Amortised Inference in Structured Generative Models with Explaining Away

Anonymous Author(s)
Affiliation
Address
email

Abstract

A key goal of unsupervised learning is to go beyond density estimation and sample generation to reveal the structure inherent within observed data. Such structure can be expressed in the pattern of interactions between explanatory latent variables captured through a probabilistic graphical model. Although the learning of structured graphical models has a long history, much recent work in unsupervised modelling has instead emphasised flexible deep-network-based generation, either transforming independent latent generators to model complex data or assuming that distinct observed variables are derived from different latent nodes. Here, we extend the output of amortised variational inference to incorporate structured factors over multiple variables, able to capture the observation-induced posterior dependence between latents that results from "explaining away" and thus allow complex observations to depend on multiple nodes of a structured graph. We show that appropriately parameterised factors can be combined efficiently with variational message passing in elaborate graphical structures. We instantiate the framework based on Gaussian Process Factor Analysis models, and empirically evaluate its improvement over existing methods on synthetic data with known generative processes. We then fit the structured model to high-dimensional neural spiking time-series from the hippocampus of freely moving rodents, demonstrating that the model identifies latent signals that correlate with behavioural covariates.

1 Introduction

A central goal of unsupervised learning is the extraction of low-dimensional latent structure that captures dependence between the component variables of high-dimensional observations. Recent developments in deep learning have led to a family of deep generative models (DGM) that have enjoyed remarkable success in density estimation and high-fidelity image/text generation [21, 42, 13, 14, 35]. However, much of the focus with DGMs has been on increasing the expressiveness and accuracy of generation, often overlooking the importance of uncovering the dependency structure inherent within the observed data.

Structured variational inference, by including conditional dependence between the set of latent variables, improves the expressiveness of the variational approximation for the latent posterior distribution. The structured variational autoencoder [SVAE; 18] connects the DGM framework with one element of such inference, combining a structured prior distribution with fully factored mean-field (and functionally conjugate) amortised likelihood potentials. However many generative models of interest involve observations that depend on more than one latent variable. This induces joint potentials in the likelihood, coupling latents even if they are independent *a priori*, a phenomenon sometimes known as "explaining away" [38]. Our goal here is to extend the amortised inference setting to capture such observation-induced dependence in the DGM setting.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

We propose *structured amortised inference* where a structured recognition potential is utilised in combination within the structured autoencoding framework [18]. We instantiate the framework by extending Gaussian Process Factor Analysis (GPFA) models that are widely used in neuroscience, to incorporate deep network-based nonlinear mappings from the latent Gaussian process values to observations. The corresponding likelihood functions induce posterior "explaining-away" correlations between the processes. We develop "amortised explaining-away sparse variational GPFA" (AEA-svGPFA), where we introduced a novel variational approximation and achieve scalable amortised inference of the full-covariance posterior distribution within the sparse inducing-point GP framework. We empirically evaluate AEA-svGPFA on various synthetic and real dataset, including a spike-count dataset collected from a population of neurons in the hippocampal complex, and show that the unsupervised approach automatically captures dimensions underlying neural firing that reflect relevant behavioural correlates.

2 Background

2.1 Variational inference and structured variational autoencoders

We consider the variational inference (VI) problem of approximating the posterior distribution of the latent variables z given observations y. Consider a general generative process

$$z \sim p(z|\theta), \quad y \sim p(y|z;\gamma)$$

where θ and γ are the parameters of the prior and conditional likelihood distributions, respectively. Note that we assume deterministic θ for simplicity, but hyperpriors on θ can be incorporated in the context of Variational Bayes (VB). The posterior distribution, $p(z|y) \propto p(z|\theta)p(y|z,\gamma)$, is often analytically intractable. Hence it is common practice to seek a tractable variational approximation to the true posterior, $q(z)$, which is learned by minimising the KL-divergence with respect to the true posterior distributions, and is equivalent to the maximisation of the variational free energy [47]:

$$q^* = \operatorname*{argmin}_{q\in\mathcal{Q}} \mathrm{KL}[q(z)||p(z|y,\theta)] = \operatorname*{argmax}_{q\in\mathcal{Q}} \langle(\log p(y,z|\gamma,\theta) - \log q(z))\rangle_{q(z)} \tag{1}$$

The Variational Autoencoder (VAE) extends VI by modelling the variational parameters using a recognition network that takes as input the observed values, and is able to jointly train both generative and recognition network parameters with stochastic optimisation of the free energy objective [21, 42].

Although the mean-field variational distribution is often utilised in both standard and amortised VI owing to its simplicity, the assumption of fully factorised latent distributions can be too restrictive to capture the complexity of the true posterior [44, 16]. A number of approaches have been proposed for improving the expressiveness of variational approximation [5, 46, 28, 26]. Here we specifically look at structured VAE [SVAE; 18]. Unlike standard VAEs, SVAE assumes the generative prior distribution to be specified by a structured probabilistic graphical model (PGM), $p(z|\theta) \propto \prod_{c\in C}\psi_c(z_c)$. Moreover, the amortised inference network outputs factors, $r(z|y,\phi)$, that approximate the generative likelihood function rather than the variational distribution itself. These are combined with the structured prior distribution to obtain the variational posterior.

$$q_{\text{SVAE}}(z|y,\theta,\phi) \approx \operatorname*{argmax}_{q\in\mathcal{Q}} \left\langle \log\left(\frac{p(z|\theta)r(z|y,\gamma)}{q(z)}\right)\right\rangle_{q(z)} \propto r(z|y,\phi)p(z|\theta) \tag{2}$$

Johnson et al. [18] assumed that the amortised inference outputs fully factorised factor potentials (that are conjugate with $p(z|\theta)$). Comparing to the mean-field VI in standard VAE, Eq. 2 is able to carry the dependency structure in the prior distribution to the variational posterior distribution. We provide a more in-depth review of the SVAE framework in Appendix.

2.2 Gaussian Process Factor Analysis

Gaussian Process Factor Analysis (GPFA) is a widely used model in neural data analysis for inferring the dynamic latent structure underlying high-dimensional population spike trains [51]. Standard GPFA assumes the following generative model.

$$\begin{aligned} f^k(\cdot) &\sim \mathcal{GP}(m^k_\theta(\cdot), \kappa^k_\theta(\cdot,\cdot)), \text{ for } k = 1,\dots,K, \\ h_n(\cdot) &= \sum_{k=1}^{K} c_{nk} f^k(\cdot) + d_n, \text{ for } n = 1,\dots,N, \\ y_n(t) &= p(y_n(t)|g(h_n(t))), \text{ for } t = 1,\dots,T, \end{aligned} \tag{3}$$

where c_{jk} and d_n are parameters defining the affine mapping from the latent space to the neural space (or a subspace of the neural space in more general likelihood), and $g(\cdot)$ is some (point-wise) smooth function. We assume $m_\theta^k(\cdot) = 0$ unless stated otherwise. The temporal correlation in the high-dimensional neural space are fully captured by the latent temporal correlations modelled by the latent processes. The PGM describing the GPFA generative model is shown in Fig 1a. The parameters of GPFA can be fitted with (variational) expectation-maximisation [9, 51].

However, the applicability of GPFA models is limited by the poor scalability of exact posterior inference of Gaussian processes (GP), which scales cubically with the number of observations [40]. To this end, an extension of the sparse variational approximation for GPs is applied to the GPFA framework, leading to the svGPFA model [45, 1, 12]. Details of svGPFA formulations are included in our discussion of AEA-svGPFA in Section 3.2.

3 Structured amortised inference and explaining away

3.1 Structured amortised inference with SVAE

Non-linear generative likelihood induces additional latent dependency structure that does not exist a priori (we provide a simple example in Appendix B illustrating such effect). Despite introducing the structured prior distribution, the mean-field recognition potential assumption of SVAE only injects singleton factor potential to the prior distribution, prohibiting accurate inference over the latents with changed graphical structure (relative to the prior PGM).

Instead of enforcing a fully factorised form, we assume the recognition network outputs structured factor potentials over z, providing amortised estimates of the additional dependencies induced by "explaining away". Following the SVAE formulation, the "Amortised Explaining Away" (AEA) variational approximation takes the following form.

$$q_{\text{AEA}}(z|y, \phi, \theta) \propto \prod_{c \in C} \psi_c(z_c; \theta) \prod_{c \in C_r} \xi_c(z_c|y; \phi) \tag{4}$$

where C_r is the set of recognition factors and ξ_c is the associated factor potential. We assume conjugate factors unless otherwise stated. Hence the analytical (approximate) form of $q_{AEA}(z|x, \phi, \theta)$ can be computed with (loopy, if necessary) belief propagation. In the most general form, we could let $r(z|y, \phi)$ be a single joint factor potential over the complete set of latent variables.

The free energy objective has the identical functional form as the standard antuencoding free energy objective [21, 42], but with the amortised structured variational approximation, q_{AEA}, as the variational distribution (same applies with the SVAE free energy, with q_{SVAE} (Eq. 2)).

$$\mathcal{F}_{\text{AEA}}(\theta, \phi) = \langle \log p(z|\theta) + \log p(y|z, \gamma) - \log q_{\text{AEA}}(z|y, \phi, \theta) \rangle_{q_{\text{AEA}}(z|y,\phi,\theta)} \tag{5}$$

The training of the model follows standard VAE-style stochastic optimisation to update the parameters of the recognition and generative networks, as well as (optionally) the parameters of the prior PGM.

The following proposition justifies the utility of using structured amortised inference (the proof can be found in appendix D).

Proposition 3.1. *The AEA-SVAE objective function provides a tighter lower bound to the free energy than the SVAE objective function.*

$$\max_{q(z)} \mathcal{F}[q(z)] \geq \mathcal{F}_{AEA}(\theta, \phi) \geq \mathcal{F}_{SVAE}(\theta, \phi) \tag{6}$$

Below we describe one specific instantiation of the AEA framework based on the GPFA models, where we show that our novel variational approximation allows scalable inference of the posterior latent distribution with full covariance structure, hence enabling arbitrary posterior dependency structure to be captured. The AEA framework can be additionally combined with general latent variable models, which we provide more examples in Appendix A.

3.2 Structured autoencoding nonlinear GPFA

Neural computations are thought to be implemented by the joint activity of large populations of neurons. Unsupervised methods have been developed to identify the dynamic covariates that underlie

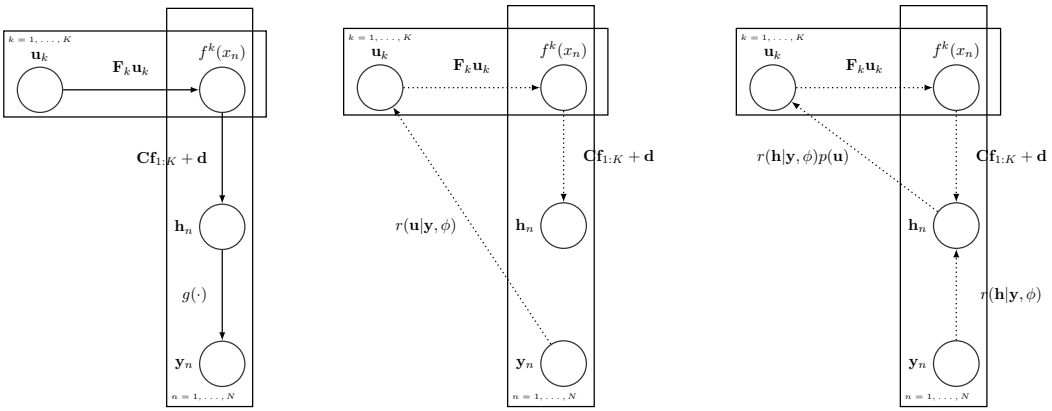

(a) (b) (c)

Figure 1: **Graphical models for the sparse variational GPFA models**. **(a)** Generative model of GPFA model with sparse inducing point approximation; **(b)** Standard GPFA inference with sparse amortised variational approximation; **(c)** Structured amortised inference of AEA-svGPFA model.

the population neuronal spiking trains, and relate these to behaviour or environmental variables so as to gain understanding of the functional and cognitive aspects of neural computations. One popular model is GPFA [51]. In order to alleviate the scalability issue, sparse variational approximation inference for GPFAs with inducing points were proposed [45, 1, 12]. For $k = 1, \ldots, K$, the inducing points $\mathbf{u}_k$ represents the function values of the latent process f^k at the M_k inducing locations, $\mathbf{z}_k$. We then augment the GPFA generative model (Eq. 3) with the inclusion of the inducing points.

$$p(\mathbf{u}_k|\mathbf{z}_k) = \mathcal{N}(0, K^k_{\mathbf{z}_k\mathbf{z}_k}), \quad p(f_k(\cdot)|\mathbf{u}_k) = \mathcal{GP}(\mathbf{F}_k(\cdot)\mathbf{u}_k, \kappa^k_\theta(\cdot,\cdot) - \mathbf{F}_k(\cdot)K^k_{\mathbf{z}_k\mathbf{z}_k}\mathbf{F}^T_k(\cdot)), \tag{7}$$

where $\mathbf{F}_k(\cdot) = \kappa^k_\theta(\cdot, \mathbf{z}_k)(K^k_{\mathbf{z}_k\mathbf{z}_k})^{-1}$ is the linear operator that maps $\mathbf{u}_k$ to $f_k(\cdot)$; $\kappa^k_\theta(\cdot, \mathbf{z}_k)$ is a vector-valued function with $\kappa^k_\theta(x, \mathbf{z}_k) = [\kappa^k_\theta(x, z_{k1}), \kappa^k_\theta(x, z_{k2}), \ldots, \kappa^k_\theta(x, z_{kM_k})]$; and $K^k_{\mathbf{z}_k\mathbf{z}_k}$ is the Gram matrix of evaluation of the kernel function κ^k_θ at the inducing locations $\mathbf{z}_k$.

By introducing a variational distribution over the inducing points and the latent functions, and utilising the generative model (Fig 1a), the free energy of the sparse variational GPFA (svGPFA) can be expressed as following.

$$\mathcal{F}(\theta, \phi, \gamma, \mathbf{C}, \mathbf{d}, \mathbf{Z}) = \sum_{n=1}^{N} \left\langle \log \frac{p(\mathbf{y}_n|\mathbf{h}_n; \gamma)p(\mathbf{h}_n|\mathbf{f}_n; \mathbf{C}, \mathbf{d})p(\mathbf{f}_n|\mathbf{U}, \mathbf{X}, \mathbf{Z})p(\mathbf{U}|\mathbf{Z}; \theta)}{q(\mathbf{f}, \mathbf{U}; \phi)} \right\rangle_{q(\mathbf{f}, \mathbf{U}; \phi)} \tag{8}$$

It is common practice to use the mean-field variational approximation of the joint posterior distribution of the latent process and the inducing points [45, 1], $q(f_{1:K}, \mathbf{u}_{1:K}) = \prod_{k=1}^{K} p(f_k|\mathbf{u}_k)q(\mathbf{u}_k)$, where $q(\mathbf{u}_k) = \mathcal{N}(\mathbf{m}_k, \mathbf{S}_k)$. Hence the mean-field free energy objective simplifies as following.

$$\mathcal{F}_{\text{MF}} = \sum_n \langle \log p(\mathbf{y}_n|\mathbf{h}_n) \rangle_{q(\mathbf{h}_n)} - \sum_{k=1}^{K} \text{KL}[q(\mathbf{u}_k)||p(\mathbf{u}_k|\mathbf{z}_k)]$$

where $q(\mathbf{h}_n) = \int d\mathbf{U} p(\mathbf{h}_n|\mathbf{U})q(\mathbf{U})$. It is easy to derive that $q(\mathbf{h}_n) \sim \mathcal{GP}(m^h_n(\cdot), \nu^h_n(\cdot,\cdot))$ is itself a GP with mean and kernel function given by the following expressions.

$$m^h_n(x) = \sum_{k=1}^{K} C_{nk}\mathbf{F}_k(x)\mathbf{m}_k + d_n, \quad \nu^h_n(x, x') = \sum_{k=1}^{K} C^2_{nk}(\kappa_k(x, x') + \mathbf{F}_k(x)(\mathbf{S}_k - K_{\mathbf{z}_k\mathbf{z}_k})\mathbf{F}_k(x')^T)$$

Hence the computational cost reduces from $\mathcal{O}(T^3)$ down to $\mathcal{O}(\bar{M}^3)$, where $\bar{M} = \frac{1}{K}\sum_{k=1}^{K} M_k$, and scales linearly with respect to the number of observations.

By introducing a recognition network into the svGP framework, we could perform amortised inference over the inducing points, hence improving the scalability of the model. However, naive application of the amortised mean-field VI, $q(\mathbf{U}) = \mathcal{N}(\mathbf{U}|\mathbf{m}(\mathbf{y}; \phi), \mathbf{\Psi}(\mathbf{y}; \phi)$, is inappropriate for GP models since the strong *a priori* temporal correlation will be carried into the posterior distribution. In order to

preserve the within-process latent correlations, Ashman et al. [2] introduced the SGP-VAE, where the variational distribution is computed by combining the prior GP distribution with an amortised approximation of the likelihood function given a set of observation pair, $\{(x_b, \mathbf{y}_b)\}_{b=1}^B$.

$$q(\mathbf{f}, \mathbf{U}) = \prod_k p(f_k(\cdot)|\mathbf{u}_k) q(\mathbf{u}_k), \text{ where } q(\mathbf{u}_k) \propto p(\mathbf{u}_k) \prod_b r(\mathbf{u}_k|\mathbf{y}_b, x_b, \mathbf{z}_k), \tag{9}$$

where $p(\mathbf{U}) = \prod_k \mathcal{N}(\mathbf{m}_\theta^k(\mathbf{z}_k), \mathbf{k}_\theta^k(\mathbf{z}_k, \mathbf{z}_k))$. The graphical model for SGP-VAE inference is shown in Figure 1b, with $\mathbf{C} = \mathbf{I}$ and $\mathbf{d} = 0$. However, we note that this approach still fails to capture correlations *between* the latent processes that arise in the posterior through "explaining away", hence leading to sub-optimal inference and learning.

Our solution is to recast the output of the amortised inference to the likelihood of $\mathbf{h}$ rather than $\mathbf{U}$, and propose the following structured variational distribution.

$$q(\mathbf{f}, \mathbf{U}) = \left[\prod_k p(f_k(\cdot)|\mathbf{u}_k)\right] q(\mathbf{U}), \text{ where } q(\mathbf{U}) \propto \int d\mathbf{h} p(\mathbf{U}) p(\mathbf{h}|\mathbf{U}) \prod_b r(\mathbf{h}|\mathbf{y}_b) \tag{10}$$

where $r(\mathbf{h}|\mathbf{y}_b) = \mathcal{N}(\mathbf{h}|\mathbf{m}(\mathbf{y}_b; \phi), \psi(\mathbf{y}_b; \phi)) = \mathcal{N}(\mu_b^h, \Psi_b^h)$, where Ψ_b^h is diagonal. We note that here we assume $g(\cdot)$ is some general non-linear mapping instead of being a point-wise function as in the standard GPFA model, which could greatly increase the expressiveness of the general modelling. Given the linear-Gaussian relationship between $\mathbf{U}$ and $\mathbf{h}$ (Eq. 11, Fig 1a), the integral can be computed in closed form. Furthermore, despite that the recognition potential on $\mathbf{h}$, $r(\mathbf{h}|\mathbf{y}_b)$, is assumed to be fully factorised over the dimensions of $\mathbf{h}$ for each $\mathbf{y}_b$, the variational distribution on $\mathbf{U}$ includes coupling between latent processes induced by the affine mixing coefficients. Hence $q(\mathbf{U})$ captures the correlations given both by between the latent processes and by the latent temporal dynamics (through combination with the GP prior under the SVAE formulation).

Specifically, given the general expression of the GPFA generative model and the sparse approximation with inducing points, we have the linear Gaussian relationship between $\mathbf{h}$ and $\mathbf{U}$.

$$p(\mathbf{h}|\mathbf{U}, x_b) = \mathcal{N}(\mathbf{C}\mathbf{F}(x_b)\mathbf{U} + \mathbf{d}, \mathbf{C}(\mathbf{K}_x - \mathbf{F}(x_b)\mathbf{K}_\mathbf{U}\mathbf{F}(x_b)^T)\mathbf{C}^T), \tag{11}$$

where

$$\mathbf{F}(x) = \begin{bmatrix} \mathbf{F}_1(x) & & \\ & \ddots & \\ & & \mathbf{F}_K(x) \end{bmatrix}, \quad \mathbf{K}_\mathbf{U} = \begin{bmatrix} \mathbf{K}_{\mathbf{z}_1\mathbf{z}_1} & & \\ & \ddots & \\ & & \mathbf{K}_{\mathbf{z}_K\mathbf{z}_K} \end{bmatrix}$$

with $\mathbf{U} = [\mathbf{u}_1, \ldots, \mathbf{u}_K]$ being the collection of all inducing points for all latents, and $\mathbf{K}_x = \text{diag}[\kappa_\theta^k(x, x)]$. Given the fully factorised recognition potential on $\mathbf{h}_b$ (Eq. 10), we could compute the structured variational distribution on $\mathbf{U}$ given the linear-Gaussian relationship $p(\mathbf{h}|\mathbf{U})$ (Eq. 11).

$$q(\mathbf{U}) = \mathcal{N}(\mathbf{m}_\mathbf{U}, \mathbf{S}_\mathbf{U}) \propto \prod_b p(\mathbf{U}) \mathcal{N}(\mathbf{C}\mathbf{F}(x_b)\mathbf{U}|\mu_b^h, \mathbf{\Psi}_b^h), \text{ with}$$

$$\mathbf{S}_\mathbf{U}^{-1} = \mathbf{K}_\mathbf{U}^{-1} + \sum_b \mathbf{F}(x_b)^T \mathbf{C}^T (\mathbf{\Psi}_b^h)^{-1} \mathbf{C}\mathbf{F}(x_b), \quad \mathbf{m}_\mathbf{U} = \mathbf{S}_\mathbf{U} \left(\sum_b \mathbf{F}(x_b)^T \mathbf{C}^T (\mathbf{\Psi}_b^h)^{-1} (\mu_b^h - \mathbf{d}) \right) \tag{12}$$

Hence giving the $b-$th datapoint, $x_b = x$, we have the following:
$q(\mathbf{h}(x)) \propto p(\mathbf{h}|\mathbf{U}) q(\mathbf{U}) = \mathcal{N}(\mathbf{C}\mathbf{F}(x)\mathbf{m}_\mathbf{U} + \mathbf{d}, \mathbf{C}(\mathbf{K}_n + \mathbf{F}(x)(\mathbf{S}_\mathbf{U} - \mathbf{K}_\mathbf{U})\mathbf{F}(x)^T)\mathbf{C}^T) = \mathcal{N}(\mathbf{m}_b^h, \mathbf{S}_b^h)$

Then the complete (reparameterised) Monte Carlo estimate of the free energy objective given a batch of data, $\{(x_b, \mathbf{y}_b\}_{b=1}^B$, takes the following expression.

$$\mathcal{F}(\theta, \gamma, \phi, \mathbf{C}, \mathbf{d}, \mathbf{Z}) = \sum_{b=1}^B \frac{1}{S} \sum_{s=1}^S \log p(\mathbf{y}_b|\mathbf{m}_b^h + \mathbf{L}_b^h \epsilon_s) - \text{KL}[q(\mathbf{U})||p(\mathbf{U})] \tag{13}$$

where $\epsilon_s \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, θ is the set of kernel parameters, $\mathbf{L}_n^h$ is the lower-triangular matrix such that $\mathbf{S}_n^h = \mathbf{L}_n^h (\mathbf{L}_n^h)^T$, and $\mathbf{Z} = [\mathbf{z}_1, \ldots, \mathbf{z}_K]$ is the set of all inducing locations.

We note that we have retained the scalability of amortised inference by outputting a diagonal amortised Gaussian potential on $\mathbf{h}$, and due to the relationship between the inducing points $\mathbf{U}$ and $\mathbf{h}$

($\mathbf{F}$-projection and affine transformation, Eq. 11), we are able to compute a full-covariance variational Gaussian approximation for $\mathbf{U}$, allowing the computation of posterior correlations between latent processes induced by "explaining away". We term our model as Amortised Explaining Away svGPFA (AEA-svGPFA), which has the same generative model as svGPFA (Figure 1a) but replacing $g(\cdot)$ with some general non-linear mapping modelled by a neural network. The complete inference procedure for AEA-svGPFA described above is graphically illustrated in Figure 1c.

Instead of directly outputting a recognition potential on the inducing points, the variational distribution over $\mathbf{U}$ in AEA-svGPFA is computed implicitly given the recognition potential on $\mathbf{h}$ and the linear-Gaussian relationship between the two variables. Hence it is possible to change the inducing locations and the number of inducing points for inference over varying timescales upon training complete, given by an addition structured svGP inference step (Eq. 12). This is particularly useful in situations when inference over testing dataset of varying length upon training is required, and when the latent processes underlying a particularly long time-span are to be inferred. The latter situation is more important since in practice, the number of inducing points is usually set small relative to the entire set of timepoints for scalability. For instance, when we wish to examine the inferred latent dynamics over the entire time-span, it is very likely we would observe local inconsistencies and abrupt variations that are not reflected in the data, and inference with more inducing points could alleviate such issue.

4 Related Works

Structured Deep Generative Models. Ample literature has studied the inclusion of latent graphical structures into DGM [22, 10, 18, 29, 27, 37]. The prominent work by Johnson et al. [18] proposed SVAE, where the recognition network outputs fully factorised conjugate factor potentials with respect to the structured latent prior, and the variational distribution is computed through partial optimisation of the surrogate free energy given the recognition potentials. The model parameters are jointly optimised by maximising the free energy. In another closely related work, Lin et al. [29] proposed structured inference network, which approximates the variational distribution as the combination of the recognition potential and a separate structured latent distribution independent of the prior, the independent structured latent distribution can have non-conjugate factors to improve the expressiveness of the model. Our model can be interpreted as a strict generalisation of the SVAE framework, where we allow arbitrary structured amortised potential, broadening the flexibility of variational approximation given SVAE. However, the current framework cannot handle models with latent priors containing non-conjugate factors, which we leave as future work. In practice, we extend the recognition potential to contain factor potentials over multiple latent variables whereas Johnson et al. [18], Lin et al. [29] have only exhibited success with mean-field recognition potentials.

svGPFA and Extensions. High-dimensional time-series data, such as population neural firing pattern, usually lie on a low-dimensional manifold that can be characterised by the small number of dynamical covariates that fully explain the temporal correlation exhibited in the observations. Extraction of the latent covariates is particularly essential in analysing neural data, where the set of external measurable variables and internal mental states are potentially entangled and go through a highly non-linear process to influence the neuronal firing. Hence latent variable models such as GPFA are usually adopted for extracting the relevant latent covariates that contains non-trivial temporal correlations [51, 8, 12]. Recent advances have proposed sparse variational inference and methods for dealing with non-conjugate likelihood function to promote the scalability and applicability of GPFA models [1, 12, 19]. There also exists prior works of combining DGM (e.g., VAEs) with GP latents [7, 6, 2]. However, AEA-svGPFA differs from all prior works as it is able to combine svGPFA with amortised inference to allow scalable variational inference of a "free-form" GP posterior (i.e., with full-covariance). The closest model to AEA-svGPFA is the SGP-VAE [2], which similarly incorporate the latent temporal dynamics into posterior inference through the SVAE formulation, and assumes the same generative model as AEA-svGPFA (consider setting $\mathbf{C} = \mathbf{I}$ and $\mathbf{d} = 0$ in Eq. 3), but with different inference procedure from AEA-svGPFA as discussed in Section 3.2. However, since SGP-VAE directly imposes a fully factorised recognition distribution on the inducing points, it i) cannot capture the posterior correlations between the latent processes; ii) does not allow additional svGP inference step with changed inducing points.

5 Results

Here we empirically examine the performance of AEA-svGPFA and compare to relevant baselines on various synthetic and real datasets. Empirical results on the instantiations of the AEA framework with other latent variable models can be found in Appendix E.

5.1 Synthetic and EEG Dataset

We firstly empirically evaluate whether AEA-svGPFA outperforms methods that do not capture the intra-latent posterior correlation. We use a synthetic dataset and a small-scale EEG dataset for comparison. Note that we assume Squared Exponential (SE) kernel, $\kappa(x, x') = \lambda \exp(-\frac{||x-x'||^2}{\tau^2})$, unless stated otherwise, where λ and τ are the scale and length-scale parameters, respectively.

Baselines Our main baseline model is SGP-VAE [2], which we introduce and provide discussions relative to AEA-svGPFA in Section 3.2 and 4, respectively. For completeness, we additional compare against a vanilla VAE with fully factorised Gaussian variational distribution [21, 42], and SVAE with linear dynamical system (SVAE-LDS) latent prior [18]. All models are implemented with the same recognition and generative network architectures (see Appendix F for implementation details).

Synthetic Dataset We generate from the GPFA generative model (Eq. 3), with $g(\cdot) = \sigma(\Phi(\cdot))$, where $\Phi(\cdot)$ represents the functional mapping through a fixed, randomly-initialised 2-layer MLP with ReLU hidden non-linearity, and σ is the sigmoid function.

EEG Dataset We follow the experimental procedure as described in Requeima et al. [41], and consider an EEG measurement dataset spanning 1 second at $256Hz$ frequency when images are shown to the patients [52]. Each datapoint consists of the voltage readings from 7 electrodes positioned on the patient's scalp. We assume all data are observed, contrary to the settings in [41, 2] (we include results and further discussion on partial observability in appendix E.3).

	synthetic dataset		**EEG datset**		**Population spike dataset**
Methods	SMSE	NLL	SMSE	NLL	SMSE
AEA-svGPFA	$\mathbf{0.31 \pm 0.02}$	$\mathbf{1.69 \pm 0.04}$	$\mathbf{0.27 \pm 0.05}$	$\mathbf{1.81 \pm 0.21}$	$\mathbf{0.47 \pm 0.18}$
SGP-VAE [2]	0.36 ± 0.05	1.76 ± 0.07	0.35 ± 0.05	2.17 ± 0.17	0.55 ± 0.19
Vanilla-VAE [21, 42]	1.05 ± 0.02	13.60 ± 3.58	0.57 ± 0.09	3.45 ± 0.87	3.19 ± 1.98
SVAE-LDS [18]	0.933 ± 0.02	9.37 ± 1.94	3.04 ± 0.38	11.71 ± 2.66	2.52 ± 1.31

Table 1: Quantitative comparison of performance between AEA-svGPFA and baseline models on synthetic dataset, EEG dataset and population spiking dataset (Section 5.2) with respect to standarised mean squared error (SMSE) and negative log-likelihood (NLL). All evaluations are averaged over 5 random seeds.

Given the evaluation comparison in Table 1, we see that comparing to all selected baseline models, AEA-svGPFA achieves stronger performance on both the synthetic and EEG dataset, with respect to both the SMSE and NLL metric. Hence the empirical results provide evidence for the advantages of the inclusion of structured amortised inference for capturing the additional posterior correlations that do not exist *a priori*. Moreover, the computational complexity of AEA-svGPFA is on the same order as SGP-VAE, see appendix E.3 for further analysis and discussion.

5.2 Population Neuronal Firing Data

It is well-known that the hippocampal CA1 place cells and medial Entorhinal Cortex (mEC) grid cells are spatially modulated [34, 32], and previous work has also established speed and directional modulation of CA1 and mEC neurons [31, 43, 23]. Moreover, there are often non-trivial correlations between the behavioural covariates. For instance, one might reasonably expect correlation between speed and spatial location as an animal runs fastest in the central area of the environment and slower near the boundaries. Neither svGPFA or SGP-VAE can capture the posterior correlations between latent variables given their mean-field assumptions, whereas AEA-svGPFA, given its free-form variational approximation with full covariance structure, is able to capture arbitrary correlation between the latents whilst utilising the same mutually independent GP priors, and hence is an ideal model for unsupervised learning of the latent structure that correctly reflect the behavioural covariates underlying the neural firing pattern.

We use the single-cell spiking data of neurons in the hippocampal CA1 and mEC regions of rats recorded during exploration in a Z-shaped track reported by Ólafsdóttir et al. [53]. The dataset contains 28 experiment sessions, each spanning 10 minutes. Exemplary neural firing patterns of the population are shown in Figure 11. Note that due to the nature of the task, we can interpret the spatial location of the rat being standard and "unfolded" distances along the track ("unfolded" meaning the

sign of the distance value is dependent on the direction of movement). Other considered behavioural correlates are speed, direction of movement (inbound v.s. outbound), and head-direction.

We again use the same generative model in Eq. 3, but with Poisson observation likelihood.

$$\mathbf{y}(t) = \prod_{d=1}^{D} \text{Poisson}(y_d(t)|\lambda_{\text{NN}}(\mathbf{h})_d) \tag{14}$$

where $\lambda_{\text{NN}}(\cdot)$ is a neural-network mapping from the GPFA features, $\mathbf{h}$, to the rate of Poisson-distributed firing counts for each neuron over a bin-size of $100ms$.

We apply AEA-svGPFA and SGP-VAE to the dataset, In Table 1 we show that AEA-svGPFA quantitatively outperforms SGP-VAE (and other baselines) with respect to SMSE of prediction.

Having trained the model, and given the firing pattern observations, $(x_n, \mathbf{y}_n)$, we could immediately compute the variational posterior latent distribution over the latent processes.

$$q(\mathbf{f}_n) \propto p(\mathbf{f}|\mathbf{U})q(\mathbf{U}) = \mathcal{N}(\mathbf{m}_n^f, \mathbf{S}_n^f), \quad \mathbf{m}_n^f = \mathbf{F}(x_n)\mathbf{m}_{\mathbf{U}}, \quad \mathbf{S}_n^f = \mathbf{K}_n + \mathbf{F}(x_n)(\mathbf{S}_{\mathbf{U}} - \mathbf{K}_{\mathbf{U}})\mathbf{F}(x_n)^T \tag{15}$$

where $q(\mathbf{U})$ is computed as described in Section 3.2 (Eq. 12). Note that due to computational constraints and the stochastic mini-batch nature of the training process, the number of inducing points is much smaller than the total number of observations (64 v.s. 6000 in this case), hence causing temporal chunking artifacts in the posterior estimation. As discussed in Section 3.2, we note that AEA-svGPFA can resolve the temporal chunking with an additional svGP inference step over a much larger time window (i.e., increasing the number of inducing points) with the weights of the recognition network fixed at the final state of the training process (see further discussion in Section E.3.2).

We perform the two-dimensional Canonical Correlation Analysis (CCA) on the variational posterior means and the behavioural correlates over the entire trajectory [15]. Here we choose one session for in-depth investigation and mostly show the qualitative predictions and comparison, more qualitative and quantitative results across sessions can be found in Appendix E.3.2.

The heatmap of the correlation coefficients between the canonical correlates of the posterior means ($CCX\{1, 2\}$) and each of the behavioural correlates is shown in Figure 2a. We observe an overall pattern of high-correlations, indicating the low-dimensional manifold parameterised by the conjunctive set of behavioural correlates can be accurately captured by the latent variables learned with AEA-svGPFA solely from neural spikes.

In order to verify that if the learned latent structure indeed contain decodable information with respect to the behavioural covariates, we extract the direction-modulated neurons predicted by the trained model (details in appendix E.3.2). In Figure 2b, we compare the direction-modulation of the such model-predicted neurons against that of the rest of the neurons, where the direction-modulation is defined as the correlation between raw single-cell spike counts and the direction values. We observe that the model-predicted direction-modulated neurons exhibits significantly greater direction-modulation than the other neurons (p-value=5.56×10^{-4}). For the neuron with the strongest predicted direction-modulation, we show the firing profiles of the neuron when the animal is travelling along both directions, Figure 2b shows that the model-predicted neuron has different firing fields corresponding to different travel direction, hence indeed exhibiting strong direction-modulation. Similar comparison holds for spatial and speed modulations as well, indicating AEA-svGPFA is able to learn the latent structure that coheres well with the ground-truth behavioural covariates, given solely the population spike trains, even for neurons that exhibit conjunctive coding mechanism.

We plot $CCX\{1, 2\}$ against the (x, y) location in Figure 2d, where we observe clear clustered latent structure with respect to the direction of movement, indicating the latent dimensions learned by AEA-svGPFA exhibit disentanglement of direction from spatial location given simple linear projection, despite the conjunctive coding mechanism exhibited in the CA1 neurons. By comparing to the same plot with SGP-VAE, we observe from Figure 13b that the spatial and direction variables are more intertwined, hence qualitatively indicating the superiority of AEA-svGPFA.

Figure 2e shows the temporal evolution of the latent CCs for AEA-svGPFA and SGP-VAE, as well as that of the distance along the track . We observe that $CCX1$ of AEA-svGPFA consistently tracks the dynamics of the distance, whereas neither $CCX1$ or $CCX2$ of SGP-VAE could accurately represents the distance. Moreover, given the additional svGP inference step over the entire trajectory (with

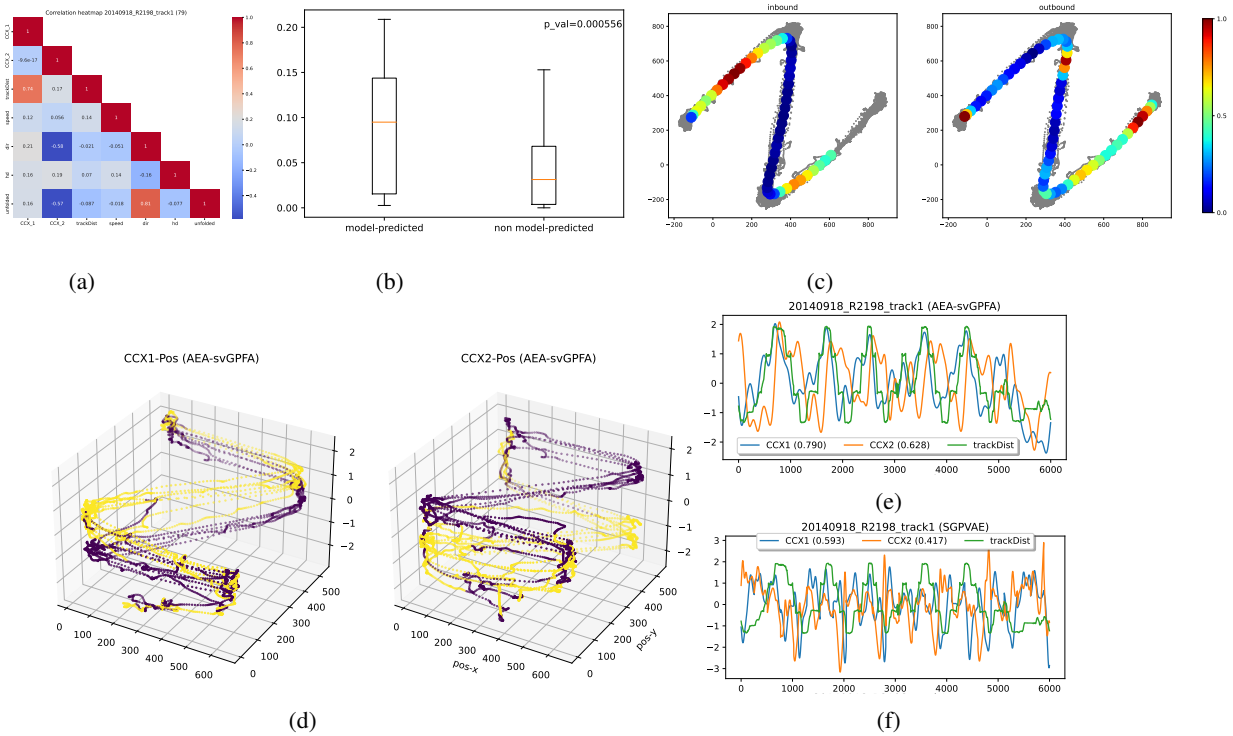

Figure 2: **Analysis of AEA-svGPFA posterior latent structure**. **(a)** Heatmap of the correlation between the CCs and the behavioural correlates for the selected session; **(b)** Directional-modulation comparison between the model-predicted direction-modulated neurons (appendix E.3.2) against other neurons; **(c)** Firing fields of the model-predicted direction-modulated neuron during inbound v.s. outbound movement; **(d)** Plot of $CCX\{1,2\}$ of AEA-svGPFA posterior latents against (x, y) position, color indicates direction of movement (yellow: inbound, magenta: outbound); **(e-f)** Temporal evolution of (standardised values of) $CCX\{1,2\}$ and distance along the track for AEA-svGPFA and SGP-VAE.

changed inducing points, see Section 3.2), $CCX\{1,2\}$ of AEA-svGPFA exhibits much less temporal chunking than that of SGP-VAE.

6 Discussion

In this paper, we proposed structured amortised inference, which can be viewed as a strict generalisation of the SVAE framework [18], where our model allows structured amortised factor potentials to capture the posterior correlations induced by the "explaining away" effect. We theoretically prove that the AEA framework yields an upper bound on the SVAE free energy objective.

We introduce AEA-svGPFA, a combination of the proposed structured amortised inference framework and VAEs with GPFA latents, where scalable free-form variational Gaussian approximation can be achieved. AEA-svGPFA supports a fully non-linear, non-conjugate extension to standard GPFA [51, 1, 12], which is particularly well-suited to settings where posterior inference must be scaled to additional data. We show that AEA-svGPFA outperforms the baseline on all presented tasks, both quantitatively and qualitatively. By applying to population spiking data of hippocampal neurons, we show that AEA-svGPFA is able to identify latent signals that exhibit strong correlation with the behavioural covariates that are well-known to be conjunctively encoded by the recorded neurons. While we have focused on temporal correlation exhibited in time-series data, we note that the sparse amortised variational approximation presented here can also be straightforwardly applied to other GP-latent models, such as GP-LVM and KISS-GP [24, 50, 48]. We note that here we have only looked at single-cell spiking data, which could in principle be combined with the latent processes found by applying our model to EEG or calcium imaging datasets recorded during the same session, possibly with different likelihoods functions. Such joint analysis of latent structures found by analysing different datasets under the same behaviour could potentially discover new findings with respect to, e.g., replay, theta sweeps, etc.

References

[1] Adam, V., Hensman, J., and Sahani, M. Scalable transformed additive signal decomposition by non-conjugate gaussian process inference. In *2016 IEEE 26th international workshop on machine learning for signal processing (MLSP)*, pp. 1–6. IEEE, 2016.

[2] Ashman, M., So, J., Tebbutt, W., Fortuin, V., Pearce, M., and Turner, R. E. Sparse gaussian process variational autoencoders. *arXiv preprint arXiv:2010.10177*, 2020.

[3] Attias, H. Inferring parameters and structure of latent variable models by variational bayes. *arXiv preprint arXiv:1301.6676*, 2013.

[4] Bishop, C. M. *Pattern recognition and machine learning*, volume 4. Springer, 2006.

[5] Blei, D. M. and Jordan, M. I. Variational inference for dirichlet process mixtures. *Bayesian analysis*, 1(1):121–143, 2006.

[6] Campbell, A. and Liò, P. tvgp-vae: Tensor-variate gaussian process prior variational autoencoder. *arXiv preprint arXiv:2006.04788*, 2020.

[7] Casale, F. P., Dalca, A., Saglietti, L., Listgarten, J., and Fusi, N. Gaussian process prior variational autoencoders. *Advances in neural information processing systems*, 31, 2018.

[8] Cunningham, J. P. and Byron, M. Y. Dimensionality reduction for large-scale neural recordings. *Nature neuroscience*, 17(11):1500–1509, 2014.

[9] Dempster, A. P., Laird, N. M., and Rubin, D. B. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1): 1–22, 1977.

[10] Dilokthanakul, N., Mediano, P. A., Garnelo, M., Lee, M. C., Salimbeni, H., Arulkumaran, K., and Shanahan, M. Deep unsupervised clustering with gaussian mixture variational autoencoders. *arXiv preprint arXiv:1611.02648*, 2016.

[11] Dorta, G., Vicente, S., Agapito, L., Campbell, N. D., and Simpson, I. Training vaes under structured residuals. *arXiv preprint arXiv:1804.01050*, 2018.

[12] Duncker, L. and Sahani, M. Temporal alignment and latent gaussian process factor inference in population spike trains. *Advances in Neural Information Processing Systems*, 31, 2018.

[13] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[14] Gregor, K., Danihelka, I., Graves, A., Rezende, D., and Wierstra, D. Draw: A recurrent neural network for image generation. In *International Conference on Machine Learning*, pp. 1462–1471. PMLR, 2015.

[15] Hardoon, D. R., Szedmak, S., and Shawe-Taylor, J. Canonical correlation analysis: An overview with application to learning methods. *Neural computation*, 16(12):2639–2664, 2004.

[16] Hoffman, M. D. and Blei, D. M. Structured stochastic variational inference. In *Artificial Intelligence and Statistics*, pp. 361–369, 2015.

[17] Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.

[18] Johnson, M. J., Duvenaud, D. K., Wiltschko, A., Adams, R. P., and Datta, S. R. Composing graphical models with neural networks for structured representations and fast inference. *Advances in neural information processing systems*, 29:2946–2954, 2016.

[19] Keeley, S., Zoltowski, D., Yu, Y., Smith, S., and Pillow, J. Efficient non-conjugate gaussian process factor models for spike count data using polynomial approximations. In *International Conference on Machine Learning*, pp. 5177–5186. PMLR, 2020.

[20] Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[21] Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[22] Krishnan, R. G., Shalit, U., and Sontag, D. Deep kalman filters. *arXiv preprint arXiv:1511.05121*, 2015.

[23] Kropff, E., Carmichael, J. E., Moser, M.-B., and Moser, E. I. Speed cells in the medial entorhinal cortex. *Nature*, 523(7561):419–424, 2015.

[24] Lawrence, N. Gaussian process latent variable models for visualisation of high dimensional data. *Advances in neural information processing systems*, 16, 2003.

[25] LeCun, Y., Cortes, C., and Burges, C. Mnist handwritten digit database. *ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist*, 2, 2010.

[26] Lee, D. B., Min, D., Lee, S., and Hwang, S. J. Meta-gmvae: Mixture of gaussian vae for unsupervised meta-learning. In *International Conference on Learning Representations*, 2020.

[27] Li, C., Welling, M., Zhu, J., and Zhang, B. Graphical generative adversarial networks. *Advances in neural information processing systems*, 31, 2018.

[28] Li, Y., Turner, R. E., and Liu, Q. Approximate inference with amortised mcmc. *arXiv preprint arXiv:1702.08343*, 2017.

[29] Lin, W., Hubacher, N., and Khan, M. E. Variational message passing with structured inference networks. *arXiv preprint arXiv:1803.05589*, 2018.

[30] Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.

[31] McNaughton, B. L., Barnes, C. A., and O'Keefe, J. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental brain research*, 52(1):41–49, 1983.

[32] Moser, E. I., Kropff, E., and Moser, M.-B. Place cells, grid cells, and the brain's spatial representation system. *Annu. Rev. Neurosci.*, 31:69–89, 2008.

[33] Nieh, E. H., Schottdorf, M., Freeman, N. W., Low, R. J., Lewallen, S., Koay, S. A., Pinto, L., Gauthier, J. L., Brody, C. D., and Tank, D. W. Geometry of abstract learned knowledge in the hippocampus. *Nature*, 595(7865):80–84, 2021.

[34] O'keefe, J. and Nadel, L. *The hippocampus as a cognitive map*. Oxford university press, 1978.

[35] Oord, A. v. d., Kalchbrenner, N., Vinyals, O., Espeholt, L., Graves, A., and Kavukcuoglu, K. Conditional image generation with pixelcnn decoders. *arXiv preprint arXiv:1606.05328*, 2016.

[36] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, 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 Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019. URL `http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf`.

[37] Paulus, M., Choi, D., Tarlow, D., Krause, A., and Maddison, C. J. Gradient estimation with stochastic softmax tricks. *Advances in Neural Information Processing Systems*, 33:5691–5704, 2020.

[38] Pearl, J. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan kaufmann, 1988.

[39] Rainforth, T., Kosiorek, A., Le, T. A., Maddison, C., Igl, M., Wood, F., and Teh, Y. W. Tighter variational bounds are not necessarily better. In *International Conference on Machine Learning*, pp. 4277–4285. PMLR, 2018.

[40] Rasmussen, C. E. and Williams, C. K. I. *Gaussian processes for machine learning.* MIT Press, 2005.

[41] Requeima, J., Tebbutt, W., Bruinsma, W., and Turner, R. E. The gaussian process autoregressive regression model (gpar). In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1860–1869. PMLR, 2019.

[42] Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pp. 1278–1286. PMLR, 2014.

[43] Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B. L., Witter, M. P., Moser, M.-B., and Moser, E. I. Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, 312(5774):758–762, 2006.

[44] Saul, L. and Jordan, M. Exploiting tractable substructures in intractable networks. *Advances in neural information processing systems*, 8, 1995.

[45] Titsias, M. Variational learning of inducing variables in sparse gaussian processes. In *Artificial intelligence and statistics*, pp. 567–574. PMLR, 2009.

[46] Tran, D., Ranganath, R., and Blei, D. M. The variational gaussian process. *arXiv preprint arXiv:1511.06499*, 2015.

[47] Wainwright, M. J. and Jordan, M. I. *Graphical models, exponential families, and variational inference*. Now Publishers Inc, 2008.

[48] Wilson, A. and Nickisch, H. Kernel interpolation for scalable structured gaussian processes (kiss-gp). In *International conference on machine learning*, pp. 1775–1784. PMLR, 2015.

[49] Winn, J., Bishop, C. M., and Jaakkola, T. Variational message passing. *Journal of Machine Learning Research*, 6(4), 2005.

[50] Wu, A., Roy, N. A., Keeley, S., and Pillow, J. W. Gaussian process based nonlinear latent structure discovery in multivariate spike train data. *Advances in neural information processing systems*, 30, 2017.

[51] Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S., Shenoy, K. V., and Sahani, M. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Advances in neural information processing systems*, 21, 2008.

[52] Zhang, X. L., Begleiter, H., Porjesz, B., Wang, W., and Litke, A. Event related potentials during object recognition tasks. *Brain research bulletin*, 38(6):531–538, 1995.

[53] Ólafsdóttir, H. F., Carpenter, F., and Barry, C. LFP and single-unit data published in Olafsdottir, Carpenter and Barry (2016) Nature Neuroscience 19: 792-794, October 2021. URL `https://doi.org/10.5281/zenodo.5566548`.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]