PREDICTIVE CODING WITH APPROXIMATE LAPLACE MONTE CARLO

Anonymous authors
Paper under double-blind review

ABSTRACT

Predictive coding (PC) accounts of perception now form one of the dominant computational theories of the brain. Despite this, they have enjoyed little export to the broader field of machine learning, where comparative generative models have flourished. In part, this has been due to the poor performance of models trained with standard implementations of PC when evaluated by both sample quality and marginal likelihood. By adopting the perspective of PC as a variational Bayes algorithm under the Laplace approximation, we identify the source of these deficits to lie in the exclusion of an associated Hessian term in the standard PC objective function. To remedy this, we make three primary contributions: we begin by suggesting a simple Monte Carlo estimated evidence lower bound which relies on sampling from the Hessian-parameterised variational posterior. We then derive a novel block diagonal approximation to the full Hessian matrix that has lower memory requirements and favourable mathematical properties. Lastly, we present an algorithm that combines our method with standard PC to reduce memory complexity further. We evaluate models trained with our approach against the standard PC framework on image benchmark datasets. Our approach produces higher log-likelihoods and qualitatively better samples that more closely capture the diversity of the data-generating distribution.

1 INTRODUCTION

In the last two decades, conceptions of the brain as an organ actively engaged in Bayesian inference have become exceedingly prominent in cognitive neuroscience (Pouget et al., 2013; Clark, 2013; Kanai et al., 2015). Under this paradigm, the brain maintains a probabilistic generative model of the world, with perception corresponding to inference over latent states, and learning to the inference over its parameters. Predictive coding (PC) (Rao and Ballard, 1999; Friston, 2018), arguably the most notable instantiation of this perspective, posits a hierarchical generative model composed of normally distributed random variables with possibly complex and highly non-linear parameterisations governing their conditional distributions. Concomitant with this scheme has been a litany of associated empirical predictions and correspondences in neurological circuitry and cognitive phenomena (Mumford, 1992; Hosoya et al., 2005; Hohwy et al., 2008; Bastos et al., 2012; Shipp, 2016; Feldman and Friston, 2010; Fountas et al., 2022), that have furnished the theory with considerable weight.

From a machine learning perspective, PC bares a close mathematical relationship to techniques such as the variational auto-encoder (VAE) (Kingma and Welling, 2014) which rely on optimising an evidence lower bound (ELBO); with a key advantage over VAEs ostensibly being in PC's use of non-amortised inference (Cremer et al., 2018). Furthermore, PC also benefits from design principles inherited from its origins as a theory of cognitive function - namely asynchronous and local error computation (Whittington and Bogacz, 2019), suggesting a far greater amenability to implementation on energy-efficient neuromorphic hardware.

Given these properties, and the ubiquity of PC as a model of perception, it seems pertinent to ask what has prevented its export into the toolbox of generative modelling techniques present in machine learning. Two key arguments for this omission include clear deficits in the performance of generative models trained with PC, when evaluated by both log-likelihood and sample quality, as well as practical difficulties with PC-based training. Both those issues are present in practice but have not found

significant coverage in the existing literature. Here, we discuss these issues, and isolate their origin to the structure of the traditional PC loss function - which we show, fails to regularise for the sharpness of the probability landscape or, equivalently, fails to account for uncertainty over latent states.

To do so, we begin by adopting the perspective of PC as a variational Bayes algorithm under the Laplace approximation (Friston, 2003; 2005; 2008). Under this approximation, quadratic assumptions over the log joint density of a generative model result in a Gaussian variational posterior with precision equal to the Hessian matrix - or curvature - of the negative log joint with respect to its latent states. We present a simple objective function that accounts for this curvature - and thus the uncertainty over latent states - using samples from the Laplace-optimal variational posterior.

To improve upon the memory complexity of computing the full Hessian matrix required for this objective, we then present a novel block diagonal approximation to the Hessian that has lower memory complexity and is guaranteed positive semi-definite (PSD) - ensuring its associated variational posterior can always be sampled from. Finally, to further remove the dependency of memory complexity on the output image dimensionality, we present a combined model, in which the final layer of our generative model is trained with standard PC, and all higher layers are trained with approximate Laplace Monte Carlo. The resultant method has memory complexity reduced to $O(n_L^2)$, from $O(N^2)$ - where n_L, and N are the dimensionalities of the largest latent layer, and all latent layers combined respectively - while retaining improved log likelihoods and sample quality.

2 Related Work and the Laplace Approximation

The Laplace approximation has historically been derived in two contexts. The first and earlier context adopted Laplace's method for the computation of the ordinarily intractable marginalised model evidence after the maximum a posteriori (MAP) value of the latent states had already been identified (Kass and Raftery, 1995; Tierney and Kadane, 1986). The second context adopted the Laplace approximation in the context of variational inference, wherein, under quadratic assumptions for the log joint, it can be shown that the Gaussian variational posterior which minimises the evidence lower bound has inverse covariance equal to the Hessian of the negative log joint probability evaluated at the variational mode (Friston et al., 2007).

Adopting this quadratic approximation over the log joint, and plugging in this optimal posterior into the evidence lower bound results in the following analytical expression (See: Appendix A.2 for a recounting of the full derivation):

$$\log P(x|\theta) \geq \log P(x, \mu_z|\theta) + \frac{1}{2} \log \left[(2\pi)^N \det \mathrm{He}(\theta)^{-1} \right] \tag{1}$$

Here, x and z are a set of observed and latent random variables respectively, μ_z is the mean of the variational Gaussian posterior over latent states, and He is the Hessian associated with the negative log joint ($-\log P(x, z)$) with respect to z, evaluated at the variational mode.

In much of the PC literature, it has been common practice to ignore this second log determinant Hessian term as a further simplification. This is either done explicitly (Buckley et al., 2017; Millidge et al., 2020; Buckley et al., 2017; Whittington and Bogacz, 2017), or implicitly by adopting a point mass for the posterior density (Friston, 2005). In doing so, one optimises the same function during both inference and learning - the log joint - providing us with the following simple recipe for inference and learning over a generative model:

1. Define a graphical model over hidden (z) and observed (x) states with parameters θ (i.e. $\log P(x, z|\theta)$
2. For $x \sim \mathcal{D}$, where $\mathcal{D}$ is the data-generating distribution
   (a) Optimise our variational modes (μ_z) by enacting a gradient descent on $\log P(x, \mu_z|\theta)$
   (b) Update the parameters θ using stochastic gradient descent with respect to $\log P(x, \mu_{\text{MAP}}|\theta)$

It is worthwhile to note that the determinant Hessian of a function evaluated at a critical point is equal to the Gaussian curvature of that function. When this determinant is well-defined (i.e. the Hessian is

positive semi-definite), its eigenvalues are of the same sign, and thus the value of this log determinant is monotonically related to the sharpness (the maximum eigenvalue of the Hessian), a metric directly related to the risk of divergence in gradient descent (See Cohen et al. (2022) for it's use in the context of neural network training). It is unsurprising then that one of the practical difficulties that arise with PC is managing the risk of divergence during the inference procedure as training progress, which in practice is ameliorated by either finely tuning the inference learning rate, or using an adaptive step size. This increased risk of divergence, can therefore be elegantly explained as a failure to regularise for the curvature of the log joint due to the exclusion of the log determinant Hessian term in the variational Laplace Bayes objective.

In the statistical literature, where one is unbeholden to the constraints of neurobiology, this hessian term has been comparatively less neglected. Attempts in this arena have optimised model parameters (θ) with respect to equation 1 or approximations thereof (Bell, 2001). Alternative approaches have used purpose-built auto-differentiation packages to first compute the Hessian, and subsequently computed gradients of model parameters with respect to Laplace importance sampling estimate of the marginal likelihood, (Skaug, 2002; Skaug and Fournier, 2006; Kristensen et al., 2016). The disadvantage of these methods is that their accuracy and variance (Chatterjee and Diaconis, 2018) respectively, are strongly dependent upon the validity of the quadratic approximation. Comparatively, the validity of the objective function we present, as an evidence lower bound, is independent of the validity of the quadratic approximation we make.

3 METHOD

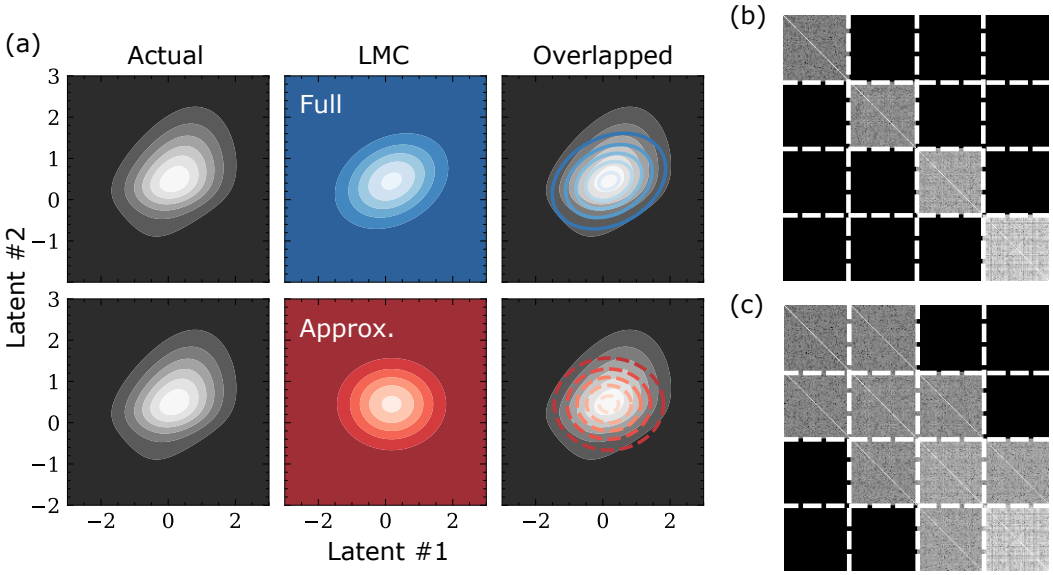

Figure 1: (a) Visualisation of the joint probability, $P(x, z_1, z_2)$ (left), under quadratic (middle) and factorised quadratic approximation for the log joint of a model with two latent states (z_1, z_2) connected by leaky ReLU and a linear coefficient. We can see the approximate quadratic formulation trades off the accuracy of our variational posterior with tractability. (b) Log absolute values of our approximate block-diagonal Hessian for a 4 layer, 64 dimension per layer generative model. (c) Log absolute values for the full Hessian for the same model. Layers are demarcated with dashed white lines.

3.1 THE LAPLACE MONTE CARLO ELBO

One can begin by considering the following decomposition of the ELBO into an expected energy term and an entropy term:

$$\log P(x) \geq \mathrm{E}_{Q(z;\mu_{\text{MAP}}),\Sigma_q}\left[\log P(x,z|\theta)\right] + \mathrm{H}[Q] \tag{2}$$

Since we are already aware of what the optimal approximate posterior is under the Laplace approximation, and since our model parameters are independent of the entropy of this approximate posterior, we may optimise our model parameters with respect to the free-energy by simply considering the first term - the expectation of the log joint probability with respect to our variational posterior. We may then approximate this term by taking Monte Carlo samples from our Laplace-optimal variational posterior and optimise our model parameters via ordinary automatic differentiation.

$$\text{LMC}(\theta) = \frac{1}{K}\sum_i \left[\log P(x, z_i|\theta)\right] \tag{3}$$

Where we are sampling z from our Laplace optimal approximate posterior

$$z_i \sim Q(Z; \mu_{\text{MAP}}, \Sigma_q = \text{He}^{-1}) \tag{4}$$

The advantage of this approach is that while we are sampling from the optimal posterior under the Laplace approximation, unlike the analytical expression (equation 1) described in section 2 we do not require this approximation to be true for the resultant objective to still be an evidence lower bound; with the validity of this approximation instead only impacting the tightness of this bound. Furthermore, in section 4 we demonstrate empirically that despite not optimising with respect to the log determinant Hessian directly, our objective function nonetheless successfully regularises for the sharpness of the probability landscape. We denote models trained with the objective in equation 13 as Laplace Monte Carlo (LMC) models.

3.2 Approximating the Hessian

There remain a number of difficulties present however when working with the Hessian under this approach. First, we require the Hessian of the log-joint to be positive semi-definite as its inverse forms the covariance matrix for the variational Posterior under the Laplace approximation. Second, computing the Hessian has a strict lower-bound computational and memory complexity of $O(N^2)$ where N is the total dimensionality of our latent states across the entire network - which for deep networks can approach the order of tens of thousands to millions.

To circumvent these issues we present an approximation to the full Hessian that retains only curvature information within a layer, resulting in a variational posterior that is factorised across layers. The resultant approximate Hessian also has the desirable property of being guaranteed PSD and thus Monte Carlo estimates of our ELBO are always computable.

We consider a general probabilistic model consisting of a set of latent (unobserved) random variables Z, and observed random variables X. We define a generative model under these random variables factorised such that disjoint subsets of our random variables, $\{z^{(j)}|z^{(j)} \subset Z\}$ and $\{x^{(i)}|x^{(i)} \subset X\}$, have associated with them a multivariate Gaussian conditional distribution with fixed or learnt diagonal covariance matrices $\Sigma^{(i)} = \text{diag}(\sigma^{(i)})$, $\Sigma^{(j)} = \text{diag}(\sigma_j)$; and means $\mu^{(i)}$, $\mu^{(j)}$ parameterised by a function of a subset of the remaining random variables, which we denote with $\mathcal{P}a(z^{(j)})$ and $\mathcal{P}a(x^{(i)})$.

For a single set of observed and latent variables, the negative log joint of this general model can then be defined as follows.

$$\begin{aligned} -\log P(x,z) = &\frac{1}{2}\sum_{\{x^{(i)}\subset X\}} \left(x^{(i)} - f_i(\mathcal{P}a(x^{(i)}),\theta_i)\right)^T \Sigma^{(i)} \left(x^{(i)} - f_i(\mathcal{P}a(x^{(i)}),\theta_i)\right) \\ &+ \frac{1}{2}\sum_{\{z^{(j)}\subset Z\}} \left(z^{(j)} - f_j(\mathcal{P}a(z^{(j)}),\theta_j)\right)^T \Sigma^{(j)} \left(z^{(j)} - f_j(\mathcal{P}a(z^{(j)}),\theta_j)\right) + C \end{aligned} \tag{5}$$

We focus on an approximation to the Hessian of this log-joint where we only consider the second order relations of random variables within the same layer, and not between layers. As such the approximate Hessian derived here will be a block diagonal matrix, which we may then guarantee to be positive semi-definite if the constituent blocks on its diagonal are also guaranteed to be positive semi-definite.

We can ask ourselves what one of these blocks looks like, by first applying the gradient operator with respect to the latent variables of one layer to our unnormalised log joint probability. Note that, solely for the sake of clarity during this derivation, we will exclude any contributions from observed child random variables, i.e. $X^{(k)} \in \mathcal{C}h(z^{(j)})$, for which the steps below follow identically.

$$-\nabla_{z^{(j)}} \log P(X, Z) = \Sigma^{(j)} \underbrace{\left(z^{(j)} - f_j(\mathcal{P}a(z^{(j)}))\right)}_{\text{dim}=N_j * 1} - \sum_{z^{(k)} \in \mathcal{C}h(z^{(j)})} \underbrace{\left(\frac{\partial f_k}{\partial z^{(j)}}\right)^T}_{\text{dim}=N_j * N_k} \Sigma^{(k)} \underbrace{\left(z^{(k)} - f_k(\mathcal{P}a(z^{(k)})\right)}_{\text{dim}=N_k * 1} \tag{6}$$

We can now apply the Jacobian operator to the resultant gradient vector field. But first we decompose the the matrix vector product inside the summation as a sum of vectors multiplied by scalars.

$$-\nabla_{z^{(j)}} \log P(X, Z) = \Sigma^{(j)} \left(z^{(j)} - f_j(\mathcal{P}a(z^{(j)}))\right) - \sum_{z^{(k)} \in \mathcal{C}h(z^{(j)})} \sum_{i \in N_k} \sigma_i^{(k)} \left(\frac{\partial f_k}{\partial z^{(j)}}^T\right)_{*,i} \left(z^{(k)} - f_k(\mathcal{P}a(z^{(k)})\right)_i \tag{7}$$

Now, applying the Jacobian operator.

$$-\mathbf{J}(\nabla_{z^{(j)}} \log P(X, Z)) = \Sigma^{(j)} - \sum_{z^{(k)} \in \mathcal{C}h(z^{(j)})} \sum_{i \in N_k} \sigma_i^{(k)} \left(\frac{\partial^2 f_k}{\partial {z^{(j)}}^2}\right)_{i,*,*} \left(z^{(k)} - f_k(\mathcal{P}a(z^{(k)})\right)_i - \sigma_i^{(k)} \left(\frac{\partial f_k}{\partial z^{(j)}}^T\right)_{*,i} \left(\frac{\partial f_k}{\partial z^{(j)}}\right)_{i,*} \tag{8}$$

We may rewrite the second part of the summation as product of the Jacobian tranopose, diagonal covariance, and Jacobian.

$$-\mathbf{J}(\nabla_{z^{(j)}} \log P(X, Z)) = \Sigma^{(j)} + \sum_{z^{(k)} \in \mathcal{C}h(z^{(j)})} \left(\frac{\partial f_k}{\partial z^{(j)}}\right)^T \Sigma^{(k)} \left(\frac{\partial f_k}{\partial z^{(j)}}\right) - \sum_{z^{(k)} \in \mathcal{C}h(z^{(j)})} \sum_{i \in N_k} \sigma_i^{(k)} \left(\frac{\partial^2 f_k}{\partial {z^{(j)}}^2}\right)_{i,*,*} \left(z^{(k)} - f_k(\mathcal{P}a(z^{(j)})\right)_i \tag{9}$$

We can then choose to either approximate this by ignoring all terms involving second order derivatives, or alternatively, if the functions f_k are piece-wise linear - as is the case of for activations functions such as the (leaky) ReLU - these terms are guaranteed to be 0, resulting in the following:

$$-\mathbf{J}(\nabla_{z^{(j)}} \log P(X, Z)) = \Sigma^{(j)} + \sum_{z^{(k)} \in \mathcal{C}h(z^{(j)})} \left(\frac{\partial f_k}{\partial z^{(j)}}\right)^T \Sigma^{(k)} \left(\frac{\partial f_k}{\partial z^{(j)}}\right) \tag{10}$$

Including the contributions from any observed random variable log likelihood terms, for which the derivation follows identically, results in the final simplified expression:

$$-\mathbf{J}(\nabla_{z^{(j)}} \log P(X, Z)) = \Sigma^{(j)} + \sum_{z^{(k)} \in \mathcal{C}h(z^{(j)})} \left(\frac{\partial f_k}{\partial z^{(j)}}\right)^T \Sigma^{(k)} \left(\frac{\partial f_k}{\partial z^{(j)}}\right) + \sum_{x^{(i)} \in \mathcal{C}h(z^{(j)})} \left(\frac{\partial f_i}{\partial z^{(j)}}\right)^T \Sigma^{(i)} \left(\frac{\partial f_i}{\partial z^{(j)}}\right) \tag{11}$$

Since we have assumed diagonal covariance matrices ($\Sigma_i = \text{diag}(\sigma)$) throughout our generative model, the blocks of our block diagonal Hessian thus simplify to a sum of terms that are guaranteed to be PSD (see Appendix A.1 for a short proof), resulting in a Hessian approximation that is also guaranteed to be PSD. A visualisation of the resultant approximation for a 4 layer model, captured during training, can be seen in Figure 1b alongside the full Hessian (Fig. 1c). We also visualise the joint probability under quadratic, and approximate quadratic assumptions for a model with 2 latent states connected hierarchically alongside the ground truth joint probability in Figure 1a.

$$\mathbf{He}_{\text{Approx}} = \begin{bmatrix} -\mathbf{J}(\nabla_{Z_1}) & 0 & \cdots & 0 \\ 0 & -\mathbf{J}(\nabla_{Z_2}) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & -\mathbf{J}(\nabla_{Z_J}) \end{bmatrix} \tag{12}$$

Because we are only computing curvature information with respect to each layer individually, our memory complexity is also reduced to $O(\max(n_L^2, n_F * n_O))$, where n_L, n_F, n_O are the dimensionalities of the largest latent layer, final latent layer and observation layer respectively, which will generally be significantly lower than $O(N^2)$ for a full Hessian. Note that our memory complexity is not $O(n_L^2)$ as one might expect because of the Jacobians associated with our observed random variable log likelihood terms, a fact that we will address in the next section.

3.3 Combination Models

Naively adopting the aforementioned Hessian approximation may still have impractically high memory requirements as the Jacobians associated with our observed random variables can be exceptionally large if the dimensionality of our observations are high even if the final gram Jacobian matrix is quite small - transiently resulting in high memory requirements. To use an example to illustrate, the Jacobian associated with a small penultimate layer of size 32 predicting an output image with the dimensionality of CelebA ($178 * 218 * 3$) will require temporarily storing a matrix of size $116412 * 32$, even though the final $(\text{Jacobian})^T(\text{Jacobian})$ gram matrix will only be of dimensionality $32 * 32$.

To address this issue and reduce memory complexity further, we take an approach inspired by the recent class of latent diffusion models (Rombach et al., 2022). The key insight here is that one may reduce the untenably high memory requirements of more complex generative models by instead training them on the smaller latent space of an autoencoder trained with the objective of producing good reconstructions. In our experiments, while standard PC networks failed to produce good generative models, they frequently produced excellent reconstructions with very few training samples. Thus we hypothesised that training a combined model, in which lower level/s are trained with the standard PC objective and higher layers are trained with our Approximate-LMC objective, would allow us to combine the strengths of curvature aware training with the reduced memory complexity of the ordinary PC objective.

Unlike the latent diffusion modelling approach, we train our combination models simultaneously and end-to-end. Mathematically, this approach may still be described by expression 13, with lower layers being "sampled" from a dirac delta variational posterior and higher layers being sampled from our (approximate) Laplace-optimal Gaussian posterior, both centred at the MAP estimates found at the

end of inference. Resulting in the following expression, where we have now segmented our latent states into those trained in accordance with LMC (z^L) and those with PC (z^P):

$$\text{LMC}(\theta) = \frac{1}{K} \sum_i \left[\log P(x, z_i^{\text{L}}, \mu_{\text{MAP}}^{\text{P}} | \theta) \right] \tag{13}$$

Where we are sampling z^L from our Laplace optimal approximate posterior:

$$z_i^L \sim Q(Z^L; \mu_{\text{MAP}}^L, \Sigma_q = \text{He}^{-1}) \tag{14}$$

4 EXPERIMENTS

To test the proposed objectives, we trained hierarchical models composed of layers of latent states connected via a non-linearity (leaky ReLU or tanh) followed by an affine transformation - as well as skip connections for adjacent layers with equal dimensionality. The decision to have the non-linearity precede the affine transformation was to ensure that the predicted means for each layer were unbounded - in parity with the unbounded support of our Gaussian variational posterior. Our tanh models were trained with variances fixed to the identity matrix - an architecture frequently adopted by the prior literature, while our leaky ReLU models incorporated learnt variances. To initialise our variational modes we adopted the amortisation scheme described by Tschantz et al. (2022), wherein a feedforward amortisation network is trained alongside our generative model to initialise close to the non-amortised MAP estimates found at the end of inference. All the results we present here adopt the combined approach we describe in 3.3, with $K = 20$, unless otherwise stated, with the final latent layer trained with the standard PC objective, and all higher layers with the LMC objectives. For more extensive details on the model architectures we refer the reader to appendix A.3.

As we note in section 3.2, a significant difficulty with naively using the Hessian is the inability to guarantee positive semi-definiteness - a property which, in our experiments, could result in up to 90% of the samples in a batch having non-PSD associated Hessians in the early stages of training. To accommodate for this issue we experimented with skipping the relevant samples but found that if done so training quickly deteriorates and the issue of non-PSD Hessians exacerbate as training progressed. Thus, for the experiments discussed in this paper Hessians for non-PSD samples were instead replaced with identity matrices, which we found stemmed the continued presence of the problem.

All Hessians and Jacobians in this text were computed using the general purpose implementations of higher-order automatic differentiation in the PyTorch based functorch library (Horace He, 2021). Due to the 7th digit round-off error associated with single-precision floating point numbers, it was possible for the computed Hessians (or approximate Hessian matrices) to, on occasion, be asymmetric; all Hessians (and approximate Hessians) were therefore symmetrised after their computation.

Log Likelihood. We evaluate our models by estimating log marginal likelihoods using Laplace importance sampling of the kind described in (Kuk, 1999). To avoid reporting erroneously inflated log-likelihood values stemming from modelling discrete pixel intensities with continuous models, we follow best practices and dequantize the data by adding uniform noise for both training and evaluation, (Uria et al., 2014); the resultant log-likelihoods are thus guaranteed to be lower bounds on the true log-likelihood on the discrete data distribution (Theis et al., 2016). The resultant log likelihoods (reported as bits per dimension) can be found in Table 4.

Sharpness. We track the log determinant Hessian for 3200 unseen samples throughout training, see Figure 2. We find that Laplace Monte Carlo objectives (approximate or otherwise) regularise for the curvature (log determinant Hessian) despite not explicitly optimising with respect to the term.

Interpolations and Samples. By fixing any particular hierarchical layer of our generative model and enacting a feed-forward top-down pass through it we may visualise what the latent states at each hierarchical layer appear to be representing. We do this while interpolating across the latent embeddings for two images and find that LMC models were more likely to exhibit semantically meaningful hierarchical structure, with higher layers representing global features such as background, hair colour and gender. An example can be seen in Figure 4. Additionally, examples of samples obtained via ancestral sampling can be seen in Figure 4.

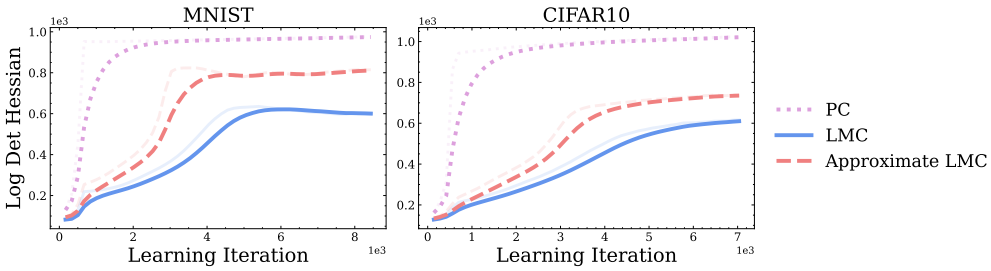

Figure 2: Log determinant Hessian (curvature) for 3200 unseen samples across the training of our models, for all 3 objective functions. Lighter curves are original, and darker curves have exponential smoothing applied.

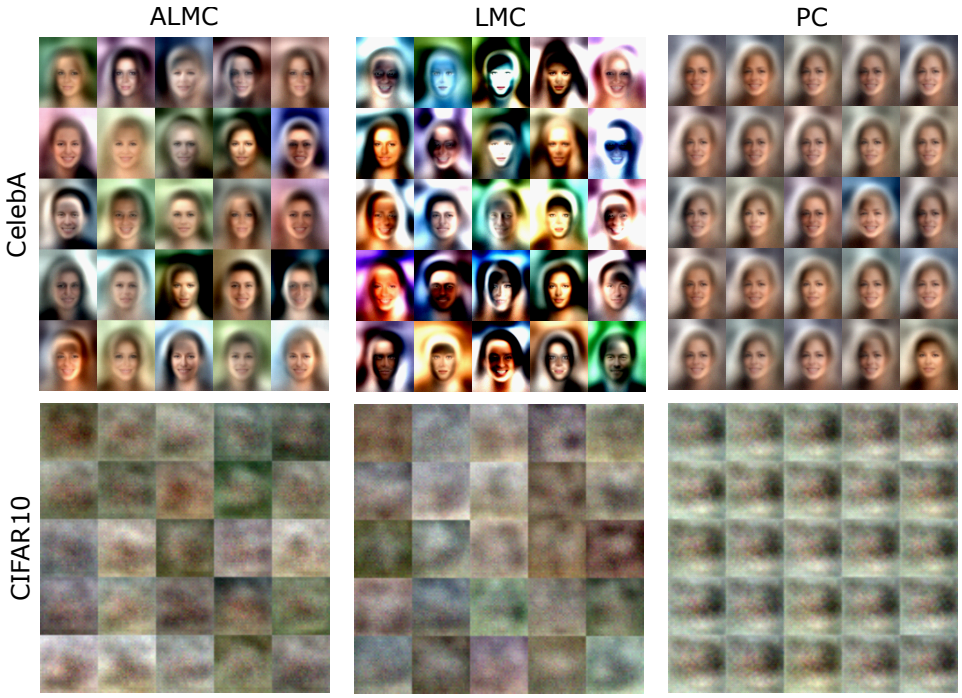

Figure 3: Samples (temp=1) for models trained on CelebA (top) and CIFAR10 (bottom) with the ALMC (left), LMC (center), and standard PC objective (left). While the fidelity of these samples are clearly lacking due to the small size of our tested models, LMC and ALMC capture far more of the training set diversity, while standard PC produces highly uniform samples.

5 Conclusion

While the models evaluated in this work are undoubtedly smaller than contemporary deep learning approaches, they serve as a clear demonstration that the PC objective function as commonly presented is insufficient for producing generative models that capture the diversity of the data generating distribution. We argue that the source of these deficits lie in a failure to accommodate for the uncertainty over our latent states as determined by the curvature of our model's log joint probability.

The simple LMC objective we present here accommodates for this by sampling latent states from a Gaussian centred upon the MAP estimates found at the end of inference, with covariances equal to the inverse Hessian of the negative log joint - an optimal approach under the Laplace approximation. The resultant objective produces consistent improvements in terms of log marginal likelihood and sample

Method	Dataset	BPD
PC	MNIST	6.785
LMC	MNIST	6.731
ALMC	MNIST	**6.727**
PC	CIFAR10	7.008
LMC	CIFAR10	6.931
ALMC	CIFAR10	**6.898**
PC	CelebA	7.1
LMC	CelebA	7.1
ALMC	CelebA	7.1

Method	Dataset	BPD
PC	MNIST	9.69
LMC	MNIST	9.581
ALMC	MNIST	**9.539**
PC	CIFAR10	9.43
LMC	CIFAR10	9.405
ALMC	CIFAR10	**9.392**
PC	CelebA	9.407
LMC	CelebA	9.392
ALMC	CelebA	**9.309**

Method	Dataset	BPD
PC	MNIST	9.69
LMC	MNIST	9.535
ALMC	MNIST	**9.374**
PC	CIFAR10	9.43
LMC	CIFAR10	9.362
ALMC	CIFAR10	**9.352**

Table 1: Bits per dimension (BPD) for various model configurations: leaky ReLU with variance optimisation (left), leaky ReLU with fixed variance (right), and tanh with fixed variance (bottom)

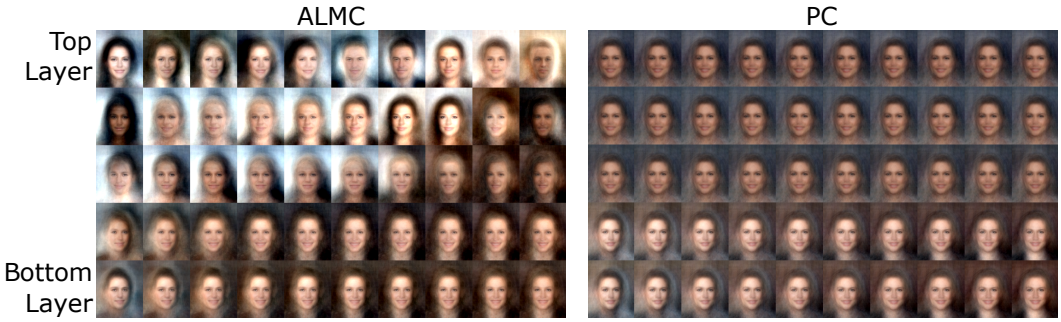

Figure 4: Linear interpolations in the latent states of different hierarchical layers (rows) for ALMC (left) and PC (right). Images refer to a feed-forward top-down pass through our generative model after fixing the latent states at a particular layer.

diversity, but is marred with difficulties that make it impractical such as the inability to guarantee positive semi-definiteness for the Hessian, and poorly scaling memory requirements.

We resolve this by deriving a novel block diagonal approximation to the Hessian that has better memory complexity, and is guaranteed positive semi-definite. Perhaps unexpectedly, the resultant approximate LMC objective consistently performs better than the LMC objective. We attribute this effect to the aforementioned inability to guarantee positive semi-definiteness when utilising the full Hessian, resulting in having to either skip training samples, or as employed here, using a fixed alternative Hessian in its stead.

We also demonstrate that accommodating for this curvature can be interpreted as regularising for the sharpness of the loss landscape, an effect that we empirically verify our LMC and approximate LMC objectives exhibit despite not explicitly optimising with respect to the log determinant Hessian. This has important implications for practical implementations of PC which are frequently ailed by gradient divergence, an effect which is the direct consequent of a loss landscape that is too sharp relative to the step size of ones descent.

REPRODUCIBILITY STATEMENT

Detailed information on all the necessary experimental details required to reproduce our results can be found in section 4 and appendix A.3. Additionally, source code associated with the experiments will be made available after the review process.

REFERENCES

Alexandre Pouget, Jeffrey M. Beck, Wei Ji Ma, and Peter E. Latham. Probabilistic brains: knowns and unknowns. *Nature Neuroscience*, 16(9):1170–1178, September 2013. ISSN 1546-1726. doi:10.1038/nn.3495. URL https://www.nature.com/articles/nn.3495. Number: 9 Publisher: Nature Publishing Group.

Andy Clark. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3):181–204, June 2013. ISSN 0140-525X, 1469-1825. doi:10.1017/S0140525X12000477. Publisher: Cambridge University Press.

Ryota Kanai, Yutaka Komura, Stewart Shipp, and Karl Friston. Cerebral hierarchies: predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1668):20140169, May 2015. doi:10.1098/rstb.2014.0169. URL https://royalsocietypublishing.org/doi/10.1098/rstb.2014.0169. Publisher: Royal Society.

Rajesh P. N. Rao and Dana H. Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87, January 1999. ISSN 1546-1726. doi:10.1038/4580. URL https://www.nature.com/articles/nn0199_79. Number: 1 Publisher: Nature Publishing Group.

Karl Friston. Does predictive coding have a future? *Nature Neuroscience*, 21(8):1019–1021, August 2018. ISSN 1546-1726. doi:10.1038/s41593-018-0200-7. URL https://www.nature.com/articles/s41593-018-0200-7. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 8 Primary_atype: News & Views Publisher: Nature Publishing Group Subject_term: Computational neuroscience;Neural circuits;Neuronal physiology Subject_term_id: computational-neuroscience;neural-circuit;neuronal-physiology.

D. Mumford. On the computational architecture of the neocortex. II. The role of cortico-cortical loops. *Biological Cybernetics*, 66(3):241–251, 1992. ISSN 0340-1200. doi:10.1007/BF00198477.

Toshihiko Hosoya, Stephen A. Baccus, and Markus Meister. Dynamic predictive coding by the retina. *Nature*, 436(7047):71–77, July 2005. ISSN 1476-4687. doi:10.1038/nature03689.

Jakob Hohwy, Andreas Roepstorff, and Karl Friston. Predictive coding explains binocular rivalry: an epistemological review. *Cognition*, 108(3):687–701, September 2008. ISSN 0010-0277. doi:10.1016/j.cognition.2008.05.010.

Andre M. Bastos, W. Martin Usrey, Rick A. Adams, George R. Mangun, Pascal Fries, and Karl J. Friston. Canonical microcircuits for predictive coding. *Neuron*, 76(4):695–711, November 2012. ISSN 1097-4199. doi:10.1016/j.neuron.2012.10.038.

Stewart Shipp. Neural Elements for Predictive Coding. *Frontiers in Psychology*, 7:1792, 2016. ISSN 1664-1078. doi:10.3389/fpsyg.2016.01792. URL https://www.frontiersin.org/article/10.3389/fpsyg.2016.01792.

Harriet Feldman and Karl Friston. Attention, Uncertainty, and Free-Energy. *Frontiers in Human Neuroscience*, 4:215, 2010. ISSN 1662-5161. doi:10.3389/fnhum.2010.00215. URL https://www.frontiersin.org/article/10.3389/fnhum.2010.00215.

Zafeirios Fountas, Anastasia Sylaidi, Kyriacos Nikiforou, Anil K. Seth, Murray Shanahan, and Warrick Roseboom. A Predictive Processing Model of Episodic Memory and Time Perception. *Neural Computation*, 34(7):1501–1544, June 2022. ISSN 0899-7667. doi:10.1162/neco_a_01514. URL https://doi.org/10.1162/neco_a_01514.

Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. *arXiv:1312.6114 [cs, stat]*, May 2014. URL http://arxiv.org/abs/1312.6114. arXiv: 1312.6114.

Chris Cremer, Xuechen Li, and David Duvenaud. Inference Suboptimality in Variational Autoencoders. *arXiv:1801.03558 [cs, stat]*, May 2018. URL http://arxiv.org/abs/1801.03558. arXiv: 1801.03558.

James C.R. Whittington and Rafal Bogacz. Theories of Error Back-Propagation in the Brain. *Trends in Cognitive Sciences*, 23(3):235–250, March 2019. ISSN 13646613. doi:10.1016/j.tics.2018.12.005. URL https://linkinghub.elsevier.com/retrieve/pii/S1364661319300129.

Karl Friston. Learning and inference in the brain. *Neural Networks: The Official Journal of the International Neural Network Society*, 16(9):1325–1352, November 2003. ISSN 0893-6080. doi:10.1016/j.neunet.2003.06.005.

Karl Friston. A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456):815–836, April 2005. ISSN 0962-8436. doi:10.1098/rstb.2005.1622. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1569488/.

Karl Friston. Hierarchical Models in the Brain. *PLOS Computational Biology*, 4(11): e1000211, November 2008. ISSN 1553-7358. doi:10.1371/journal.pcbi.1000211. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1000211. Publisher: Public Library of Science.

Robert E. Kass and Adrian E. Raftery. Bayes Factors. *Journal of the American Statistical Association*, 90(430):773–795, 1995. ISSN 0162-1459. doi:10.2307/2291091. URL https://www.jstor.org/stable/2291091. Publisher: [American Statistical Association, Taylor & Francis, Ltd.].

Luke Tierney and Joseph B. Kadane. Accurate Approximations for Posterior Moments and Marginal Densities. *Journal of the American Statistical Association*, 81(393):82–86, March 1986. ISSN 0162-1459. doi:10.1080/01621459.1986.10478240. URL https://www.tandfonline.com/doi/abs/10.1080/01621459.1986.10478240. Publisher: Taylor & Francis _eprint: https://www.tandfonline.com/doi/pdf/10.1080/01621459.1986.10478240.

Karl Friston, Jérémie Mattout, Nelson Trujillo-Barreto, John Ashburner, and Will Penny. Variational free energy and the Laplace approximation. *NeuroImage*, 34(1):220–234, January 2007. ISSN 1053-8119. doi:10.1016/j.neuroimage.2006.08.035. URL https://www.sciencedirect.com/science/article/pii/S1053811906008822.

Christopher L. Buckley, Chang Sub Kim, Simon McGregor, and Anil K. Seth. The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 81:55–79, December 2017. ISSN 0022-2496. doi:10.1016/j.jmp.2017.09.004. URL https://www.sciencedirect.com/science/article/pii/S0022249617300962.

Beren Millidge, Alexander Tschantz, and Christopher L. Buckley. Predictive Coding Approximates Backprop along Arbitrary Computation Graphs. *arXiv:2006.04182 [cs]*, October 2020. URL http://arxiv.org/abs/2006.04182. arXiv: 2006.04182.

James C. R. Whittington and Rafal Bogacz. An Approximation of the Error Backpropagation Algorithm in a Predictive Coding Network with Local Hebbian Synaptic Plasticity. *Neural Computation*, 29(5):1229–1262, May 2017. ISSN 1530-888X. doi:10.1162/NECO_a_00949.

Jeremy Cohen, Simran Kaur, Yuanzhi Li, J. Zico Kolter, and Ameet Talwalkar. Gradient Descent on Neural Networks Typically Occurs at the Edge of Stability. February 2022. URL https://openreview.net/forum?id=jh-rTtvkGeM.

Bradley M. Bell. Approximating the marginal likelihood estimate for models with random parameters. *Applied Mathematics and Computation*, 119(1):57–75, March 2001. ISSN 0096-3003. doi:10.1016/S0096-3003(99)00230-1. URL https://www.sciencedirect.com/science/article/pii/S0096300399002301.

Hans J Skaug. Automatic Differentiation to Facilitate Maximum Likelihood Estimation in Nonlinear Random Effects Models. *Journal of Computational and Graphical Statistics*, 11(2): 458–470, June 2002. ISSN 1061-8600. doi:10.1198/106186002760180617. URL https://doi.org/10.1198/106186002760180617. Publisher: Taylor & Francis _eprint: https://doi.org/10.1198/106186002760180617.

Hans J. Skaug and David A. Fournier. Automatic approximation of the marginal likelihood in non-Gaussian hierarchical models. *Computational Statistics & Data Analysis*, 51(2):699–709, November 2006. ISSN 0167-9473. doi:10.1016/j.csda.2006.03.005. URL https://www.sciencedirect.com/science/article/pii/S0167947306000764.

Kasper Kristensen, Anders Nielsen, Casper W. Berg, Hans Skaug, and Brad Bell. TMB: Automatic Differentiation and Laplace Approximation. *Journal of Statistical Software*, 70(5), 2016. ISSN 1548-7660. doi:10.18637/jss.v070.i05. URL http://arxiv.org/abs/1509.00660. arXiv:1509.00660 [stat].

Sourav Chatterjee and Persi Diaconis. THE SAMPLE SIZE REQUIRED IN IMPORTANCE SAMPLING. *The Annals of Applied Probability*, 28(2):1099–1135, 2018. ISSN 1050-5164. URL https://www.jstor.org/stable/26542331. Publisher: Institute of Mathematical Statistics.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models, April 2022. URL http://arxiv.org/abs/2112.10752. arXiv:2112.10752 [cs].

Alexander Tschantz, Beren Millidge, Anil K. Seth, and Christopher L. Buckley. Hybrid Predictive Coding: Inferring, Fast and Slow. *arXiv:2204.02169 [cs, q-bio]*, April 2022. URL http://arxiv.org/abs/2204.02169. arXiv: 2204.02169.

Richard Zou Horace He. functorch: JAX-like composable function transforms for PyTorch, 2021. URL https://github.com/pytorch/functorch.

Anthony Y. C. Kuk. Laplace Importance Sampling for Generalized Linear Mixed Models. *Journal of Statistical Computation and Simulation*, 63(2):143–158, April 1999. ISSN 0094-9655. doi:10.1080/00949659908548522. URL https://doi.org/10.1080/00949659908548522. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00949659908548522.

Benigno Uria, Iain Murray, and Hugo Larochelle. RNADE: The real-valued neural autoregressive density-estimator, January 2014. URL http://arxiv.org/abs/1306.0186. arXiv:1306.0186 [cs, stat].

Lucas Theis, Aäron van den Oord, and Matthias Bethge. A note on the evaluation of generative models, April 2016. URL http://arxiv.org/abs/1511.01844. arXiv:1511.01844 [cs, stat].

Arak Mathai and Serge Provost. Quadratic Forms in Random Variables: Theory and Applications. In *Journal of the American Statistical Association*, volume 87. December 1992. doi:10.2307/2290674. Journal Abbreviation: Journal of the American Statistical Association.

A.V. Lazo and P. Rathie. On the entropy of continuous probability distributions (Corresp.). *IEEE Transactions on Information Theory*, 24(1):120–122, January 1978. ISSN 1557-9654. doi:10.1109/TIT.1978.1055832. Conference Name: IEEE Transactions on Information Theory.

A APPENDIX

A.1 PROOF OF PSD

Consider the matrix $M = A^T \text{diag}(\sigma) A$, with sigma being a vector of positive real values. We may rewrite this matrix as follows:

$$M = A^T \text{diag}(\sigma^{-1}) A \tag{15}$$
$$= A^T \text{diag}(\sigma^{-\frac{1}{2}}) \text{diag}(\sigma^{-\frac{1}{2}}) A \tag{16}$$
$$= A^T \text{diag}(\sigma^{-\frac{1}{2}})^T \text{diag}(\sigma^{-\frac{1}{2}}) A \tag{17}$$
$$= (\text{diag}(\sigma^{-\frac{1}{2}}) A)^T (\text{diag}(\sigma^{-\frac{1}{2}}) A) \tag{18}$$

Call $B = \text{diag}(\sigma^{-\frac{1}{2}}) A$, so we can write:

$$M = B^T B \tag{19}$$

Thus M is a real gram matrix and guaranteed to be positive semi-definite.

A.2 DERIVING THE VARIATIONAL LAPLACE OBJECTIVE

Consider a generative model, $\log P(x, z|\theta)$, defined over a set of observed (x) and latent states (z), with model parameters given by θ. We are then concerned with optimising our model parameters (θ) with respect to the marginal likelihood of observations x. For the sake of clarity during this derivation we will drop the dependency on model parameters θ.

$$\log P(x) = \log \left[\int P(x, z) dz \right] = \log \left[\int \frac{Q(z)}{Q(z)} P(x, z) dz \right] = \log \mathbb{E}_{Q(z)} \left[\frac{P(x, z)}{Q(z)} \right] \tag{20}$$

Adopting Jensen's inequality allows us to define the following lower bound.

$$\geq \mathbb{E}_{Q(z)} \left[\log \frac{P(x, z)}{Q(z)} \right] = \mathbb{E}_{Q(z)} \left[\log P(x, z) \right] - \mathbb{E}_{Q(z)} \left[\log Q(z) \right] \tag{21}$$

While the entropy (second) term in equation 21 may have an analytically tractable form in the case of our Gaussian assumptions for Q(z), the first expectation term may not as it is dependent on the exact form of the density function specifying our generative model.

The Laplace approximation ignores the specifics of the log joint density and instead substitutes it with a quadratic approximation using a taylor series centred around the mode of our Gaussian variational distribution μ_z:

$$\mathbb{E}_{Q(z)} \left[\log P(x, z) \right] = \int Q(z) \log P(x, z) dz \tag{22}$$
$$= \int Q(x) \left[\log P(x, \mu_z) + (z - \mu_z) \nabla \log P(x, \mu_z) + \frac{1}{2} (z - \mu_z)^T \text{He}(z - \mu_z) \right] \tag{23}$$
$$= \log P(x, \mu_z) + \underbrace{\mathbb{E}_{Q(z)} \left[(z - \mu_z) \right]}_{=0} \nabla \log P(x, \mu_z) + \frac{1}{2} \mathbb{E}_{Q(z)} \left[(z - \mu_z)^T \text{He}(z - \mu_z) \right] \tag{24}$$

Using a well-known identity on the expectation of a quadratic form (Mathai and Provost, 1992) we may simplify the third term, such that we obtain a significantly simplified expression for our expected energy.

$$= \log P(x, \mu_z) + \frac{1}{2} \text{Tr} \left[\text{He} \Sigma_q \right] \tag{25}$$

We can then reintroduce this simplified expected energy term into our original expression for the free energy, while plugging in the well-known analytic expression for the differential entropy of the Gaussian distribution for the second term (Lazo and Rathie, 1978):

$$\log P(x) \geq \log P(x, \mu_z) + \frac{1}{2}\text{Tr}\left[\text{He}\Sigma_q\right] + \frac{1}{2}\log\left[(2\pi)^N \det \Sigma_q\right] \tag{26}$$

By differentiating this expression with respect to the variational covariance matrix Σ_q we find that the optimal covariance matrix is equal to the inverse of the negative hessian of our log joint probability:

$$\Sigma_q = -\text{He}^{-1} \tag{27}$$

Plugging this into our expression for the free-energy we obtain the following significantly simplified expression for the free-energy:

$$\log P(x) \geq \log P(x, \mu_z) + \frac{1}{2}\log\left[(2\pi)^N \det \Sigma_q\right] \tag{28}$$

If one then optimises this bound with respect to the variational mode, we then obtain:

$$\log P(x) \geq \log P(x, \mu_z) + \frac{1}{2}\log\left[(2\pi)^N \det -\text{He}^{-1}\right] \tag{29}$$

A.3 Model Architecture and Experimental Details

Model Architecture. For the image dataset CelebA, we evaluate models with 5 layers of latent states of dimensionality: [40, 32, 32, 32, 32]. For MNIST and CIFAR10 we evaluate models with 5 layers of dimensionality: [10, 64, 64, 64, 64].

Each layer of the aforementioned models parameterises the mean of the subsequent layer via a non-linearity (either tanh or leaky ReLU), following by an affine transformation. Where adjacent layers have equal dimensionality we also use skip connections.

Amortisation. Our amortisation models adopt the same basic architecture but in reverse, i.e. for a MNIST model, we use an ANN with layer sizes [64,64,64,64,10]. Each layer of the amortisation model is then trained via SGD (with momentum) on a MSE loss between its feed-forward predictions and the MAP latent states identified at the end of inference.

Optimisation of θ. We use batch sizes of 32 for CelebA and 64 for MNIST and CelebA. We use SGD learning rates of 0.02 and 0.0001 for tanh and leaky relu models respectively. We train all generative models using SGD, with momentum set to 0.9. Additionally, for models with learnt variances, we constrain variances to lie between $1e-3$ and 2.

Optimisation of μ_z. We use inference step sizes of 0.05 and 0.001 for tanh and leaky ReLU models respectively and 150 inference steps per batch. Step sizes were reduced by 10% on a per sample basis if an increase in the log joint probability was observed. Furthermore, for experiments with learnt variances we rescale the step size (SS) as max(1e-5, SS*(Minimum Variance)).