $\mathcal{N}_P(\mathbf{0}, \gamma_0)$ and letting $\gamma_0 \to \infty$, i.e. the initial belief mean will always be $\mu_0 = \mathbf{0}$. With Lemma 3.4, this gives us

$$q(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{x}, \lambda) = \mathcal{N}_{P}\left(\frac{\lambda - \lambda_{0}}{\lambda} \, \boldsymbol{x}, \frac{\lambda^{2}}{\lambda - \lambda_{0}}\right). \tag{20}$$

If we now define $\alpha = \lambda - \lambda_0$, choose the initial precision $\lambda_0 = 1$ and write the Normal distribution in variance format, we see that

$$q(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{x}, \lambda) = \mathcal{N}\left(\frac{\alpha}{1+\alpha}\boldsymbol{x}, \frac{\alpha}{(1+\alpha)^{2}}\right), \tag{21}$$

which equals the BFN flow distribution $p_F(\theta \mid x;t)$ (Graves et al., 2023, Equation (76)) if we parametrize λ (and therefore α) in terms of $t \in [0,1]$ as in Section 5.2.

Since a comprehensive description of BFN would go beyond the scope of this work, we will only point out the correspondence between terms from Section 3 and their BFN counterparts without explaining them in detail. For a complete description, we refer the reader to the original work (Graves et al., 2023).

The current belief (μ_i, λ_i) is equivalent to the input distribution p_I (Graves et al., 2023, Equation (43)). Lemma 2.1 is the equivalent of the Bayesian update function h (Graves et al., 2023, Section 4.2). A noisy measurement $y \sim \mathcal{N}_P(x,\alpha)$ corresponds to the sender distribution p_S (Graves et al., 2023, Equation (86)), while a noisy measurement $y \sim \mathcal{N}_P(\hat{x},\alpha)$ of the model's current prediction \hat{x} of the true sample corresponds to the receiver distribution p_R (Graves et al., 2023, Equation (88)). The output distribution p_O and the Bayesian update distribution p_U are just intermediate terms to derive the model and appear neither in the final training nor sampling algorithm.

Fixing the initial belief to $\mu_0=0$ with infinite precision for BFN recovers the behavior described by Graves et al. (2023, Figures 3 and 4) and shown in Eq. (21) that the precision $(1+\alpha)^2/\alpha$ of the flow / encoding distribution $q(\mu_{\lambda}\mid \boldsymbol{x},\lambda)$ in the ELBO first falls and then rises again as α grows. In contrast, with our belief prior $p(\mu_0)=\mathcal{N}_P(\mathbf{0},\lambda_0)$ of the same precision as the initial belief (μ_0,λ_0) as we choose it in Section 3.2, the precision of $q(\mu_{\lambda}\mid \boldsymbol{x},\lambda)$ grows linearly in λ (and α) in lockstep with the precision of the belief (μ_i,λ_i) . We hypothesize that this makes learning for the model easier, because the noise level in its input varies linearly instead of non-linearly across noise levels. Furthermore, in BSI, the first sampling step will already contribute to drawing a random sample, since the initial input μ_0 to f_θ is random. In BFN, the initial belief is fixed to 0, which makes the first sampling step deterministic and equal across all samples.

In Section 3.2, we have argued that the reasonable range of prior precisions γ_0 is $[\lambda_0, \infty]$. BSI and BFN occupy the two extremes of this range with BSI using the least informed prior $\gamma_0 = \lambda_0$, i.e. making the fewest assumptions, and BFN the most informed one $\gamma_0 = \infty$. Note that these extremes are the only choices in the reasonable range for which the precision $\lambda^2(\lambda-\lambda_0+\lambda_0^2/\gamma_0)^{-1}$ of the encoder q in Lemma 3.4 simplifies, i.e. to just λ for BSI and $\lambda^2(\lambda-\lambda_0)^{-1}$ for BFN.

In our comparison to DMs in Section A.2, we see that BSI and BFN also differ in their associated noising process. While BSI's noising process, i.e. how one could go from a more precise measurement back to a less precise one, does not form a Markov chain, BFN's does, making BFN more similar to DMs.

In Section A.2, we exploit that BFN can be represented as a special case of BSI to derive a Markovian forward process for BFN as DMs.

A.2 DIFFUSION MODELS

DMs are a widely used class of generative models built on the concept of inverting a diffusion process (Sohl-Dickstein et al., 2015; Ho et al., 2020). Given a sample x, they define a Markov chain of increasingly noisy versions x_1, x_2, \ldots of x where $x_0 = x$ and

$$p(\boldsymbol{x}_i \mid \boldsymbol{x}_{i-1}) = \mathcal{N}(\alpha_i \boldsymbol{x}_{i-1}, \beta_i)$$
(22)

for some coefficients α_i and β_i . In training, a model learns to invert this Markov chain, which lets you finally generate data by sampling from a noise distribution and stepping along the learned, reverse Markov chain until you reach the data distribution.

While DMs initially achieved prominence in image generation (Dhariwal & Nichol, 2021), they have since been applied successfully across a variety of domains, such as text-to-image mapping (Saharia et al., 2022), fluid simulations (Lienen et al., 2024; Saydemir et al., 2024), adversarial attacks (Kollovieh et al., 2024a), temporal (Lüdke et al., 2023) and general point processes (Lüdke et al., 2024), molecular dynamics (Lewis et al., 2024), molecular structure generation (Ayadi et al., 2024), and time series forecasting (Kollovieh et al., 2023; 2024b).

DMs and BSI are remarkably similar at first glance. Both revolve around the concept of iteratively transforming noise into data samples, though DMs work with Langevin dynamics and BSI uses posterior inference. For training, both models aim to align a parametric distribution $p_{\theta}(x'' \mid x')$ with a distribution $p_{\theta}(x'' \mid x')$ that describes a less noisy version x'' of a noisy sample x' given that the true sample is x.

However, conceptually, they approach sampling from two different perspectives. DMs start with the so-called forward process, where signal is iteratively converted into noise forming a Markov chain of intermediate states as in Eq. (22). Then, they revert this chain to derive the reverse process that enriches noise with data. In contrast, BSI defines the reverse process directly in the form of Lemma D.1 and never uses the associated forward process directly.

We can revert BSI's process to derive its "noising" process. This will let us see what BSI would look like as a DM and thus understand the relationship between the two. Assume that our current belief is $(\mu, \lambda = \lambda_0 + \alpha)$ and we want to denoise further based on a sample x and measurement precision α' , i.e. update our belief to $(\mu', \lambda' = \lambda_0 + \alpha + \alpha')$. The denoising process described by Lemma D.1 tells us that

$$p(\boldsymbol{\mu}' \mid \boldsymbol{\mu}, \boldsymbol{x}) = \mathcal{N}_{P}(1/\lambda' [\lambda \boldsymbol{\mu} + \alpha' \boldsymbol{x}], {\lambda'}^{2}/\alpha').$$
 (23)

To find the noising process, we revert this and get

$$p(\boldsymbol{\mu} \mid \boldsymbol{\mu}', \boldsymbol{x}) = \mathcal{N}\left(\xi^{-1} \left[\frac{\lambda \lambda'}{\alpha'} \boldsymbol{\mu}' + \lambda \left(\frac{\alpha}{\alpha + \lambda_0^2 / \gamma_0} - 1 \right) \boldsymbol{x} \right], \xi \right)$$
(24)

where $\xi = \lambda^2 \left(\left(\alpha + \lambda_0^2 / \gamma_0 \right)^{-1} + {\alpha'}^{-1} \right)$ and γ_0 is the precision of the initial belief prior $p(\mu_0) = \mathcal{N}(\mathbf{0}, \gamma_0)$. Find the proof at the end of this section.

Plugging in $\gamma_0 = \lambda_0$, we get that the noising process of BSI is

$$p(\boldsymbol{\mu} \mid \boldsymbol{\mu}', \boldsymbol{x}) = \mathcal{N}\left(\xi^{-1} \left[\frac{\lambda \lambda'}{\alpha'} \boldsymbol{\mu}' - \lambda_0 \boldsymbol{x}\right], \xi\right) \quad \text{where} \quad \xi = \lambda \left(1 + \frac{\lambda}{\alpha'}\right). \tag{25}$$

Note that this distributions depends on x since $\lambda_0 > 0$. Therefore, BSI's forward process would not be Markov, i.e. you cannot add more noise to a belief state without knowing the sample x that the belief state originated from. While DMs with non-Markov forward processes exist (Song et al., 2021a; Chen et al., 2024), they are uncommon. In conclusion, we see that BSI can be represented as a DM, though with a rather complex, non-Markovian forward process.

As we have shown in Section A.1, BFN are a special case of our generative framework in Section 3 if we choose $\gamma_0 = \infty$. Curiously, Eq. (24) shows that this is the only prior on μ_0 for which the associated forward process is Markov as the coefficient of x becomes 0. This agrees with Xue et al. (2024), who have shown that BFN admit a formulation based on stochastic differential equations, like score-based DMs.

Proof of Eq. (24). We know from Lemma 3.4 that

$$q(\boldsymbol{\mu} \mid \boldsymbol{x}, \lambda) = \mathcal{N}_{P} \left(\frac{\lambda - \lambda_{0}}{\lambda} \, \boldsymbol{x}, \frac{\lambda^{2}}{\lambda - \lambda_{0} + \lambda_{0}^{2} / \gamma_{0}} \right) = \mathcal{N}_{P} \left(\frac{\alpha}{\lambda} \, \boldsymbol{x}, \frac{\lambda^{2}}{\alpha + \lambda_{0}^{2} / \gamma_{0}} \right)$$
(26)

and from Lemma D.1 that

$$p(\boldsymbol{\mu}' \mid \boldsymbol{\mu}, \boldsymbol{x}) = \mathcal{N}_{P}(1/\lambda' [\lambda \boldsymbol{\mu} + \alpha' \boldsymbol{x}], {\lambda'}^{2}/\alpha'). \tag{27}$$

Therefore, $p(\mu, \mu' \mid x)$ is a Gaussian linear system and we can use (Murphy, 2012, Equation (4.125)) to see that

$$p(\boldsymbol{\mu} \mid \boldsymbol{\mu}', \boldsymbol{x}) = \mathcal{N}_{P}(\boldsymbol{\nu}, \boldsymbol{\xi}) \tag{28}$$

with

$$\xi = \lambda^2 \left(\alpha + \frac{\lambda_0^2}{\gamma_0} \right)^{-1} + \left(\frac{\lambda}{\lambda'} \right)^2 \frac{{\lambda'}^2}{\alpha'} = \lambda^2 \left((\alpha + \lambda_0^2/\gamma_0)^{-1} + {\alpha'}^{-1} \right)$$
 (29)

and

$$\boldsymbol{\nu} = \xi^{-1} \left[\frac{\lambda}{\lambda'} \frac{{\lambda'}^2}{\alpha'} \left(\boldsymbol{\mu}' - \frac{\alpha'}{\lambda'} \boldsymbol{x} \right) + \lambda^2 (\alpha + {\lambda_0^2}/{\gamma_0})^{-1} \alpha / \lambda \, \boldsymbol{x} \right]$$
(30)

$$= \xi^{-1} \left[\frac{\lambda \lambda'}{\alpha'} \mu' + \lambda \left(\frac{\alpha}{\alpha + \lambda_0^2 / \gamma_0} - 1 \right) x \right]. \tag{31}$$

B ELBO IN BITS PER DIMENSION

A common metric in probabilistic modeling is the negative log-likelihood of unseen data. The benefits of this metric are that it is theoretically motivated by the probabilistic framework and it can be computed across domains regardless of data modality. If the negative log-likelihood is small, the generative model assigns high likelihood to the unseen data and can thus be regarded as a good model (though likelihood and sample quality are not necessarily the same thing (Theis et al., 2016)). For models that come with an ELBO like BSI, we can use it to upper bound the negative log-likelihood to compare against other ELBO-based or exact-likelihood models.

The negative log-likelihood is usually reported in bits per pixel, per color channel or, in general, per dimension. This unit comes from the fact that an entropy coder could use the model to encode samples $x \in \mathbb{S}^d$ from a finite symbol alphabet \mathbb{S} from the data distribution asymptotically in $-\log_2 p_{\theta}(x)/d$ bits per dimension (Duda et al., 2015). Note that the underlying space \mathbb{S} must be discrete. If it were continuous, $p_{\theta}(x)$ would be a density and the theory would predict that we could compress x into a negative number of bits.

The discreteness requirement is a natural fit for many domains. While, for example, images are usually treated as tensors with continuous color values, the colors are actually stored as discrete values in the range $[0, 2^8 - 1]$ for 8-bit images. Similarly, audio data is a sequence of discrete values in, for example, a 16-bit range.

Let's say that \mathbb{S} is the set of integers $\{0, \dots, r-1\}$. Then we can compute an upper bound on the bits needed to encode $x \in \mathbb{S}^d$ by

$$-\log_2 p(\mathbf{x}) \le \log(2)(\mathcal{L}_{R}' + \mathcal{L}_{M}^{\infty}) \tag{32}$$

as per Theorems 3.1 and 3.2. The multiplication by $\log(2)$ converts the logarithms in \mathcal{L}_R' and \mathcal{L}_M^∞ to base 2. \mathcal{L}_R' is the same as \mathcal{L}_R but with a discretized Normal likelihood to account for the discrete nature of \boldsymbol{x} , i.e.

$$\mathcal{L}'_{R} = \mathbb{E}_{q(\boldsymbol{\mu}_{\lambda_{M}}|\boldsymbol{x},\lambda_{M})} \left[-\log \mathcal{N}'_{P}(\boldsymbol{x} \mid \hat{\boldsymbol{x}}_{\lambda_{M}}, \alpha_{R}) \right]$$
(33)

where

$$\mathcal{N}_{P}'(x_{j} \mid \hat{\boldsymbol{x}}_{\lambda_{M}}, \alpha_{R}) = \Phi(r_{j} \mid \hat{\boldsymbol{x}}_{\lambda_{M}}, \alpha_{R}) - \Phi(l_{j} \mid \hat{\boldsymbol{x}}_{\lambda_{M}}, \alpha_{R}). \tag{34}$$

 $\Phi(r_j \mid \hat{x}_{\lambda_{\rm M}}, \alpha_{\rm R})$ is the CDF of $\mathcal{N}(\hat{x}_{\lambda_{\rm M}}, \alpha_{\rm R})$ and l_j and r_j are the boundaries of the discretization interval containing x_j , i.e.

$$l_{j} = \begin{cases} -\infty & \text{if } x_{j} < \frac{1}{2} \\ r - \frac{3}{2} & \text{if } x_{j} \geq r - \frac{3}{2} \\ \lfloor x_{j} - \frac{1}{2} \rfloor + \frac{1}{2} & \text{otherwise} \end{cases} \quad \text{and} \quad r_{j} = \begin{cases} \infty & \text{if } x_{j} \geq r - \frac{3}{2} \\ \frac{1}{2} & \text{if } x_{j} < \frac{1}{2} \\ \lfloor x_{j} + \frac{1}{2} \rfloor - \frac{1}{2} & \text{otherwise.} \end{cases}$$
(35)

 $\mathcal{L}_{\mathrm{M}}^{\infty}$ is usually not discretized during ELBO computation as the latent variables only enter as a mean squared error instead of a log-likelihood. In a practical implementation, the latent variable distributions would need to be discretized as well, decreasing the ELBO slightly (Kingma et al., 2023; Townsend et al., 2019). If x is discretized to a different set of discrete symbols, e.g. numbers between -1 and 1 instead of the integers \mathbb{S} , the boundaries of the discretization intervals and bin widths in the discretized Normal distribution have to be adapted accordingly.

C PRECONDITIONING DERIVATION

We will assume in this section that the data is normalized such that $\mathbb{E}[x] = 0$ and $\operatorname{Var}[x] = I$.

Assume that we have a current belief (μ, λ) . We derive the parameters c_{skip} , c_{out} and c_{in} of the preconditioned model

$$f_{\theta}(\mu, \lambda) = c_{\text{skip}}\mu + c_{\text{out}}f'_{\theta}(c_{\text{in}}\mu, \lambda)$$
(36)

analogously to Karras et al. (2022). However, while we proceed in the same way, the resulting parameters for BSI differ from Karras et al. (2022) because BSI is not included in the family of DMs that Karras et al. (2022) consider, see Section A.2.

First, we require that $\operatorname{Var}_{\boldsymbol{x}}[c_{\text{in}}\boldsymbol{\mu}] = \boldsymbol{I}$ for all λ . We know from Corollary 3.5 that

$$q(\boldsymbol{\mu} \mid \boldsymbol{x}, \lambda) = \mathcal{N}_{P}((\lambda - \lambda_{0})/\lambda \, \boldsymbol{x}, \lambda). \tag{37}$$

Therefore, $p(x, \mu)$ is a Gaussian linear system and (Murphy, 2012, Equation (4.126)) tells us that the variance of the marginal distribution of μ is

$$\operatorname{Var}_{\boldsymbol{x}}[\boldsymbol{\mu}] = \left(\lambda^{-1} + \frac{(\lambda - \lambda_0)^2}{\lambda^2}\right) \boldsymbol{I}.$$
 (38)

By plugging this into our requirement

$$\operatorname{Var}_{\boldsymbol{x}}[c_{\operatorname{in}}\boldsymbol{\mu}] = c_{\operatorname{in}}^{2} \operatorname{Var}_{\boldsymbol{x}}[\boldsymbol{\mu}] = \boldsymbol{I}, \tag{39}$$

we get immediately that

$$c_{\rm in} = \left(\lambda^{-1} + \frac{(\lambda - \lambda_0)^2}{\lambda^2}\right)^{-1/2} = \left(\underbrace{1 + \frac{(\lambda - \lambda_0)^2}{\lambda}}_{=:\kappa}\right)^{-1/2} \lambda^{1/2} = \sqrt{\lambda/\kappa}.$$
 (40)

Next, we want to have the actual prediction target of f'_{θ} during training to have unit variance, too. In training, we optimize the ELBO from Theorem 3.2, which comes down to minimizing

$$\left\| \boldsymbol{x} - f_{\boldsymbol{\theta}}(\boldsymbol{\mu}, \lambda) \right\|_{2}^{2} \tag{41}$$

up to constant factors only depending on λ . If we plug in Eq. (36) and isolate f'_{θ} , this distance becomes

$$\|\boldsymbol{x} - c_{\text{skip}}\boldsymbol{\mu} - c_{\text{out}}f'_{\boldsymbol{\theta}}(c_{\text{in}}\boldsymbol{\mu}, \lambda)\|_{2}^{2} = c_{\text{out}}^{2} \|f'_{\boldsymbol{\theta}}(c_{\text{in}}\boldsymbol{\mu}, \lambda) - c_{\text{out}}^{-1}(\boldsymbol{x} - c_{\text{skip}}\boldsymbol{\mu})\|_{2}^{2}.$$
 (42)

From this, we identify $c_{\mathrm{out}}^{-1}(\boldsymbol{x}-c_{\mathrm{skip}}\boldsymbol{\mu})$ as the actual training target for $f_{\boldsymbol{\theta}}'$. For the rest of this derivation, we denote use the shorthand $\alpha=\lambda-\lambda_0$ for the measurement precision accumulated in our belief $(\boldsymbol{\mu},\lambda)$. After Corollary 3.5, we can write $\boldsymbol{\mu}$ as $\alpha/\lambda\,\boldsymbol{x}+\boldsymbol{z}$ where $\boldsymbol{z}\sim\mathcal{N}_{\mathrm{P}}(\boldsymbol{0},\lambda)$ and find that the variance of the training target is

$$\operatorname{Var}_{\boldsymbol{x},\boldsymbol{z}}[c_{\text{out}}^{-1}(\boldsymbol{x} - c_{\text{skip}}\boldsymbol{\mu})] = c_{\text{out}}^{-2} \operatorname{Var}_{\boldsymbol{x},\boldsymbol{z}} \left[\boldsymbol{x} - c_{\text{skip}} \left(\frac{\alpha}{\lambda} \, \boldsymbol{x} + \boldsymbol{z} \right) \right]$$

$$= c_{\text{out}}^{-2} \operatorname{Var}_{\boldsymbol{x},\boldsymbol{z}} \left[\left(1 - c_{\text{skip}} \frac{\alpha}{\lambda} \right) \boldsymbol{x} - c_{\text{skip}} \boldsymbol{z} \right]$$

$$= c_{\text{out}}^{-2} \left[\left(1 - c_{\text{skip}} \frac{\alpha}{\lambda} \right)^{2} + c_{\text{skip}}^{2} \lambda^{-1} \right] \boldsymbol{I}$$

$$(43)$$

If we now require the effective training target to have unit variance, we see that

$$c_{\text{out}}^2 = \left(1 - c_{\text{skip}} \frac{\alpha}{\lambda}\right)^2 + c_{\text{skip}}^2 \lambda^{-1} = \left[1 + \frac{\alpha^2}{\lambda}\right] \frac{1}{\lambda} c_{\text{skip}}^2 - 2\frac{\alpha}{\lambda} c_{\text{skip}} + 1. \tag{44}$$

Following Karras et al. (2022), we now choose c_{skip} to minimize the impact of errors in the output of f'_{θ} by minimizing the magnitude of c_{out} . c^2_{out} is a polynomial in c_{skip} with positive leading coefficient, so we can find the minimizer as the root of

$$\frac{1}{2} \frac{\mathrm{d}c_{\text{out}}^2}{\mathrm{d}c_{\text{skip}}} = \left[1 + \frac{\alpha^2}{\lambda}\right] \frac{1}{\lambda} c_{\text{skip}} - \frac{\alpha}{\lambda},\tag{45}$$

which is at

$$c_{\text{skip}} = \left[1 + \frac{\alpha^2}{\lambda}\right]^{-1} \alpha = \kappa^{-1} \alpha = \frac{\alpha}{\kappa}.$$
 (46)

Finally, we can plug $c_{\rm skip}$ into Eq. (44) to get

$$c_{\text{out}}^2 = \kappa \kappa^{-2} \frac{a^2}{\lambda} - 2 \frac{\alpha}{\lambda} \kappa^{-1} \alpha + 1 = \kappa^{-1} \left(\frac{a^2}{\lambda} - 2 \frac{\alpha^2}{\lambda} + \left[1 + \frac{\alpha^2}{\lambda} \right] \right) = \kappa^{-1}$$
 (47)

and consequently $c_{\rm out} = \kappa^{-1/2} = \sqrt{1/\kappa}$.

D PROOFS

D.1 Proof of Theorem 3.1

We will begin with some auxiliary insights. First, we consider the marginal distribution of the updated belief (μ', λ') . This means that our current belief about a sample x is (μ, λ) and now we want to know the distribution of μ' after updating μ with Lemma 2.1 marginalized over all possible noisy measurements y with precision α . Note that λ' is deterministic as it neither depends on x nor y.

Lemma D.1 (Update Marginal). Let $x, \mu \in \mathbb{R}^n$ and $\lambda, \alpha \in \mathbb{R}_+$. Then the distribution of the posterior belief mean μ' marginalized over all measurements y made with precision α is

$$p(\boldsymbol{\mu}' \mid \boldsymbol{\mu}, \boldsymbol{x}, \alpha) = \mathbb{E}_{\boldsymbol{y} \sim \mathcal{N}_{P}(\boldsymbol{x}, \alpha \boldsymbol{I})} \left[p(\boldsymbol{\mu}' \mid \boldsymbol{\mu}, \boldsymbol{x}, \alpha, \boldsymbol{y}) \right] = \mathcal{N}_{P} \left(\frac{1}{\lambda'} \left[\lambda \boldsymbol{\mu} + \alpha \boldsymbol{x} \right], \frac{\lambda'^{2}}{\alpha} \right).$$
(48)

Proof. The noisy measurement is a Normal random variable $y \sim \mathcal{N}_P(x, \alpha)$ and the mean of our posterior belief (μ', λ') after observing y is the deterministic linear transformation

$$\boldsymbol{\mu}' = 1/\lambda' \left[\lambda \boldsymbol{\mu} + \alpha \boldsymbol{y} \right] \tag{49}$$

of this random variable. The statement follows immediately by the linear transformation property of the Normal distribution. \Box

From this, we can see that the update marginal from multiple intermediate measurements is the same as from a single measurement with the combined precision of the intermediate measurements.

Lemma D.2. Let $x, \mu, \mu', \mu'' \in \mathbb{R}^n$ and $\lambda, \alpha, \alpha' \in \mathbb{R}_+$. μ' is the posterior belief mean after a measurement with precision α and μ'' the posterior belief mean after a second, subsequent measurement with precision α' . Then we have that the marginal distribution of the second update is

$$\mathbb{E}_{\mathbf{p}(\boldsymbol{\mu}'|\boldsymbol{\mu},\boldsymbol{x},\alpha)}[\mathbf{p}(\boldsymbol{\mu}''\mid\boldsymbol{\mu}',\boldsymbol{x},\alpha')] = \mathbf{p}(\boldsymbol{\mu}''\mid\boldsymbol{\mu},\boldsymbol{x},\alpha+\alpha'). \tag{50}$$

Proof. We know from Lemma D.1 that μ' is a random variable

$$p(\boldsymbol{\mu}' \mid \boldsymbol{\mu}, \boldsymbol{x}, \alpha) = \mathcal{N}_{P}\left(\underbrace{\frac{1/\lambda' \left[\lambda \boldsymbol{\mu} + \alpha \boldsymbol{x}\right]}_{=:\nu}, \underbrace{\frac{\lambda'^{2}/\alpha}{\lambda'^{2}}}_{=:\xi}\right)}_{=:\xi}$$
(51)

and μ'' is a random variable that depends linearly on μ'

$$p(\boldsymbol{\mu}'' \mid \boldsymbol{\mu}', \boldsymbol{x}, \alpha') = \mathcal{N}_{P}(1/\lambda'' [\lambda' \boldsymbol{\mu}' + \alpha' \boldsymbol{x}], \lambda''^{2}/\alpha').$$
 (52)

As such, they jointly form a Gaussian linear system for which the marginal distribution of μ'' is (Murphy, 2012, Equation (4.126))

$$\mathbb{E}_{\mathbf{p}(\boldsymbol{\mu}'|\boldsymbol{\mu},\boldsymbol{x},\alpha)}[\mathbf{p}(\boldsymbol{\mu}''\mid\boldsymbol{\mu}',\boldsymbol{x},\alpha')] = \mathcal{N}\left(\frac{1}{\lambda''}\left[\lambda'\boldsymbol{\nu} + \alpha'\boldsymbol{x}\right], \frac{\alpha'}{\lambda''^2} + \frac{{\lambda'}^2}{\lambda''^2\xi}\right). \tag{53}$$

Plugging ν into the mean expression and simplifying yields the marginal mean

$$^{1/\lambda^{\prime\prime}}[\lambda\boldsymbol{\mu} + (\alpha + \alpha^{\prime})\boldsymbol{x}]. \tag{54}$$

Similarly, plugging ξ into the covariance expression and simplifying yields the marginal covariance

$$\frac{\alpha + \alpha'}{\lambda''^2}. (55)$$

If we now recall from Lemma 2.1 that

$$\lambda' = \lambda + \alpha \quad \text{and} \quad \lambda'' = \lambda' + \alpha' = \lambda + \alpha + \alpha',$$
 (56)

we can identify Eq. (53) as $p(\mu'' \mid \mu, x, \alpha + \alpha')$.

This trivially generalizes to any finite sequence of measurements, which can be collapsed into a single measurement with the total precision instead.

We will furthermore need to know the KL divergence between the update marginal distributions of the same belief but based on two different samples x and x'.

Lemma D.3. Let $x, x', \mu \in \mathbb{R}^n$ and $\lambda, \alpha \in \mathbb{R}_+$. Then

$$D_{KL}(p(\boldsymbol{\mu}' \mid \boldsymbol{\mu}, \boldsymbol{x}, \alpha), p(\boldsymbol{\mu}' \mid \boldsymbol{\mu}, \boldsymbol{x}', \alpha)) = 1/2 \alpha \|\boldsymbol{x} - \boldsymbol{x}'\|_{2}^{2}.$$
 (57)

Proof. Both update marginal distributions – with x and x' – are Normal distributions of equal precision $\xi := \frac{\lambda'^2}{\alpha}$ as given by Lemma D.1 and respective means of

$$\nu = 1/\lambda' [\lambda \mu + \alpha x]$$
 and $\nu' = 1/\lambda' [\lambda \mu + \alpha x'].$ (58)

As a consequence, the closed form solution for the KL divergence between two equal-covariance Normal distributions becomes

$$D_{KL}(p(\boldsymbol{\mu}' \mid \boldsymbol{\mu}, \boldsymbol{x}, \alpha), p(\boldsymbol{\mu}' \mid \boldsymbol{\mu}, \boldsymbol{x}', \alpha)) = \frac{1}{2} (\boldsymbol{\nu} - \boldsymbol{\nu}')^{\mathsf{T}} \xi (\boldsymbol{\nu} - \boldsymbol{\nu}')$$

$$= \frac{1}{2} (\boldsymbol{x} - \boldsymbol{x}')^{\mathsf{T}} \alpha \lambda'^{-1} \xi \lambda'^{-1} \alpha (\boldsymbol{x} - \boldsymbol{x}')$$

$$= \frac{1}{2} (\boldsymbol{x} - \boldsymbol{x}')^{\mathsf{T}} \alpha (\boldsymbol{x} - \boldsymbol{x}')$$

$$= \frac{1}{2} \alpha \|\boldsymbol{x} - \boldsymbol{x}'\|_{2}^{2}$$
(59)

Equipped with these, we can derive the ELBO.

Theorem 3.1. Let $x \in \mathbb{R}^n$ and $\alpha_R, \alpha_i \in \mathbb{R}_+, i \in [k]$. Then the log-likelihood of x is lower-bounded as

$$\log p(\boldsymbol{x}) \ge -\mathcal{L}_{R} - \mathcal{L}_{M}^{k} \tag{3}$$

by a reconstruction term \mathcal{L}_{R} and a measurement term $\mathcal{L}_{\mathrm{M}}^k$,

$$\mathcal{L}_{\mathrm{R}} = \underset{\mathbf{q}(\boldsymbol{\mu}_{k} \mid \boldsymbol{x}, \lambda_{k})}{\mathbb{E}} \left[-\log \mathcal{N}_{\mathrm{P}}(\boldsymbol{x} \mid \hat{\boldsymbol{x}}_{k}, \alpha_{\mathrm{R}}) \right] \quad and \quad \mathcal{L}_{\mathrm{M}}^{k} = \frac{k}{2} \underset{\mathbf{q}(\boldsymbol{\mu}_{i} \mid \boldsymbol{x}, \lambda_{i})}{\mathbb{E}} \left[\alpha_{i+1} \| \boldsymbol{x} - \hat{\boldsymbol{x}}_{i} \|_{2}^{2} \right] \quad (4)$$

where

$$q(\boldsymbol{\mu}_i \mid \boldsymbol{x}, \lambda_i) = \mathbb{E}_{p(\boldsymbol{\mu}_0)} \left[p(\boldsymbol{\mu}_i \mid \boldsymbol{\mu}_0, \boldsymbol{x}, \lambda_i) \right], \quad \hat{\boldsymbol{x}}_i = f_{\boldsymbol{\theta}}(\boldsymbol{\mu}_i, \lambda_i) \quad and \quad \lambda_i = \lambda_0 + \sum_{i=1}^i \alpha_i. \quad (5)$$

Proof. For any distribution p(x) and any latent variable z, i.e. any choice of prior p(z), encoding distribution $q(z \mid x)$ and likelihood $p(x \mid z)$, we have the variational lower bound

$$\log p(\boldsymbol{x}) \ge - \underset{q(\boldsymbol{z}|\boldsymbol{x})}{\mathbb{E}} [-\log p(\boldsymbol{x} \mid \boldsymbol{z})] - D_{KL}(q(\boldsymbol{z} \mid \boldsymbol{x}), p(\boldsymbol{z}))$$
(60)

on $\log p(x)$ (Kingma & Welling, 2013). In particular, we can choose our sequence of beliefs as the latent variable $z = \{\mu_0, \dots, \mu_k\}$ and define the likelihood of x under this latent variable as

$$p(\boldsymbol{x} \mid \boldsymbol{z}) = \mathcal{N}_{P}(\boldsymbol{x} \mid \hat{\boldsymbol{x}}_{k}, \alpha_{R}). \tag{61}$$

Remember that $\hat{x}_k = f_{\theta}(\mu_k, \lambda_k)$ is the model's estimate of x.

Since the belief means μ_1, \ldots, μ_k are updated only based on their predecessor after Lemma 2.1, they form a Markov chain conditional on x and we can write the encoding distribution as

$$q(\boldsymbol{z} \mid \boldsymbol{x}) = p(\boldsymbol{\mu}_0) \prod_{i=1}^{k} p(\boldsymbol{\mu}_i \mid \boldsymbol{\mu}_{i-1}, \boldsymbol{x}, \alpha_i).$$
 (62)

Each $p(\mu_i \mid \mu_{i-1}, x, \alpha_i)$ is the update marginal of μ_{i-1} over all possible noisy measurements of x with precision α_i from Lemma D.1. Our encoding distribution is ignorant about the influence of x on the initial belief μ_0 , because there is no closed form for $p(\mu_0 \mid x)$. Since we can choose any encoding, not encoding x in μ_0 at all is valid.

If we now plug Eq. (62) into the first term of Eq. (60), we get

$$\mathbb{E}_{\mathbf{q}(\boldsymbol{z}|\boldsymbol{x})}[-\log p(\boldsymbol{x} \mid \boldsymbol{z})] = \mathbb{E}_{\mathbf{p}(\boldsymbol{\mu}_0) \ p(\boldsymbol{\mu}_1 \mid \boldsymbol{\mu}_0, \boldsymbol{x}, \alpha_1)} \mathbb{E}_{\mathbf{p}(\boldsymbol{\mu}_k \mid \boldsymbol{\mu}_{k-1}, \boldsymbol{x}, \alpha_k)}[-\log p(\boldsymbol{x} \mid \boldsymbol{z})]. \tag{63}$$

The intermediate expectations collapse into a single measurement with the sum of all precisions $\bar{\alpha}_i = \sum_{j=1}^i \alpha_j$ according to Lemma D.2, because μ_1, \dots, μ_{k-1} do not appear in the inner log-likelihood, and we are left with

$$\mathbb{E}_{\mathbf{q}(\boldsymbol{z}|\boldsymbol{x})}[-\log p(\boldsymbol{x} \mid \boldsymbol{z})] = \mathbb{E}_{\mathbf{p}(\boldsymbol{\mu}_0) \ p(\boldsymbol{\mu}_k \mid \boldsymbol{\mu}_0, \boldsymbol{x}, \bar{\alpha}_k)} \mathbb{E}_{-\log p(\boldsymbol{x} \mid \boldsymbol{z})].$$
(64)

Since $\lambda_i = \lambda_0 + \sum_{j=1}^i \alpha_j = \lambda_0 + \bar{\alpha}_i$, we can define

$$p(\boldsymbol{\mu}_k \mid \boldsymbol{\mu}_0, \boldsymbol{x}, \lambda) := p(\boldsymbol{\mu}_k \mid \boldsymbol{\mu}_0, \boldsymbol{x}, \alpha = \lambda - \lambda_0) = p(\boldsymbol{\mu}_k \mid \boldsymbol{\mu}_0, \boldsymbol{x}, \bar{\alpha}_k). \tag{65}$$

If we now define

$$q(\boldsymbol{\mu}_k \mid \boldsymbol{x}, \lambda_k) := \mathbb{E}_{p(\boldsymbol{\mu}_0)} [p(\boldsymbol{\mu}_k \mid \boldsymbol{\mu}_0, \boldsymbol{x}, \lambda_k)], \tag{66}$$

we can rewrite Eq. (64) as

$$\mathbb{E}_{\mathbf{q}(\boldsymbol{z}|\boldsymbol{x})}[-\log p(\boldsymbol{x} \mid \boldsymbol{z})] = \mathbb{E}_{\mathbf{q}(\boldsymbol{\mu}_k \mid \boldsymbol{x}, \lambda_k)}[-\log p(\boldsymbol{x} \mid \boldsymbol{\mu}_k)]$$
(67)

which equals the definition of \mathcal{L}_{R} after plugging in Eq. (61).

Next, we investigate the KL-divergence in Eq. (60). We begin by defining the latent prior p(z) autoregressively as

$$p(\boldsymbol{z}) = p(\boldsymbol{\mu}_0) \prod_{i=1}^{k} p(\boldsymbol{\mu}_i \mid \boldsymbol{\mu}_{i-1}, \hat{\boldsymbol{x}}_{i-1}, \alpha_i)$$
(68)

where $\hat{x}_{i-1} = f_{\theta}(\mu_{i-1}, \lambda_{i-1})$ is the point estimate of x produced by our model based on the belief at step i-1. So the prior for μ_i is the update marginal in Lemma D.1 if \hat{x}_{i-1} were the actual sample x.

Now, we plug Eqs. (62) and (68) into the KL-divergence term from Eq. (60).

$$D_{\mathrm{KL}}(\mathbf{q}(\boldsymbol{z} \mid \boldsymbol{x}), \mathbf{p}(\boldsymbol{z})) = \underset{\mathbf{q}(\boldsymbol{z} \mid \boldsymbol{x})}{\mathbb{E}} \left[\log \frac{\mathbf{q}(\boldsymbol{z} \mid \boldsymbol{x})}{\mathbf{p}(\boldsymbol{z})} \right]$$

$$= \underset{\mathbf{q}(\boldsymbol{z} \mid \boldsymbol{x})}{\mathbb{E}} \left[\log \frac{\mathbf{p}(\boldsymbol{\mu}_{0})}{\mathbf{p}(\boldsymbol{\mu}_{0})} + \sum_{i=1}^{k} \log \frac{\mathbf{p}(\boldsymbol{\mu}_{i} \mid \boldsymbol{\mu}_{i-1}, \boldsymbol{x}, \alpha_{i})}{\mathbf{p}(\boldsymbol{\mu}_{i} \mid \boldsymbol{\mu}_{i-1}, \hat{\boldsymbol{x}}_{i-1}, \alpha_{i})} \right]$$

$$= \sum_{i=1}^{k} \underset{\mathbf{q}(\boldsymbol{z} \mid \boldsymbol{x})}{\mathbb{E}} \left[\log \frac{\mathbf{p}(\boldsymbol{\mu}_{i} \mid \boldsymbol{\mu}_{i-1}, \boldsymbol{x}, \alpha_{i})}{\mathbf{p}(\boldsymbol{\mu}_{i} \mid \boldsymbol{\mu}_{i-1}, \hat{\boldsymbol{x}}_{i-1}, \alpha_{i})} \right]$$

$$= \sum_{i=1}^{k} \underset{\mathbf{p}(\boldsymbol{\mu}_{0})}{\mathbb{E}} \underset{\mathbf{p}(\boldsymbol{\mu}_{1} \mid \boldsymbol{\mu}_{0}, \boldsymbol{x}, \alpha_{1})}{\mathbb{E}} \cdots \underset{\mathbf{p}(\boldsymbol{\mu}_{i} \mid \boldsymbol{\mu}_{i-1}, \boldsymbol{x}, \alpha_{i})}{\mathbb{E}} \left[\log \frac{\mathbf{p}(\boldsymbol{\mu}_{i} \mid \boldsymbol{\mu}_{i-1}, \boldsymbol{x}, \alpha_{i})}{\mathbf{p}(\boldsymbol{\mu}_{i} \mid \boldsymbol{\mu}_{i-1}, \hat{\boldsymbol{x}}_{i-1}, \alpha_{i})} \right]$$

$$= \sum_{i=1}^{k} \underset{\mathbf{q}(\boldsymbol{\mu}_{i-1} \mid \boldsymbol{x}, \lambda_{i-1})}{\mathbb{E}} \left[D_{\mathrm{KL}} \left(\mathbf{p}(\boldsymbol{\mu}_{i} \mid \boldsymbol{\mu}_{i-1}, \boldsymbol{x}, \alpha_{i}), \mathbf{p}(\boldsymbol{\mu}_{i} \mid \boldsymbol{\mu}_{i-1}, \hat{\boldsymbol{x}}_{i-1}, \alpha_{i}) \right) \right]$$

$$= (69)$$

The intermediate expectations have collapsed again according to Lemma D.2 in the same way as for the reconstruction term.

We know the closed form for the inner KL divergences from Lemma D.3, so we can further simplify the KL-divergence term to

$$D_{\mathrm{KL}}(q(\boldsymbol{z} \mid \boldsymbol{x}), p(\boldsymbol{z})) = \frac{1}{2} \sum_{i=1}^{k} \mathbb{E}_{q(\boldsymbol{\mu}_{i-1} \mid \boldsymbol{x}, \lambda_{i-1})} \left[\alpha_{i} \| \boldsymbol{x} - \hat{\boldsymbol{x}}_{i-1} \|_{2}^{2} \right].$$
 (70)

Shifting the sum indices by 1 and replacing the sum $\sum_{i=0}^{k-1}$ with $k \mathbb{E}_{i \sim \mathcal{U}(0,k-1)}$ yields $\mathcal{L}_{\mathbf{M}}^k$.

D.2 PROOF OF THEOREM 3.2

Theorem 3.2. Let $\alpha_R, \alpha_M \in \mathbb{R}_+$. For any sequence of precision schedules $\alpha_{k,i}$ for $k \in \mathbb{N}, i \in [k]$ such that $\sum_{i=1}^k \alpha_{k,i} = \alpha_M$ and the sequence of functions $[k] \to \mathbb{R}_+ : i \mapsto \alpha_{k,i}$ converges uniformly to 0, we can take the limit of Theorem 3.1 as $k \to \infty$ to get

$$\mathcal{L}_{R} = \underset{q(\boldsymbol{\mu}_{\lambda_{M}}|\boldsymbol{x},\lambda_{M})}{\mathbb{E}} \left[-\log \mathcal{N}_{P}(\boldsymbol{x} \mid \hat{\boldsymbol{x}}_{\lambda_{M}}, \alpha_{R}) \right] \quad and \quad \mathcal{L}_{M}^{\infty} = \frac{\alpha_{M}}{2} \underset{q(\boldsymbol{\mu}_{\lambda}|\boldsymbol{x},\lambda)}{\mathbb{E}} \left[\|\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda}\|_{2}^{2} \right] \quad (6)$$

where
$$q(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{x}, \lambda) = \mathbb{E}_{p(\boldsymbol{\mu}_{0})} \left[p(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{\mu}_{0}, \boldsymbol{x}, \lambda) \right], \ \lambda_{M} = \lambda_{0} + \alpha_{M} \ \text{and} \ \hat{\boldsymbol{x}}_{\lambda} = f_{\boldsymbol{\theta}}(\boldsymbol{\mu}_{\lambda}, \lambda).$$

Proof. Since $\mathcal{L}_{\mathbb{R}}$ only depends on $\sum_{i} \alpha_{k,i}$ but not individual $\alpha_{k,i}$, the equivalence of the finite and infinite step $\mathcal{L}_{\mathbb{R}}$ is immediately apparent.

For $\mathcal{L}_{\mathrm{M}}^{k}$, we will consider its sum form from Eq. (70).

$$\mathcal{L}_{\mathbf{M}}^{k} = \frac{1}{2} \sum_{i=1}^{k} \mathbb{E}_{\mathbf{q}(\boldsymbol{\mu}_{i-1}|\boldsymbol{x}, \lambda_{i-1})} \left[\alpha_{i} \|\boldsymbol{x} - \hat{\boldsymbol{x}}_{i-1}\|_{2}^{2} \right] = \frac{1}{2} \sum_{i=1}^{k} \alpha_{i} \underbrace{\mathbb{E}_{\mathbf{q}(\boldsymbol{\mu}_{i-1}|\boldsymbol{x}, \lambda_{i-1})} \left[\|\boldsymbol{x} - \hat{\boldsymbol{x}}_{i-1}\|_{2}^{2} \right]}_{=:h(\lambda_{i-1})}$$
(71)

Note that $h(\lambda_{i-1})$ is a deterministic function of λ_{i-1} and $\lambda_0, \dots, \lambda_k$ is a partition of the interval $[\lambda_0, \lambda_0 + \alpha_m] = [\lambda_0, \lambda_{\mathrm{M}}]$ with interval lengths of α_i . It follows that Eq. (71) is a Riemann sum. Since $f_{\boldsymbol{\theta}}$ is a neural network, we can assume that $h(\lambda_{i-1})$ is continuous almost everywhere. Combined with the fact that the interval lengths $\{\alpha_i\}$ converge uniformly to 0, it follows that $\mathcal{L}_{\mathrm{M}}^k$ converges to the Riemann integral

$$\lim_{k \to \infty} \mathcal{L}_{\mathbf{M}}^{k} = \frac{1}{2} \int_{\lambda_{0}}^{\lambda_{\mathbf{M}}} \mathbb{E}_{\mathbf{q}(\boldsymbol{\mu}_{\lambda}|\boldsymbol{x},\lambda)} \left[\|(\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda})\|_{2}^{2} \right] d\lambda$$
 (72)

as $k \to \infty$. It follows trivially that

$$\lim_{k \to \infty} \mathcal{L}_{\mathrm{M}}^{k} = \frac{\alpha_{\mathrm{M}}}{2} \int_{\lambda_{0}}^{\lambda_{\mathrm{M}}} \frac{1}{\alpha_{\mathrm{M}}} \underset{\mathbf{q}(\boldsymbol{\mu}_{\lambda}|\boldsymbol{x},\lambda)}{\mathbb{E}} \left[\|(\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda})\|_{2}^{2} \right] d\lambda$$
 (73)

$$= \frac{\alpha_{\mathrm{M}}}{2} \underset{\substack{\lambda \sim \mathcal{U}(\lambda_{0}, \lambda_{\mathrm{M}}) \\ \mathbf{q}(\boldsymbol{\mu}_{\lambda} | \boldsymbol{x}, \lambda)}}{\mathbb{E}} \left[\| (\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda}) \|_{2}^{2} \right] = \mathcal{L}_{\mathrm{M}}^{\infty}.$$
 (74)

D.3 Proof of Lemma 3.3

Lemma 3.3. If h is strictly decreasing, $\mathcal{L}_{\mathrm{M}}^{\infty} < \mathcal{L}_{\mathrm{M}}^{k}$ for any k and any precision schedule $\{\alpha_{i}\}$.

Proof. In the proof of Theorem 3.2, we have established that $\mathcal{L}_{\mathrm{M}}^{k}$ is a Riemannian sum of h, where we evaluate h on the most-negative edge of each interval. Since h is a non-negative, strictly decreasing function, the discretization error on the interval $[\lambda_{i-1}, \lambda_{i}]$

$$\epsilon := \alpha_i h(\lambda_{i-1}) - \int_{\lambda_{i-1}}^{\lambda_i} h(\lambda) \, \mathrm{d}\lambda \tag{75}$$

is also non-negative. Now consider a refinement of the discretization with $\lambda' \in (\lambda_{i-1}, \lambda_i)$ and the post-refinement discretization error on that interval

$$\epsilon' := (\lambda' - \lambda_{i-1})h(\lambda_{i-1}) + (\lambda_i - \lambda')h(\lambda') - \int_{\lambda_{i-1}}^{\lambda_i} h(\lambda) \, \mathrm{d}\lambda = (\lambda' - \lambda_{i-1} - \alpha_i)h(\lambda_{i-1}) + (\lambda_i - \lambda')h(\lambda') + \epsilon.$$
(76)

Next, we express ϵ' in terms of ϵ as

$$\epsilon' = (\lambda' - \lambda_{i-1} - \alpha_i)h(\lambda_{i-1}) + (\lambda_i - \lambda')h(\lambda') + \epsilon$$

$$= (\lambda' - \lambda_i)h(\lambda_{i-1}) + (\lambda_i - \lambda')h(\lambda') + \epsilon$$

$$= (\lambda_i - \lambda')(h(\lambda') - h(\lambda_{i-1})) + \epsilon.$$
(77)

We know that $(\lambda_i - \lambda') > 0$, because $\lambda' \in (\lambda_{i-1}, \lambda_i)$, and $(h(\lambda') - h(\lambda_{i-1})) < 0$, because h is strictly decreasing. It follows that $\epsilon' < \epsilon$.

This means that any refinement of the ELBO with more steps reduces the non-negative error between the Riemannian sum $\mathcal{L}_{\mathrm{M}}^{k}$ and its limit $\mathcal{L}_{\mathrm{M}}^{\infty}$. In other words, $\mathcal{L}_{\mathrm{M}}^{\infty} < \mathcal{L}_{\mathrm{M}}^{k}$ for all k.

D.4 PROOF OF LEMMA 3.4 AND COROLLARY 3.5

The ELBO in Theorems 3.1 and 3.2 has one part that looks like it might not be so straightforward: the encoding distribution $q(\mu_{\lambda} \mid x, \lambda)$. Its definition contains a marginalization over the belief prior $p(\mu_0)$, which we still need to specify. Let's see what $q(\mu_{\lambda} \mid x, \lambda)$ becomes if we choose a zero-mean, isotropic Normal prior $p(\mu_0)$.

Lemma 3.4. Let $\lambda_0, \gamma_0 \in \mathbb{R}_+$, $p(\boldsymbol{\mu}_0) = \mathcal{N}_P(\mathbf{0}, \gamma_0)$ and $\lambda \geq \lambda_0$. Then

$$q(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{x}, \lambda) = \mathcal{N}_{P}\left(\frac{\lambda - \lambda_{0}}{\lambda} \boldsymbol{x}, \frac{\lambda^{2}}{\lambda - \lambda_{0} + \lambda_{0}^{2}/\gamma_{0}}\right). \tag{9}$$

Proof. Let $p(\mu_{\lambda} \mid \mu_0, x, \lambda)$ be the marginal distribution of μ_{λ} after a measurement of precision $\alpha = \lambda - \lambda_0$, i.e.

$$p(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{\mu}_{0}, \boldsymbol{x}, \lambda) = p(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{\mu}_{0}, \boldsymbol{x}, \alpha = \lambda - \lambda_{0}). \tag{78}$$

We know from Lemma D.1 that

$$p(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{\mu}_{0}, \boldsymbol{x}, \alpha = \lambda - \lambda_{0}) = \mathcal{N}_{P}(1/\lambda[\lambda_{0}\boldsymbol{\mu}_{0} + (\lambda - \lambda_{0})\boldsymbol{x}], \lambda^{2}/(\lambda - \lambda_{0})). \tag{79}$$

Since $p(\mu_0)$ is also Gaussian and μ_{λ} depends linearly on μ_0 , they form a Gaussian linear system for which the marginal distribution of μ_{λ} is (Murphy, 2012, Equation (4.126))

$$q(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{x}, \lambda) = \mathbb{E}_{p(\boldsymbol{\mu}_{0})} \left[p(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{\mu}_{0}, \boldsymbol{x}, \lambda) \right] = \mathcal{N} \left(\frac{1}{\lambda} \left[\lambda_{0} \mathbf{0} + (\lambda - \lambda_{0}) \boldsymbol{x} \right], \frac{\lambda - \lambda_{0}}{\lambda^{2}} + \frac{\lambda_{0}^{2}}{\lambda^{2} \gamma_{0}} \right). \quad (80)$$

By pulling λ^{-2} out of the covariance and inverting to get a precision, we get the claimed result. \Box

If we now choose $\gamma_0 = \lambda_0$, we get the simple BSI prior and the result ELBO encoder.

Corollary 3.5. Let $\lambda_0 \in \mathbb{R}_+$, $p(\mu_0) \sim \mathcal{N}_P(\mathbf{0}, \lambda_0)$ and $\lambda \geq \lambda_0$. Then

$$q(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{x}, \lambda) = \mathcal{N}_{P}\left(\frac{\lambda - \lambda_{0}}{\lambda} \boldsymbol{x}, \lambda\right). \tag{10}$$

Proof. If we choose $\gamma_0 = \lambda_0$ in Lemma 3.4, we get

$$q(\boldsymbol{\mu}_{\lambda} \mid \boldsymbol{x}, \lambda) = \mathcal{N}_{P}\left(\frac{\lambda - \lambda_{0}}{\lambda} \boldsymbol{x}, \frac{\lambda^{2}}{\lambda - \lambda_{0} + \lambda_{0}^{2}/\lambda_{0}}\right). \tag{81}$$

The precision simplifies to

$$\frac{\lambda^2}{\lambda - \lambda_0 + \lambda_0^2 / \lambda_0} = \frac{\lambda^2}{\lambda - \lambda_0 + \lambda_0} = \lambda, \tag{82}$$

proving the result. \Box

D.5 PROOF OF COROLLARY 3.6

Corollary 3.6. Let $p(\lambda)$ be a probability distribution with support $[\lambda_0, \lambda_M]$. Then we have

$$\mathcal{L}_{\mathbf{M}}^{\infty} = \frac{1}{2} \underset{\substack{\lambda \sim \mathbf{p}(\lambda) \\ \mathbf{q}(\boldsymbol{\mu}_{\lambda} | \boldsymbol{x}, \lambda)}}{\mathbb{E}} \left[\frac{1}{\mathbf{p}(\lambda)} \| \boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda} \|_{2}^{2} \right]. \tag{11}$$

Proof. We know from Eq. (72) that $\mathcal{L}_{\mathrm{M}}^{\infty}$ is the following Riemann integral.

$$\mathcal{L}_{\mathbf{M}}^{\infty} = \frac{1}{2} \int_{\lambda_0}^{\lambda_{\mathbf{M}}} \mathbb{E}_{\mathbf{q}(\boldsymbol{\mu}_{\lambda} | \boldsymbol{x}, \lambda)} \left[\| (\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda}) \|_2^2 \right] d\lambda$$
 (83)

Now we can trivially multiply by $p(\lambda)/p(\lambda)$ inside the expectation, proving the statement.

$$\mathcal{L}_{\mathbf{M}}^{\infty} = \frac{1}{2} \int_{\lambda_0}^{\lambda_{\mathbf{M}}} \mathbb{E}_{\mathbf{q}(\boldsymbol{\mu}_{\lambda} | \boldsymbol{x}, \lambda)} \left[\frac{p(\lambda)}{p(\lambda)} \| (\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda}) \|_2^2 \right] d\lambda$$
 (84)

$$= \frac{1}{2} \int_{\lambda_0}^{\lambda_{\rm M}} p(\lambda) \underset{q(\boldsymbol{\mu}_{\lambda}|\boldsymbol{x},\lambda)}{\mathbb{E}} \left[\frac{1}{p(\lambda)} \| (\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda}) \|_2^2 \right] d\lambda$$
 (85)

D.6 PROOF OF Eq. (13)

Proof. We know from Corollary 3.5 that we can write $\mu_{\lambda} = \lambda - \lambda_0 / \lambda x + 1 / \sqrt{\lambda} \varepsilon$ for Gaussian noise $\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ independent of x. Together with the assumption $f_{\theta}(\mu, \lambda) = \mu$, we can rewrite h as

$$h(\lambda) = \underset{\mathbf{q}(\boldsymbol{\mu}_{\lambda}|\boldsymbol{x},\lambda)}{\mathbb{E}} \|\boldsymbol{x} - \hat{\boldsymbol{x}}_{\lambda}\|_{2}^{2}$$

$$= \underset{\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0},\boldsymbol{I})}{\mathbb{E}} \|\boldsymbol{x} - \frac{\lambda - \lambda_{0}}{\lambda} \boldsymbol{x} + \frac{1}{\sqrt{\lambda}} \boldsymbol{\varepsilon} \|_{2}^{2}$$

$$= \underset{\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0},\boldsymbol{I})}{\mathbb{E}} \|\frac{\lambda_{0}}{\lambda} \boldsymbol{x} + \frac{1}{\sqrt{\lambda}} \boldsymbol{\varepsilon} \|_{2}^{2}$$

$$= \underset{\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0},\boldsymbol{I})}{\mathbb{E}} \left(\frac{\lambda_{0}}{\lambda}\right)^{2} \|\boldsymbol{x}\|_{2}^{2} + \frac{1}{\lambda} \|\boldsymbol{\varepsilon}\|_{2}^{2} - 2 \frac{\lambda_{0}}{\sqrt{\lambda^{3}}} \boldsymbol{x} \cdot \boldsymbol{\varepsilon}$$
(86)

If we now make use of our assumption that $\mathbb{E}[x] = 0$ and $\mathrm{Var}[x] = I$, we can distribute the expectation across terms and get

$$\mathbb{E}[h(\lambda)] = \left(\frac{\lambda_0}{\lambda}\right)^2 \underbrace{\mathbb{E}\left[\|\boldsymbol{x}\|_2^2\right]}_{=n} + \frac{1}{\lambda} \underbrace{\mathbb{E}\left[\|\boldsymbol{\varepsilon}\|_2^2\right]}_{=n} - 2\frac{\lambda_0}{\sqrt{\lambda^3}} \underbrace{\mathbb{E}\left[\boldsymbol{x} \cdot \boldsymbol{\varepsilon}\right]}_{=0} \propto \frac{\lambda_0^2}{\lambda^2} + \frac{1}{\lambda}. \tag{87}$$

E EXPERIMENT DETAILS

We trained each model on 4 H100 GPUs at a batch size of 128 on CIFAR10 and 512 on ImageNet32. Training progressed at about 26,300 steps per hour for the U-Net on CIFAR10 and 6,100 steps per hour for the DiT-L-2 backbones on ImageNet32. If we take the different batch sizes into account, the two model architectures needed about equal amounts of training time. Total training time for the 10 M step training on CIFAR10 came to about two weeks.

Furthermore, we take an exponential moving average (EMA) of model weights (Song et al., 2021b; Nichol & Dhariwal, 2021). We provide an overview of the model and training hyperparameters in Table 4, and show the U-Net and DiT parameters in Tables 5 and 6, respectively. On ImageNet32, we train the models with a cosine learning rate scheduler (with linear warm up from 1×10^{-8}) to achieve faster convergence.

To reduce the variance of the training loss further, we use low-discrepancy sampling for t in Algorithm 2 as proposed by Kingma et al. (2023). Instead of sampling b independent t for a batch size of b, we set $t_i = {}^{i-1}/b + \delta \mod 1, i \in [b]$ for a shared $\delta \sim \mathcal{U}(0,1)$ where $\mod 1$ means that we discard the integer part of the result. The marginal distribution of each t_i is $\mathcal{U}(0,1)$, but jointly they cover the [0,1] interval more uniformly than independent samples would, smoothing out the loss across batches.

Table 4. Model and training parameters of BSI on CIFAR10 and all three models on ImageNet32.

	Parameter	CIFAR10	ImageNet32	
I	$lpha_0$	1×10^{-2}		
BSI	$lpha_M$	1×10^{6}		
, ,	α_R	2×10^{6}		
•	Learning rate	2×10^{-4}	5×10^{-4}	
Optim.	LR Scheduler	None	Cosine $\downarrow 5 \times 10^{-5}$	
ďС	Weight decay	1×10^{-2}		
•	Batch size	128	512	
	Steps	10000000	2 000 000	
[A]	β	0.9999		
EMA	First update after step		1000	

Table 5. U-Net hyperparameters for CIFAR10.

Parameter	Value
Hidden dim.	128
Levels	32
Dropout	0.1
Attention heads	1
Convolution padding	Zeros

Table 6. DiT hyperparameters for ImageNet32.

Parameter	Value
Architecture	DiT-L-2
Hidden dim.	1024
Depth	24
Attention heads	16
Dropout	0.05
Patch Size	2

F GENERATED SAMPLES

Fig. 6 shows generated samples from models trained on ImageNet32 for visual reference.

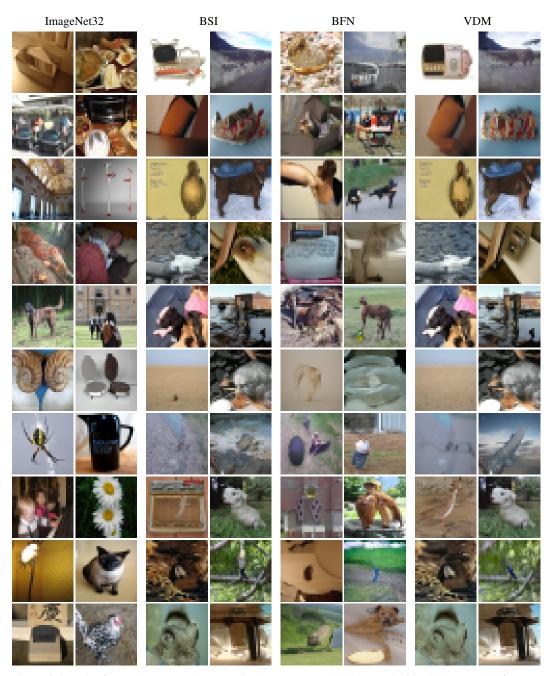


Figure 6. Samples from BSI, BFN and VDM trained on ImageNet32. Generated with 1024 steps. The first two columns show samples from the dataset for comparison.