
Decomposing feature-level variation with Covariate Gaussian Process Latent Variable Models

Kaspar Märtens¹ Kieran R Campbell^{2 3 4} Christopher Yau^{5 6}

Abstract

The interpretation of complex high-dimensional data typically requires the use of dimensionality reduction techniques to extract explanatory low-dimensional representations. However, in many real-world problems these representations may not be sufficient to aid interpretation on their own, and it would be desirable to interpret the model in terms of the original features themselves. Our goal is to characterise how feature-level variation depends on latent low-dimensional representations, external covariates, and non-linear interactions between the two. In this paper, we propose to achieve this through a structured kernel decomposition in a hybrid Gaussian Process model which we call the Covariate Gaussian Process Latent Variable Model (c-GPLVM). We demonstrate the utility of our model on simulated examples and applications in disease progression modelling from high-dimensional gene expression data in the presence of additional phenotypes. In each setting we show how the c-GPLVM can extract low-dimensional structures from high-dimensional data sets whilst allowing a breakdown of feature-level variability that is not present in other commonly used dimensionality reduction approaches.

1. Introduction

The identification of low-dimensional structure is crucial to gaining insight from complex high-dimensional data. In probabilistic models, this task can be formulated as finding a low-dimensional latent variable $\mathbf{z}_n \in \mathbb{R}^Q$ for each data

point $n = 1, \dots, N$ and a set of mappings $f^{(j)} : \mathbf{z} \mapsto \mathbf{y}^{(j)}$ for every feature $\mathbf{y}^{(j)}$, $j \in \{1, \dots, P\}$ so that $Q \ll P$. Obtaining an informative one- or two-dimensional representation of the data is often particularly desirable, allowing us to visually interpret the patterns and relationships present in the data.

There are a number of approaches for defining the mapping functions, ranging from linear models such as the pPCA (Tipping & Bishop, 1999) to non-linear models provided by neural networks such as the Variational Autoencoder (VAE) (Kingma & Welling, 2014). The former are classically well-studied, their behaviours widely explored and can be considered to be “interpretable”, but the linearity assumptions can often be too restrictive. Non-linear techniques offer more flexibility with the disadvantage that the non-trivial dependencies that are learnt can be challenging to characterise, leading to their popular description as “black box techniques”.

While the derivation of low-dimensional representations represents a powerful means of information extraction from high-dimensional data, they can be of limited utility without an explicit reference back to the original observed features. For example, in transcriptomics, latent representations of gene expression profiles can help to identify sub-populations of biological samples with similar coordinated gene behaviour. However, as the underlying biology is ultimately physically driven by variation at the level of individual genes, we would like to decompose that expression variability into a number of meaningful sub-components, in particular, as a function of latent coordinates and other observed covariates.

In this paper, we focus on dimensionality reduction in the presence of additional covariate information. This is often available in real-life applications, for example, in transcriptomics, covariate information might include categorical labels (e.g. denoting known disease sub-populations), continuous-valued measurements (e.g. biomarkers), or censored information (e.g. patient survival times). More formally, suppose that in addition to $(\mathbf{z}_n, \mathbf{y}_n)$, each data point is also associated with a C -dimensional covariate vector, $\mathbf{x}_n \in \mathbb{R}^C$ which may act to modulate the variation in the data. Our goal is not black-box predictive modelling, i.e.

¹Department of Statistics, University of Oxford, UK

²Department of Statistics, University of British Columbia, Canada

³BC Cancer Agency, Canada ⁴UBC Data Science Institute, Canada

⁵The Alan Turing Institute, UK ⁶Institute of Cancer and Genomic Sciences, University of Birmingham, UK. Correspondence to: Christopher Yau <c.yau@bham.ac.uk>.

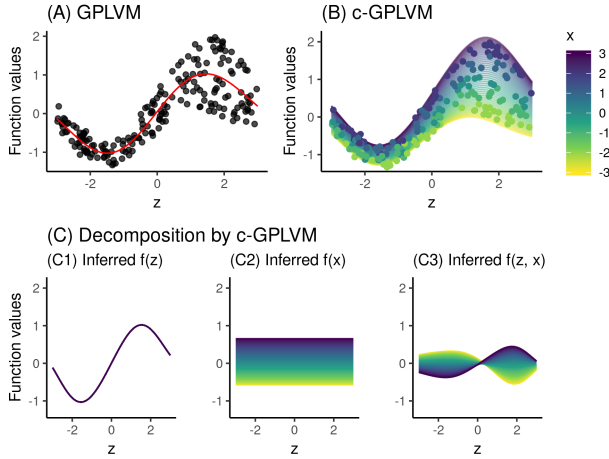


Figure 1: The presence of covariate effects confounds the mappings from the underlying latent z (on x -axis) to the feature space (selected feature on y -axis). Standard GPLVM (panel A)) ignores covariate information which is denoted by colour, whereas its inclusion in c-GPLVM (panel B)) allows us to capture a variety of nonlinear covariate effects. Furthermore, we would like to decompose the feature-level variation into covariate x and latent z additive contributions and their non-linear interaction as shown in panels (C1-C3).

we are not interested in simply learning conditional distributions $p(y|x)$ or $p(x|y)$. Instead, our work focuses on discovery applications. Our goal is two-fold:

- Learn a low-dimensional z that is covariate-adjusted.
- Characterise feature-level variation, separating what is explained by z from the contribution of x .

Thus we are interested in inferring both the posterior $p(z|y, x)$ over latent coordinates as well as the mappings $p(f^{(j)}|y, x, z)$ for every feature j . We wish to understand the nature of the feature-level variability and e.g. allow us to identify sets of features which are fully explained by covariates versus those which show complex dependence on both latent variables and covariates. We will principally do this in a high-dimensional data setting where the number of features is vastly greater than the number of covariates.

We particularly focus on the use of Gaussian Processes (GP) as a non-parametric model over the mapping functions. Our choice reflects the strong theoretical underpinnings of GP models as well as recent advances that have enabled such models to be scalable to large data sets (Hensman et al., 2013; 2015). In the context of dimensionality reduction, the Gaussian Process Latent Variable Model (GPLVM) (Lawrence, 2005) is the reference for probabilistic non-linear dimensionality reduction which has spawned numerous extensions, e.g. (Shon et al., 2006; Ek & Lawrence, 2009; Gao et al., 2011; Jiang et al., 2012; Damianou & Lawrence, 2013; Gadd et al., 2018).




Figure 2: Graphical models for (a) GPLVM, (b) a particular implementation of supervised-GPLVM, and (c) c-GPLVM.

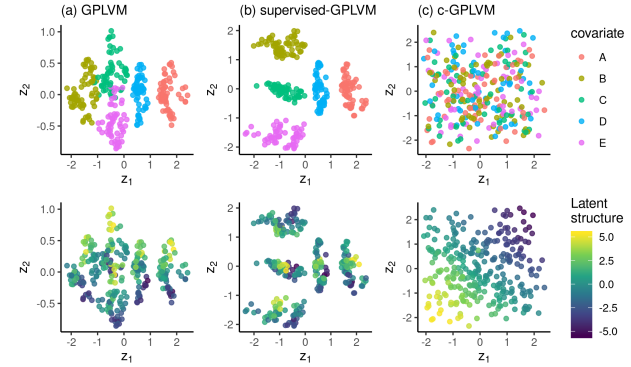


Figure 3: Here we use GPLVM, supervised-GPLVM and c-GPLVM to learn a 2D latent space (z) from a synthetic data set which contains 5 discrete classes A-E (x). Here the data generative mechanism contains both the class effects (covariate x) as well as the shared 2D latent structure (z). For all three methods, their learned 2D latent space z has been shown twice, coloured by x (top panel) and the true underlying z (bottom panel). The dominant effect of the 5 classes means a GPLVM will learn a latent space that reflects the presence of these groups. This is further exaggerated with the supervised-GPLVM which acts to increase the separation between the classes and is most useful if class discrimination is the ultimate objective. In contrast, the c-GPLVM seeks to find the common shared structure between the 5 classes, and infers a latent space which *adjusts* for the presence of the 5 classes (in linear modelling this process would be analogous to “regressing out” the class-specific effects).

Figure 1(A,B) illustrates our setting of interest where feature values vary not only over the latent coordinate, but also over a single, continuous covariate. This dependence on the covariate may confound the mappings from z to $y^{(j)}$ when applying a GPLVM. As a result, it will fail to account for the underlying latent structure in the data that is shared across all covariate values. Covariate effects can be incorporated in various ways (e.g. the supervised-GPLVM in Figure 2).

Our proposal, which we call the *covariate-GPLVM* (c-GPLVM), specifically focuses on learning a set of mappings $f^{(j)} : (z, x) \mapsto y^{(j)}$ which are defined on the joint space of z and x . This can be seen as a hybrid between the GP regression and the GP latent variable model where the input space consists of two parts: fixed covariates x and

unobserved latent coordinates \mathbf{z} . Furthermore, we embed a structured sparsity-inducing kernel decomposition which allows c-GPLVM to explicitly disentangle variation in the observed data vectors induced by the covariates and/or latent variables, and the *interaction* between the two (Figure 1C). The novelty of our approach is that the structured kernel permits both the development of a nonlinear mapping into a latent space where confounding factors are already adjusted for and feature-level variation that can be deconstructed. This construction leads to inferring a *covariate-adjusted* latent space as illustrated in Figure 3.

We demonstrate the utility of c-GPLVM on a number of simulated examples and applications in disease progression modelling from high-dimensional gene expression data in the presence of additional phenotypes. In each setting we show that the c-GPLVM is able to effectively extract low-dimensional structures from high-dimensional data sets whilst allowing a breakdown of feature-level variability that is not present in other commonly used dimensionality reduction approaches.

2. Background

2.1. Gaussian Process regression

Gaussian processes offer a principled non-parametric framework for inference over functions (Rasmussen & Williams, 2006). Consider a real valued function defined on the D -dimensional inputs $\mathbf{X} := (\mathbf{x}_1, \dots, \mathbf{x}_N)$. A function f is said to be drawn from a GP with mean $\mathbf{0}$ and covariance $k(\mathbf{x}, \mathbf{x}')$, denoted by $f(\mathbf{x}) \sim \mathcal{GP}(\mathbf{0}, k(\mathbf{x}, \mathbf{x}'))$, when

$$p(\mathbf{f}|\mathbf{X}, \theta) = \mathcal{N}(\mathbf{f}|\mathbf{0}, \mathbf{K})$$

where \mathbf{K} is the kernel matrix of all pairs of inputs which also depends on hyperparameters θ , with elements $\mathbf{K}_{ij} := k(\mathbf{x}_i, \mathbf{x}_j)$. One popular choice for $k(\cdot)$ is the squared exponential ARD kernel,

$$k(\mathbf{x}_i, \mathbf{x}_k) = \sigma^2 \exp \left(-\frac{1}{2} \sum_{j=1}^D \frac{(x_{ij} - x_{kj})^2}{l_j^2} \right)$$

where σ^2 is the kernel variance parameter and l_j are the feature-specific lengthscales. Combining the GP prior with a given likelihood $p(\mathbf{y}|\mathbf{f})$ can lead to a variety of models, e.g. the GP regression model in case of the Gaussian likelihood.

2.2. Gaussian Process Latent Variable Model

GPLVM is a latent variable model which uses GPs as latent mappings, being a non-linear extension of the probabilistic PCA (Lawrence, 2005). Suppose we have an observed data matrix \mathbf{Y} , consisting of P features ($\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(P)}$) and N data points, and our goal is to learn a low-dimensional representation $\mathbf{Z} := (\mathbf{z}_1, \dots, \mathbf{z}_N)$. We place a prior over these latent variables, $p(\mathbf{Z}) = \prod_{i=1}^N \mathcal{N}(\mathbf{z}_i|\mathbf{0}, \mathbf{I})$ and our aim

will be to infer the respective posterior. Now conditional on these latent inputs, the GPLVM is essentially a multi-output GP regression model as it specifies a GP prior for every latent mapping $\mathbf{f}^{(j)}$ and some likelihood $p(\mathbf{y}^{(j)}|\mathbf{f}^{(j)})$ for $j \in \{1, \dots, P\}$. Denoting the collection of GP function values $\mathbf{F} := (\mathbf{f}^{(1)}, \dots, \mathbf{f}^{(P)})$, the GPLVM is formulated as the following generative model

$$p(\mathbf{Y}, \mathbf{F}, \mathbf{Z}, \theta) = p(\mathbf{Y}|\mathbf{F})p(\mathbf{F}|\mathbf{Z}, \theta)p(\mathbf{Z})$$

In the special case when the emission likelihood is Gaussian, i.e. when $p(\mathbf{y}^{(j)}|\mathbf{f}^{(j)}) = \mathcal{N}(\mathbf{y}^{(j)}|\mathbf{f}^{(j)}, \sigma^2\mathbf{I})$, one can analytically integrate out the GP mappings, resulting in the marginal likelihood

$$p(\mathbf{Y}|\mathbf{Z}, \theta) = \prod_{j=1}^P \mathcal{N}(\mathbf{y}^{(j)} | \mathbf{0}, \mathbf{K}_{zz}^{(j)} + \sigma_j^2\mathbf{I}) \quad (1)$$

In this special case, we do not need to infer \mathbf{F} explicitly and now inference needs to be carried out on latent variables $\mathbf{z}_1, \dots, \mathbf{z}_N$ and kernel hyperparameters θ only.

3. Covariate-GPLVM

We now consider our extension of the GPLVM to include fixed inputs where we are specifically interested in understanding the *interaction* between the measured covariates \mathbf{x} and latent variables \mathbf{z} .

3.1. Model

Specifically, we aim to learn mappings which are defined on the joint space of \mathbf{z} and \mathbf{x} , i.e. $f^{(j)} : (\mathbf{z}, \mathbf{x}) \mapsto \mathbf{y}^{(j)}$. As discussed earlier, this is crucial for letting us to learn a *covariate-adjusted* representation \mathbf{z} (as opposed to for example modelling both \mathbf{x} and \mathbf{y} as a function of \mathbf{z}).

Different assumptions about the form of interaction between \mathbf{z} and \mathbf{x} can be made, and these can be encoded via different kernel structures. One approach would be to define the ARD kernel on this joint space:

$$k^{\text{int}}((\mathbf{x}, \mathbf{z}), (\mathbf{x}', \mathbf{z}')) :=$$

$$\sigma_{xz}^2 \exp \left[-\frac{1}{2} \sum_{j=1}^C \left(\frac{x_{ij} - x_{kj}}{l_j^{(x)}} \right)^2 - \frac{1}{2} \sum_{j=1}^Q \left(\frac{z_{ij} - z_{kj}}{l_j^{(z)}} \right)^2 \right].$$

Alternatively, one could define an additive kernel

$$k^{\text{add}}((\mathbf{x}, \mathbf{z}), (\mathbf{x}', \mathbf{z}')) := k^x(\mathbf{x}, \mathbf{x}') + k^z(\mathbf{z}, \mathbf{z}')$$

where both k^x and k^z are ARD kernels.

Using either an additive (ADD) model or an interaction (INT) model correspond to different modelling assumptions. We have illustrated the implications of this on a synthetic example in Figure 4. When trained on data with covariate values $\mathbf{x}_i \in \{-1, 0, 1\}$ and the inferred \mathbf{z}_i values cover the entire range from -3 to 3. However, training data with covariate $\mathbf{x} = 1$ is not well represented (e.g. it could be scarce




Figure 4: ADD and INT kernels have different behaviour when prediction requires extrapolation from training data. Observed data is shown as dots (\mathbf{z} on x -axis, selected feature on y -axis, coloured by covariate \mathbf{x}). Predictions by the c-GPLVM together with uncertainty for $\mathbf{x} \in \{-1, 0, 1\}$.

in practice), and thus prediction for future observations with $\mathbf{x} = 1$ may require extrapolation. Note that even though the ADD model is more restrictive, it is more data-efficient when the feature dependence w.r.t. to z is stationary, and can extrapolate well in scenarios when the non-stationarity of the INT model provides too much uncertainty.

Ideally we would like to use the simpler ADD model when it is appropriate, and capture everything that remains unexplained via the INT component. However, this is a challenging task. One might try to tackle this problem via model selection, similarly to how the Automatic Statistician (Duvinaud et al., 2013) carries out search for optimal combinations of kernels. But note that in the GPLVM setting, we would have to select between ADD and INT kernels for every feature, resulting in 2^P possible combinations which will be computationally infeasible for a high-dimensional dataset.

Instead, we propose an ADD+INT decomposition which will not require model search. Not only will it detect the presence or absence of additive and interaction effects (resulting in behaviour as shown in Figure 4 panel ADD+INT), but it will also aid interpretation via *decomposing* feature-level variation into ADD and INT effects.

3.2. ADD+INT decomposition

In classical statistics, the Latent Variable Regression (LVR) model could be used to infer the following *linear* decomposition,

$$y_i^{(j)} = \mu_j + \sum_k \alpha_k^{(j)} z_{ik} + \sum_l \beta_l^{(j)} x_{il} + \sum_{k,l} \gamma_{kl}^{(j)} z_{ik} x_{il} + \varepsilon_{ij}$$

This decomposition provides insights into how much of the variation is explained by individual inputs \mathbf{z}_k and \mathbf{x}_l , and how much is explained by their interactions.

We are interested in learning an analogous decomposition

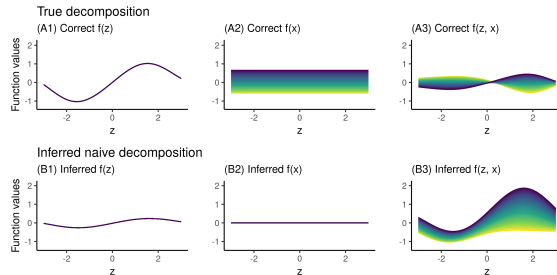


Figure 5: The naive decomposition (2) is unidentifiable. The true generative mechanism (top row) and an example of an inferred naive decomposition (bottom row).

for GPs, i.e. we would like to decompose the mapping $f^{(j)}$ as follows

$$f^{(j)}(\mathbf{z}, \mathbf{x}) = f_0^{(j)} + f_z^{(j)}(\mathbf{z}) + f_x^{(j)}(\mathbf{x}) + f_{zx}^{(j)}(\mathbf{z}, \mathbf{x}) \quad (2)$$

where $f_z^{(j)}(\mathbf{z})$ and $f_x^{(j)}(\mathbf{x})$ would capture marginal \mathbf{z} and \mathbf{x} effects, and the interaction $f_{zx}^{(j)}(\mathbf{z}, \mathbf{x})$ would capture everything that remains *unexplained* by the two.

3.2.1. UNIDENTIFIABILITY

Below we drop index j to reduce notational clutter.

A standard approach would be to place independent GP priors over f_0, f_z, f_x, f_{zx} . However this results in an unidentifiable decomposition. This is because f_0, f_z, f_x can all be seen as bivariate functions which have positive probability under the GP prior $f_{zx} \sim \mathcal{GP}(0, k(\cdot))$. Thus the set of functions f_0 is a subset of functions f_z (and f_x). Similarly the latter are a subset of functions f_{zx} . In other words, the supports of the four independent GP priors overlap.

This unidentifiability is illustrated in Figure 5, where on synthetic data we have shown an alternate realisation of the inferred decomposition. This alternate decomposition (in bottom row) would suggest that there is no additive \mathbf{x} effect and that almost all of the variation is explained by the interaction effect, which was not the generative mechanism used to synthesise the data (shown in top row).

3.2.2. ZERO-MEAN FUNCTIONAL CONSTRAINTS

To turn this into an identifiable task, we first note that for a one-dimensional $g \sim \mathcal{GP}(0, k(\cdot))$, the distribution of g conditional on $\int_a^b g(t) dt = 0$ is still a GP, but with a modified kernel (Durrande et al., 2012). For some kernels such as the squared exponential, this can be calculated in closed form, letting us easily construct a mean-zero GP prior on some interval $[a, b]$.

Furthermore this construction can be extended to a mean-zero interaction (Durrande et al., 2013). To uniquely define f_{zx} , we need to fix its marginal distributions. A natural way

would be to enforce zero-marginals $\int_a^b f_{zx}(z, x) dz = 0$ and $\int_a^b f_{zx}(z, x) dx = 0$. Similarly, we uniquely define f_z and f_x by enforcing $\int_a^b f_z(z) dz = 0$, $\int_a^b f_x(x) dx = 0$. Now the decomposition (2) has become identifiable.

This construction can be seen as a joint prior over the four functional subspaces f_0, f_z, f_x, f_{zx} which has the following two properties: First, their supports are non-overlapping. Second, the respective functional subspaces are orthogonal in L^2 . The former is sufficient for a unique decomposition, but the latter makes interpretation easier, because

$$\text{Var}(f) = \text{Var}(f_0) + \text{Var}(f_z) + \text{Var}(f_x) + \text{Var}(f_{zx})$$

due to orthogonality.

The corresponding kernel decomposition is

$$k^{\text{add+int}}((\mathbf{z}, \mathbf{x}), (\mathbf{z}', \mathbf{x}')) := \sigma_b^2 + \sigma_z^2 \tilde{k}(\mathbf{z}, \mathbf{z}') + \sigma_x^2 \tilde{k}(\mathbf{x}, \mathbf{x}') + \sigma_{zx}^2 \tilde{k}(\mathbf{z}, \mathbf{z}') \tilde{k}(\mathbf{x}, \mathbf{x}')$$

where $\tilde{k}(\cdot)$ is the mean-zero squared exponential kernel (details given in Supplementary).

Finally, we make use of Bayesian shrinkage priors on the kernel variances $\sigma_z^2, \sigma_x^2, \sigma_{zx}^2$ to encourage unnecessary components to be shrunk to zero. For this purpose, we specify $\sigma_z^2, \sigma_x^2, \sigma_{zx}^2 \sim \Gamma(1, 1)$.

3.3. Inference for c-GPLVM

The likelihood for c-GPLVM has a similar form to (1), now containing both fixed \mathbf{X} and latent \mathbf{Z} inputs,

$$p(\mathbf{Y}|\mathbf{Z}, \mathbf{X}) = \prod_{j=1}^P \mathcal{N}(\mathbf{y}^{(j)}|\mathbf{0}, \mathbf{K}_{zz}^{(j)} + \mathbf{K}_{xx}^{(j)} + \mathbf{K}_{(zx)(zx)}^{(j)} + \sigma_j^2 \mathbf{I})$$

Variational inducing point based inference (Titsias, 2009; Damianou et al., 2016) can be adopted for c-GPLVM in a straightforward way. As the kernels in c-GPLVM are defined on the extended (product) space of \mathbf{x} and \mathbf{z} , so the inducing points now lie in this space which has dimensionality $\dim(\mathbf{x}) + \dim(\mathbf{z})$. Under certain modelling assumptions this dimensionality may be reduced (e.g. under the ADD kernel we simply need inducing points in the \mathbf{x} and \mathbf{z} space separately and this can reduce the computational cost).

3.4. Inference for censored covariates

Next, we discuss inference in the scenario when some of the inputs \mathbf{x} have not been observed, but instead have been censored. In practice, this happens when we want to use patient survival times in the role of covariate, but for many individuals we only have access to their last follow-up time. We treat this as a lower bound on their true survival. In this section, we assume \mathbf{x}_i are one-dimensional and represent survival times.

Suppose for some individuals i we have observed their true

survival \mathbf{x}_i , whereas for the rest, we will treat \mathbf{x}_i as latent. For the latter, we additionally have observed a lower bound a_i which we want to incorporate into our model as this information is informative of the posterior of \mathbf{x}_i .

Following a common assumption in survival analysis, we assume *a priori* Weibull-distributed patient survival, i.e. for all individuals i we specify a prior $p(x_i) = \text{Weibull}(k_0, \lambda_0)$ with hyperparameters k_0, λ_0 . Next, we incorporate censoring information by conditioning on survival x_i exceeding a_i , and additionally upper bounding by the assumed maximum lifespan b_i (if one does not want to make such an assumption, b_i can be set to ∞). The conditional prior takes the form $p(x_i|a_i, b_i) = \text{TruncWeibull}_{[a_i, b_i]}(k_0, \lambda_0)$ which has a closed form density.

Next we construct a variational inference scheme which will allow us to fit c-GPLVM in the presence of censored covariates as well as to infer the posteriors over true survival times. We need to choose the family of approximating distributions $q(\mathbf{x}_i)$ such that they would have the same support as $p(\mathbf{x}_i)$. We choose it to be the truncated Gaussian distribution which has been restricted to the interval $[a_i, b_i]$, i.e. we choose $q(\mathbf{x}_i) = \mathcal{TN}_{[a_i, b_i]}(\mu_i, \sigma_i^2)$ where μ_i, σ_i^2 are variational parameters. In addition to having appropriate support, this choice lets us perform efficient reparameterisation-based inference (Kingma & Welling, 2014; Rezende et al., 2014).

This lets us sample from $q(\mathbf{x}_i)$ via reparameterisation:

$$\begin{aligned} F(a_i) &:= \Phi((a_i - \mu_i)/\sigma_i) \\ F(b_i) &:= \Phi((b_i - \mu_i)/\sigma_i) \\ \varepsilon &:= F(a) + u(F(b) - F(a)) \text{ with } u \sim U(0, 1) \\ x_i &\sim \Phi^{-1}(\varepsilon)\sigma_i + \mu_i \end{aligned}$$

where $\Phi(\cdot)$ is the cdf of the standard normal.

Denoting the censored set of survival times by \mathbf{X}^{cens} and fully observed set by \mathbf{X}^{obs} , we can fit c-GPLVM by optimising the following ELBO:

$$\begin{aligned} \mathcal{L} &= \mathbb{E}_{q(\mathbf{X}^{\text{cens}})} \log p(\mathbf{Y}|\mathbf{Z}, \mathbf{X}^{\text{obs}}, \mathbf{X}^{\text{cens}}) + \\ &\quad - \sum_{i=1}^N \text{KL}(q(\mathbf{z}_i)||p(\mathbf{z}_i)) - \sum_{i \in \text{cens}} \text{KL}(q(\mathbf{x}_i)||p(\mathbf{x}_i)). \end{aligned}$$

4. Related Work

A number of alternative approaches have been developed which can be said to lie on the spectrum between GPR and GPLVMs.

Latent GP regression models (LGPR) (Wang & Neal, 2012; Bodin et al., 2017) extend the input space of a GP regression model with additional latent variables that are used to modulate the covariance function. LGPR models often implicitly assume that the dimensionality of the latent variables is less than the number of regressors $\dim(\mathbf{x}_n) \gg \dim(\mathbf{z}_n)$, in

this sense being closer to a GP regression model compared to a GPLVM. Approximate marginalisation of the latent variables allows these models to consider non-stationary, multimodal behaviour, and e.g. has recently been used for conditional density estimation (Dutordoir et al., 2018). In these hybrid GP models, the inputs \mathbf{z} and \mathbf{x} are usually concatenated, resulting in the INT kernel (as in our section 3.1) suited for predictive applications.

Semi-described and semi-supervised learning in GP regression is considered by (Damianou & Lawrence, 2015) when the inputs (or outputs) are partially observed or uncertain, this formulation also leads to a hybrid model with latent variables introduced in place of missing data.

Supervised GPLVM encompass a family of related ideas seek to decompose the joint distribution $p(\mathbf{y}, \mathbf{x}, \mathbf{z})$ in different ways. For example, in (Gadd et al., 2018), the joint distribution is factorised such that the latent variables are conditionally dependent on the covariates $p(\mathbf{z}|\mathbf{x})$. Whilst in the supervised-GPLVM formulation (Gao et al., 2011) and the shared-GPLVM (Shon et al., 2006) both observations and fixed inputs are conditionally dependent on the latent inputs, $p(\mathbf{y}|\mathbf{z})p(\mathbf{x}|\mathbf{z})$. In the discriminative-GPLVM (Urtasun & Darrell, 2007), class label information is used to maximise the discriminatory power between discrete classes in the latent space. This is fundamentally different from c-GPLVM which learns a *covariate-adjusted* \mathbf{z} .

Similar decompositions as (2) were also discussed in (Duveaud et al., 2011), but as we discussed above, such a kernel decomposition is unidentifiable. This is not problematic for predictive applications they consider, but it is crucial for interpretability.

5. Experiments

Our implementation of c-GPLVM is available in <https://github.com/kasparmartens/c-GPLVM>.

5.1. Toy examples

First, we consider synthetic two-dimensional datasets as shown as displayed in Figure 6 (“rings” and “pinwheel”). In the observation space, each circle (panel A) or spoke (panel B) corresponds to a *trajectory* and we would like to understand the shared properties of how features vary *along* each trajectory irrespective of the location of each circle or the *angle* of each spoke. If we had additional, *covariate* information on the clustering structure in the form of this angle (shown by colour coding in both panels), we could use this information *jointly* to achieve our goal. Note that the “rings” data has an additive data generating mechanism, whereas the “pinwheel” features exhibit a complex non-linear interaction between the angle of the spoke and the position along the trajectory.




Figure 6: Observed 2D toy data with additional covariate information shown by colour.




Figure 7: Comparison of models on the “rings” data. Top row: data coloured according to the inferred one-dimensional \mathbf{z} values. Bottom row: Mean of the posterior predictive in the 2D observation space when varying $\mathbf{z} \in [-2, 2]$ and fixing \mathbf{x} to the five observed values.




Figure 8: Comparison of models on the “pinwheel” data.

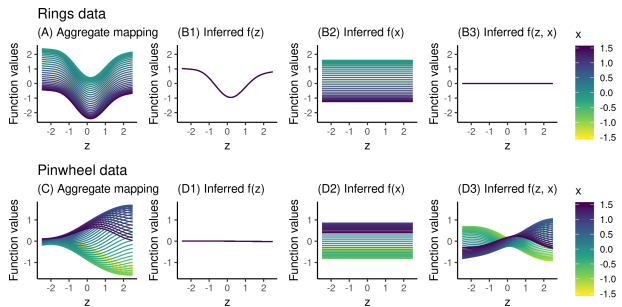


Figure 9: The ADD+INT decomposition inferred by c-GPLVM on the “rings” data (top row) and the “pinwheel data” (bottom row). The first panels (A and C) show the inferred aggregate function, remaining panels (B1-B3, D1-D3) show the breakdown into three components.

The inferred one-dimensional z values are shown by colour coding in the top row of Figures 7 and 8 (panels A1-D1 comparing four models). As expected, the z inferred by GPLVM captures global structure (ignoring covariate information) and z learned by supervised-GPLVM discriminates between the groups defined by the covariate. The Latent Variable Regression, which adjusts for the covariate in a linear model, struggles due to the highly non-linear signal present in the data.

To gain a better insight into what the models have *learned*, we have also visualised the mean of the posterior predictive distributions in the data space (bottom row of Figures 7 and 8, same models as in the top row). Here we have varied z from -2 to 2 and kept fixed x to its observed five values. We see that c-GPLVM is the only model which has identified the one-dimensional structure along the circular trajectories for the “rings” data and along the spokes for the “pinwheel” data.

Furthermore, we have visualised the ADD+INT decomposition inferred by c-GPLVM. This has been shown for the first coordinate y_1 of both data sets in Figure 9, showing both the inferred aggregate mapping (panels A and C) as well as the individual components (rest of panels). We show here the range of GP mappings obtained by varying x over a fine grid $x \in [-1.5, 1.5]$. For the “rings” data, c-GPLVM has inferred the presence of additive effects for both z and x , but no interaction whereas for the “pinwheel” data, all three components contribute.

The above “rings” and “pinwheel” exhibit highly non-linear patterns. In Supplementary we explore the behaviour of these models on more realistic synthetic examples, considering dependency structures that we expect to see in real gene expression data.

5.2. Survival toy example

To demonstrate the utility of our inference scheme for censored covariates, we carried out a synthetic experiment. Having generated data as a function of known z and x (details in Supplementary), we wanted to explore the behaviour of our approximate posterior when the true $x = 1.5$, but we vary its lower bound on a grid from 0.7 to 1.7. The resulting approximate posteriors have been shown in Figure 10. The two panels (a) and (b) represent two observations which exhibit a different degree of uncertainty (in (b) data is more informative of survival than in (a)). Note that when the lower bound becomes larger than the true survival, the posterior starts to concentrate towards the lower bound.

5.3. Real-world data experiment

We now consider a real world data set consisting of $N = 770$ breast cancers from The Cancer Genome Atlas cohort




Figure 10: Posterior inference for censored survival times on synthetic data. For two individuals (in panels (a) and (b), demonstrating two different degrees of uncertainty), we vary the lower bound from 0.7 to 1.7 (ten censoring scenarios on y -axis), lower bound value denoted by small vertical dashed lines. The true x is always 1.5 (vertical red line).

(Weinstein et al., 2013). We use gene expression measurements across 500 highly variable genes (Y) and survival information for each patient as a covariate (x) to explore the relationship between gene expression changes, survival and a unidimensional latent representation z .

5.3.1. PREDICTION OF SURVIVAL TIMES

To ensure that our modelling assumptions are aligned with the observed data, we first investigate the survival c-GPLVM in a controlled setting. Specifically we focus on its predictive ability. That is, we consider a subset of individuals whose death times have been observed ($N = 151$) and carry out artificial censoring in batches of size 5. For 5 patients at a time, we artificially censor their survival time by half a year, and fit c-GPLVM to infer the posterior of the true survivals. When considering alternative methods for the prediction of survival times, we note that most models in the field of survival analysis are non- or semi-parametric, i.e. they do not model the baseline hazard function, and thus do not provide a straightforward way for prediction. To obtain a parametric survival model, one common choice is to use the Weibull distribution. Thus, to compare c-GPLVM with a baseline method, we used the Weibull regression model with shrinkage priors as described in (Peltola et al., 2014), predicting patient survival as a function of the gene expression matrix Y . Predictions by c-GPLVM have been shown against the true values in Figure 11, detailed comparison with Weibull regression is shown in Supplementary (on repeated sampling from the posterior, the average MSE from the c-GPLVM is 9.2 whereas for the Weibull model it is 17.1). We also note that the c-GPLVM *knows when it does not know*, i.e. the mis-predictions typically have high uncertainty.

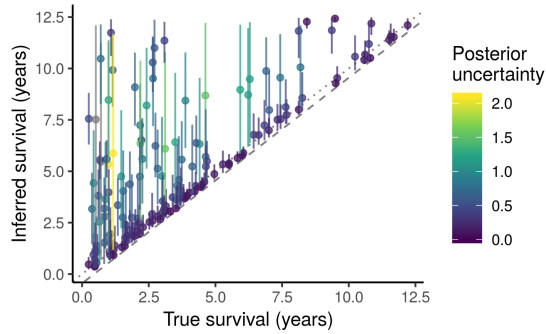


Figure 11: Predicted survival times for artificially censored individuals in TCGA data using c-GPLVM (true value on x -axis, inferred posterior mean and (5%, 95%) quantiles on y -axis, colour denotes posterior uncertainty).

5.3.2. SURVIVAL-ADJUSTED CANCER MODELLING

Next, we returned to the original data set of all 770 breast cancers and fitted standard GPLVM and c-GPLVM to the entire cohort. Fig 12a compares a feature-level fit for standard GPLVM and c-GPLVM for three genes *TSPAN1*, *KRT23* and *LPL*. First, *TSPAN1* is a gene that codes for a member of the protein family, Tetraspanins, also known as the transmembrane 4 superfamily. These are small transmembrane glycoproteins which were first described in studies of tumour associated proteins and has been reported to regulate cancer progression in many human cancers (Munkley et al., 2017). In line with this, c-GPLVM identifies *TSPAN1* as a gene whose expression increases along the latent coordinate and not with survival suggesting that the identified latent dimension could correspond to disease progression. In contrast, Keratin-23 (*KRT23*) decreases with z but is also additively modulated by the survival covariate with patients who survive longer seemingly retaining a higher expression of *KRT23*. Lipoprotein lipase (*LPL*) plays a role in breaking down fat in the form of triglycerides, which are carried from various organs to the blood by molecules called lipoproteins. This gene is identified as having interaction effects and we see that there is a range of latent input values $-2 < z < 0$ where low *LPL* expression is associated with longer survival and high *LPL* expression is associated with reduced lifespan. This suggests that during certain periods of breast cancer development, high levels of *LPL* activity could enable aggressive disease progression possibly through acting to provide a supply of fatty acids to fuel tumour growth (Kuemmerle et al., 2011). Our analysis here is only illustrative but serves to demonstrate that the kernel decomposition within c-GPLVM allowed us to “discover” three distinct gene behaviours. This could not have been readily achieved without further post-processing analysis or modification of other existing GPLVM extensions and implementations.

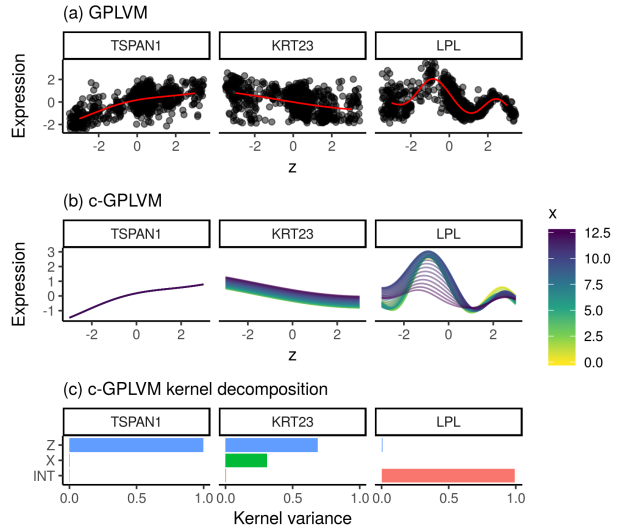


Figure 12: GPLVM (panel A) does not capture the variation driven by the censored covariate (x =survival), whereas c-GPLVM (panel B) not only captures the effects explained by survival, but in addition, the **add+int** kernel decomposition (panel C) lets us identify three sets of genes: those without any covariate effect (e.g. *TSPAN1*), genes with an additive covariate effect (e.g. *KRT23*), and genes that exhibit a complex non-linear interaction (e.g. *LPL*).

6. Discussion

We have introduced the covariate-GPLVM that integrates GP regression and latent variable modelling with a specific focus on the interaction between covariates and latent variables. The c-GPLVM learns latent spaces that reveal structure in the data which is shared across covariate values. By making use of GP mappings that are specified on the extended joint space of covariates and latent variables, we can model complex non-linear dependencies, while maintaining interpretable and decomposable mappings. The c-GPLVM is applicable for a wide range of applications where there is known structure accounting for variance attributable to this structure. This encompasses scenarios when we are interested in explicitly exploring the interactions between this covariate information and other features, as well as those where such covariates are “nuisance” variables and we would like to adjust for confounding factors (e.g. batch effects or the presence of different ancestral populations in population genetics). A natural extension of our work is to consider a deep multi-output Gaussian Process formulation (Damianou & Lawrence, 2013) in which multiple output dimensions can be coupled via shared Gaussian process mappings. However, this addition introduces many design choices with regards to how this shared structure could be embedded within our structured kernel decomposition framework. We will explore this approach in future iterations of this work.

Acknowledgements

We would like to thank Nicolas Durrande for advising us on his prior work on the mean-zero functional constraints. KM is supported by a UK Engineering and Physical Sciences Research Council Doctoral Studentship. KRC is supported by a Banting postdoctoral fellowship from the Canadian Institutes of Health Research, and postdoctoral fellowships from the Canadian Statistical Sciences Institute and the UBC Data Science Institute. CY is supported by a UK Medical Research Council Research Grant (Ref: MR/P02646X/1) and by The Alan Turing Institute under the EPSRC grant EP/N510129/1.

References

- Bodin, E., Campbell, N. D., and Ek, C. H. Latent gaussian process regression. *arXiv preprint arXiv:1707.05534*, 2017.
- Damianou, A. and Lawrence, N. Deep gaussian processes. In *Artificial Intelligence and Statistics*, pp. 207–215, 2013.
- Damianou, A. and Lawrence, N. D. Semi-described and semi-supervised learning with gaussian processes. *31st Conference on Uncertainty in Artificial Intelligence (UAI)*, 2015.
- Damianou, A. C., Titsias, M. K., and Lawrence, N. D. Variational inference for latent variables and uncertain inputs in gaussian processes. *The Journal of Machine Learning Research*, 17(1):1425–1486, 2016.
- Durrande, N., Ginsbourger, D., and Roustant, O. Additive covariance kernels for high-dimensional gaussian process modeling. *Annales de la Faculté des sciences de Toulouse : Mathématiques*, Ser. 6, 21(3):481–499, 2012. doi: 10.5802/afst.1342.
- Durrande, N., Ginsbourger, D., Roustant, O., and Carraro, L. Anova kernels and rkhs of zero mean functions for model-based sensitivity analysis. *Journal of Multivariate Analysis*, 115:57–67, 2013.
- Dutordoir, V., Salimbeni, H., Hensman, J., and Deisenroth, M. Gaussian process conditional density estimation. In *Advances in Neural Information Processing Systems*, pp. 2385–2395, 2018.
- Duvenaud, D., Lloyd, J. R., Grosse, R., Tenenbaum, J. B., and Ghahramani, Z. Structure discovery in nonparametric regression through compositional kernel search. *arXiv preprint arXiv:1302.4922*, 2013.
- Duvenaud, D. K., Nickisch, H., and Rasmussen, C. E. Additive gaussian processes. In *Advances in neural information processing systems*, pp. 226–234, 2011.
- Ek, C. H. and Lawrence, P. *Shared Gaussian process latent variable models*. PhD thesis, PhD thesis, 2009.
- Gadd, C., Wade, S., Shah, A., and Grammatopoulos, D. Pseudo-marginal bayesian inference for supervised gaussian process latent variable models. *arXiv preprint arXiv:1803.10746*, 2018.
- Gao, X., Wang, X., Tao, D., and Li, X. Supervised gaussian process latent variable model for dimensionality reduction. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 41(2):425–434, 2011.
- Hensman, J., Fusi, N., and Lawrence, N. D. Gaussian processes for big data. *Uncertainty in Artificial Intelligence*, 2013.
- Hensman, J., Matthews, A. G. d. G., and Ghahramani, Z. Scalable variational gaussian process classification. 2015.
- Jiang, X., Gao, J., Wang, T., and Zheng, L. Supervised latent linear gaussian process latent variable model for dimensionality reduction. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 42(6):1620–1632, 2012.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014.
- Kuemmerle, N. B., Rysman, E., Lombardo, P. S., Flanagan, A. J., Lipe, B. C., Wells, W. A., Pettus, J. R., Froehlich, H. M., Memoli, V. A., Morganelli, P. M., et al. Lipoprotein lipase links dietary fat to solid tumor cell proliferation. *Molecular cancer therapeutics*, pp. molcanther-0802, 2011.
- Lawrence, N. Probabilistic non-linear principal component analysis with gaussian process latent variable models. *Journal of machine learning research*, 6(Nov):1783–1816, 2005.
- Munkley, J., McClurg, U. L., Livermore, K. E., Ehrmann, I., Knight, B., McCullagh, P., Mcgrath, J., Crundwell, M., Harries, L. W., Leung, H. Y., et al. The cancer-associated cell migration protein tspan1 is under control of androgens and its upregulation increases prostate cancer cell migration. *Scientific reports*, 7(1):5249, 2017.
- Peltola, T., Havulinna, A. S., Salomaa, V., and Vehtari, A. Hierarchical bayesian survival analysis and projective covariate selection in cardiovascular event risk prediction. In *Proceedings of the Eleventh UAI Conference on Bayesian Modeling Applications Workshop-Volume 1218*, pp. 79–88. Citeseer, 2014.
- Rasmussen, C. E. and Williams, C. K. *Gaussian process for machine learning*. MIT press, 2006.

- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In Xing, E. P. and Jebara, T. (eds.), *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pp. 1278–1286, Beijing, China, 22–24 Jun 2014. PMLR.
- Shon, A., Grochow, K., Hertzmann, A., and Rao, R. P. Learning shared latent structure for image synthesis and robotic imitation. In *Advances in neural information processing systems*, pp. 1233–1240, 2006.
- Tipping, M. E. and Bishop, C. M. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- Titsias, M. Variational learning of inducing variables in sparse gaussian processes. In *Artificial Intelligence and Statistics*, pp. 567–574, 2009.
- Urtasun, R. and Darrell, T. Discriminative gaussian process latent variable model for classification. In *Proceedings of the 24th international conference on Machine learning*, pp. 927–934. ACM, 2007.
- Wang, C. and Neal, R. M. Gaussian process regression with heteroscedastic or non-gaussian residuals. *arXiv preprint arXiv:1212.6246*, 2012.
- Weinstein, J. N., Collisson, E. A., Mills, G. B., Shaw, K. R. M., Ozenberger, B. A., Ellrott, K., Shmulevich, I., Sander, C., Stuart, J. M., Network, C. G. A. R., et al. The cancer genome atlas pan-cancer analysis project. *Nature genetics*, 45(10):1113, 2013.