
BasisVAE: Translation-invariant feature-level clustering with Variational Autoencoders

Kaspar Märtens
University of Oxford

Christopher Yau
Alan Turing Institute
University of Birmingham
University of Manchester

Abstract

Variational Autoencoders (VAEs) provide a flexible and scalable framework for non-linear dimensionality reduction. However, in application domains such as genomics where data sets are typically tabular and high-dimensional, a black-box approach to dimensionality reduction does not provide sufficient insights. Common data analysis workflows additionally use clustering techniques to identify groups of similar *features*. This usually leads to a two-stage process, however, it would be desirable to construct a joint modelling framework for simultaneous dimensionality reduction and clustering of features. In this paper, we propose to achieve this through the *BasisVAE*: a combination of the VAE and a probabilistic clustering prior, which lets us learn a one-hot basis function representation as part of the decoder network. Furthermore, for scenarios where not all features are *aligned*, we develop an extension to handle translation-invariant basis functions. We show how a collapsed variational inference scheme leads to scalable and efficient inference for BasisVAE, demonstrated on various toy examples as well as on single-cell gene expression data.

1 Introduction

High-dimensional data analysis employs dimensionality reduction techniques to extract major components of variation, often using nonlinear techniques, that enable interpretation and visualisation. Additionally, cluster-

ing can also be utilised to identify structure both within samples and also across features.

In this paper, our goal is to develop a general purpose approach for joint dimensionality reduction and clustering of features within the Variational Autoencoder (VAE) framework (Kingma and Welling, 2014). Our methodology particularly applies to any *tabular* data problem where individual features have a distinct physical meaning (e.g. a gene, a physiological measurement, the position of a joint on a moving body) and there is an interest in understanding the relationship between these features. For example, in genomics there is interest in identifying groups of genes with similar behaviour, but the standard VAE framework does not provide a mechanism to cluster features. Instead, in practice, often dimensionality reduction and clustering are carried out consecutively in an *ad hoc* manner.

As a motivating example, consider the synthetic gene expression data set shown in Figure 1A. In the heatmap, the columns have been ordered to highlight the existence of five feature groups and the rows arranged to indicate that the observations have a certain ordering. In a real data set this ordering of rows and columns would be *unknown* and part of the statistical inference task would be to identify this *latent* structure, i.e. dimensionality reduction. On the other hand, there is also interest in identifying groups of similar features, as shown in Figure 1B. Typically these two tasks are solved separately, whereas we propose a VAE extension to achieve this in a joint probabilistic model.

Specifically, our approach, which we refer to as BasisVAE, assumes that the features vary as functions of latent dimensions \mathbf{z} , and that there exists a finite set of such (basis) functions. Hence, BasisVAE aims to discover the scale-invariant clustering of features shown in Figure 1B whilst simultaneously inferring the latent variable \mathbf{z} (in this case, $\mathbf{z} \in \mathbb{R}$ relates to an ordering of the observations, but in general \mathbf{z} may be multi-dimensional). Moreover, in this example, three groups of features have similar shapes which are merely shifted

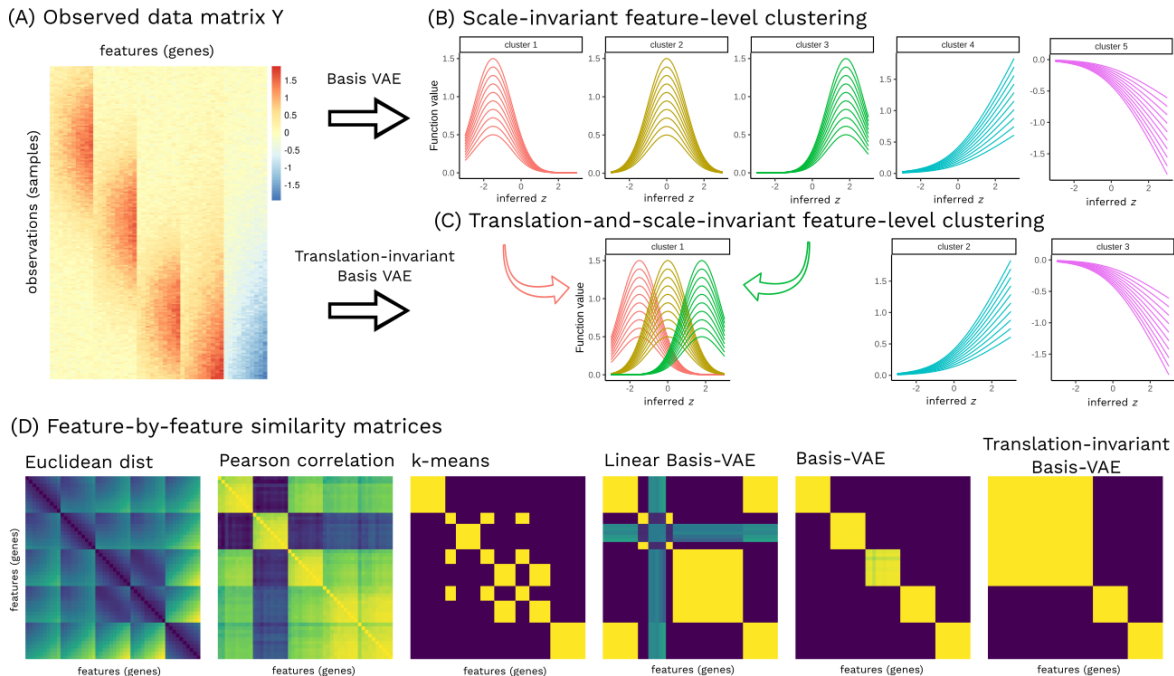


Figure 1: Illustration of our goal: Given a high-dimensional data matrix Y (heatmap in panel (A)), we want to learn a low-dimensional representation z and simultaneously cluster features based on the mappings $z \mapsto Y$ (to make it easier for the reader, we have ordered rows in panel (A) according to the inferred z). Panels (B) and (C) show the inferred clusters: clustering (B) is scale-invariant whereas clustering (C) is additionally translation-invariant. For every feature we have plotted the inferred neural network mapping $z \mapsto y^{(j)}$. These have been grouped into subpanels which correspond to inferred clusters). (D) shows feature-by-feature similarity matrices either based on distances (Euclidean, Pearson correlation), or co-clustering metrics either for k-means (with $K = 5$) or co-occurrence posterior probabilities in the BasisVAE framework.




Figure 2: BasisVAE provides feature-level interpretability: Having identified clusters as shown in Fig. 1, not only can we visualise (A) the inferred basis functions, but thanks to the *explicit* parameterisation we have access to the inferred (B) scaling parameters λ and (C) translation parameters δ (shown in colour).

versions of each other. To capture this, we introduce *translation invariance* as part of the model so it would allow these features to be further grouped together, as in Figure 1C, whilst maintaining an explicit parameterisation that would allow the unique properties of each feature to be maintained.

Feature-by-feature similarity measures shown in Figure 1D help us to further describe our desired output. Under Euclidean or Pearson correlation based distances, the structure of the features is not immediately obvious (note that the features displayed are ordered according

to their known groupings which would ordinarily be unknown) whilst k-means clustering ($k = 5$) does not uncover this structure either. A linear model is also unsuitable due to the nonlinear relationship between the features and the latent dimensions. In comparison, BasisVAE reveals the correct structure through joint dimensionality reduction and clustering. Translation invariance allows us to further determine whether to treat clusters 1-3 in Figure 1B as three distinct groups of features or as a single group. For easier interpretability, each feature has its own scale and translation parameters as shown in Figure 2.

In this paper, we endow VAEs with the capability to cluster features in the data. We propose to achieve this via introducing additional probabilistic structure within the decoder, thus combining the scalability of VAEs with a Bayesian mixture prior for the purpose of clustering features (architecture illustrated in Figure 3).

2 Background: VAEs

Variational Autoencoders (Kingma and Welling, 2014) are constructed based on particular a latent variable model structure. Let $y \in \mathbb{R}^P$ denote a data

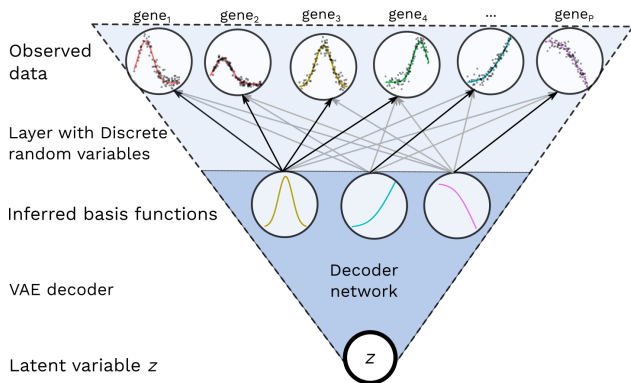


Figure 3: Diagram of the decoder architecture within BasisVAE. From top to bottom: observed data, the inferred basis functions, latent variable \mathbf{z} . A typical decoder network is used to map \mathbf{z} to the basis functions, whereas a specialised layer with Categorical random variables is used to map the latter to features. The posterior distribution of these Categorical variables provides us with the inferred clustering.

point in some potentially high-dimensional space and $\mathbf{z} \in \mathbb{R}^Q$ be a vector of associated latent variables typically in a much lower dimensional space. We will assume a prior $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Now suppose $f^\theta(\mathbf{z})$ is a family of deterministic functions, indexed by parameters $\theta \in \Theta$, such that $f : \mathbb{R}^Q \times \Theta \rightarrow \mathbb{R}^P$. In VAEs, these deterministic functions f^θ are given by deep neural networks (NNs). Given data points $\mathbf{y}_{1:N} = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$, the marginal likelihood is given by $p(\mathbf{y}_{1:N}) = \prod_{i=1}^N \int p_\theta(\mathbf{y}_i | \mathbf{z}_i) p(\mathbf{z}_i) d\mathbf{z}_i$. Here we will assume a Gaussian likelihood, i.e. $\mathbf{y}_i | \mathbf{z}_i, \theta \sim \mathcal{N}(f^\theta(\mathbf{z}_i), \sigma^2)$. Posterior inference in a VAE is carried out via *amortized* variational inference, i.e. with a parametric inference model $q_\phi(\mathbf{z}_i | \mathbf{y}_i)$ where ϕ are variational parameters. Application of standard variational inference methodology (e.g. (Blei et al., 2017)) leads to a lower bound on the log marginal likelihood, i.e. the ELBO of the form:

$$\mathcal{L} = \sum_{i=1}^N \mathbb{E}_{q_\phi(\mathbf{z}_i | \mathbf{y}_i)} [\log p_\theta(\mathbf{y}_i | \mathbf{z}_i)] - \text{KL}(q_\phi(\mathbf{z}_i | \mathbf{y}_i) || p(\mathbf{z}_i))$$

In the VAE, the variational approximation $q_\phi(\mathbf{z}_i | \mathbf{y}_i)$ is referred to as the *encoder* since it encodes the data \mathbf{y} into the latent variable \mathbf{z} whilst the *decoder* refers to the generative model $p_\theta(\mathbf{y} | \mathbf{z})$ which decodes the latent variables into observations.

Training a VAE seeks to optimise both the model parameters θ and the variational parameters ϕ jointly using stochastic gradient ascent. This typically requires a stochastic approximation to the gradients of the variational objective which is itself intractable. Typically the approximating family is assumed to be Gaus-

sian, i.e. $q_\phi(\mathbf{z}_i | \mathbf{y}_i) = \mathcal{N}(\mathbf{z}_i | \mu_\phi(\mathbf{y}_i), \sigma_\phi^2(\mathbf{y}_i))$, so that a reparametrisation trick can be used (Kingma and Welling, 2014; Rezende et al., 2014).

We write $f_{\text{decoder}} : \mathbb{R}^Q \rightarrow \mathbb{R}^P$ to denote the (multi-output) decoder. But as our interest lies in modelling tabular data where individual features have a distinct meaning, we additionally introduce notation $f_{\text{decoder}}^{(j)}$ where j indexes the j -th output dimension of the decoder network (i.e. predictions for feature j).

3 BasisVAE

The goal of our work is to extend the VAE to address the scientific problem of *identifying groups of features* with similar trends. That is, we would like to cluster features. We propose to achieve this via a mixture prior as part of the VAE decoding model. The idea is to introduce a set of “basis” functions $\{f_{\text{basis}}^{(k)} : k = 1, \dots, K\}$ parameterised by the neural network, and probabilistically assign every feature $\mathbf{y}^{(j)}$ to one of those basis functions $f_{\text{basis}}^{(k)}$. We refer to this as a one-hot-basis functional representation.

Specifically, we achieve this by replacing the decoder network $f_{\text{decoder}} : \mathbb{R}^Q \rightarrow \mathbb{R}^P$ with a basis decoder network $f_{\text{basis}} : \mathbb{R}^Q \rightarrow \mathbb{R}^K$ (where the number of basis functions $K < P$) whose output is mapped to the data via Categorical random variables, that is, for every data dimension $j \in \{1, \dots, P\}$

$$f_{\text{decoder}}^{(j)}(\mathbf{z}) := \sum_{k=1}^K w^{j,k} \cdot f_{\text{basis}}^{(k)}(\mathbf{z}) \quad (1)$$

where $(w^{j,1}, \dots, w^{j,K}) \sim \text{Categorical}(\pi_1, \dots, \pi_K)$. Note that this can be seen as a mixture model on the *features* rather than observations. The neural architecture of the BasisVAE decoder has been illustrated in Figure 3.

3.1 Translation and scale invariance

In the use of basis function representations, a question arises in how to represent two signals which have the same functional shape, but which may be shifted or scaled relative to one another. In a pure basis representation (as in equation (1)), two different basis functions would be learnt, one for each signal, but this may not always be desirable. Sometimes, we would instead like the basis function to only capture the *shape* of the pattern. This can be thought of as an additional layer of clustering in terms of functional similarity across features which is scale- and translation-invariant.

This behaviour could be achieved via introducing feature-specific scale and translation parameters to (1). For scale-invariance, we introduce scaling parameters

$\lambda_{j,k} \geq 0$ as follows

$$f_{\text{BasisVAE}}^{(j)}(\mathbf{z}) := \sum_{k=1}^K w^{j,k} \cdot \lambda_{j,k} \cdot f_{\text{basis}}^{(k)}(\mathbf{z}) \quad (2)$$

For additional translation-invariance, we introduce $\delta_{j,k}$

$$f_{\text{BasisVAE}}^{\text{translation-inv}}^{(j)}(\mathbf{z}) := \sum_{k=1}^K w^{j,k} \cdot \lambda_{j,k} \cdot f_{\text{basis}}^{(k)}(\mathbf{z} + \delta_{j,k}) \quad (3)$$

Throughout the paper, we refer to (2) as BasisVAE (i.e. we default to scale invariance) and to (3) as translation-invariant BasisVAE. Note that we have restricted $\lambda_{j,k}$ to be positive. This is because in the genomics applications that are our main interest, we want to distinguish between gene upregulation and downregulation, as these correspond to different biological processes (for this reason, we would not like to merge clusters 4 and 5 in Figure 1B).

The added value from scale and translation invariance is two-fold:

- **Interpretability:** Feature-specific scale and/or translation parameters lets us explicitly interpret how features within a cluster relate to each other.
- **Increased model capacity:** In model given by equation (1), a large number of basis functions K may be needed to capture the range of observed patterns, whereas (2) and (3) can provide the same flexibility with a significantly smaller K .

3.2 Specifying the number of basis functions

So far, we have presented BasisVAE as if the number of underlying basis functions was known. However, this is rarely the case in practice, usually we do not know K a priori. Ideally, we would like to be able to over-specify the value of K , so that that the unnecessary extra components would be left empty.

In the mixture model framework outlined above with a Categorical(π_1, \dots, π_K) prior, this would require careful specification of $\boldsymbol{\pi} = (\pi_1, \dots, \pi_K)$ which would be problem dependent. Hierarchical Bayesian models offer a way to get around this by placing a distribution over $\boldsymbol{\pi}$. One possible construction would be as follows

$$\begin{aligned} \boldsymbol{\pi} | \boldsymbol{\alpha} &\sim \text{Dirichlet}(\boldsymbol{\alpha}) \\ (w^{j,1}, \dots, w^{j,K}) | \boldsymbol{\pi} &\sim \text{Categorical}(\pi_1, \dots, \pi_K) \end{aligned} \quad (4)$$

This hierarchical construction can be useful, because for $\alpha < 1$ values the draws from $p(\boldsymbol{\pi} | \boldsymbol{\alpha})$ will be sparse, thus providing a mechanism to *learn* an appropriately sparse $\boldsymbol{\pi}$ rather than pre-specify its value. This particular setup has been popular for finite Dirichlet mixture models as well as their Dirichlet Process counterparts, where conjugacy made it possible to develop efficient (collapsed) Gibbs samplers, see e.g. (Neal, 2000). MCMC-based inference has perhaps been most popular

for such mixture models, but also variational inference techniques have been derived (e.g. analytic update rules for exponential families (Blei and Jordan, 2006)).

However, for non-linear latent variable models, such as VAEs, analytic updates are not available in closed form any more. Also it is not straightforward to exploit conjugacy in eq. (4). This imposes a challenge: on the one hand we want to utilise the computational benefits (scalability, modularity, computational efficiency) of the VAE-framework, while on the other hand we would like to make use of the sparsity properties that are present in more classical (“pre automatic differentiation era”) hierarchical Bayesian mixture models.

Existing work in the context of VAEs¹ have either relied on using a fixed value of $\boldsymbol{\pi} = (1/K, \dots, 1/K)$ (Dilokthanakul et al., 2016) or developed inference schemes where variational approximation is introduced on the joint space of $(\boldsymbol{\pi}, \mathbf{w})$ assuming factorisation $q(\mathbf{w}, \boldsymbol{\pi}) = q(\mathbf{w})q(\boldsymbol{\pi})$ (Nalisnick et al., 2016). The former approach is undesirable, as it has no mechanism to encourage sparsity. The latter could in principle work well, but its performance could suffer due to ignoring the (potentially strong) dependence between cluster allocations \mathbf{w} and the $\boldsymbol{\pi}$.

Next, we show how our attempts to adapt these ideas lead to *inefficient* inference schemes, before proposing a modern adaptation of the *collapsed* inference scheme within the VAE framework.

3.3 Generative model and inference

The generative model for the scale-and-translation-invariant BasisVAE is as follows

$$\begin{aligned} \mathbf{z}_i &\sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \\ \boldsymbol{\pi} | \boldsymbol{\alpha} &\sim \text{Dirichlet}(\boldsymbol{\alpha}) \\ (w^{j,1}, \dots, w^{j,K}) | \boldsymbol{\pi} &\sim \text{Categorical}(\pi_1, \dots, \pi_K) \\ y_i^{(j)} | \mathbf{w}^j, \mathbf{z}_i, \theta, \lambda, \delta &\sim \mathcal{N} \left(\sum_{k=1}^K w^{j,k} \lambda_{j,k} f_{\text{basis}}^{(k)}(\mathbf{z}_i + \delta_{jk}), \sigma_j^2 \right) \end{aligned}$$

where $i = 1, \dots, N$ indexes data points, $j = 1, \dots, P$ indexes features, and $k = 1, \dots, K$ basis components. We additionally place a prior over $\delta_{jk} \sim \mathcal{N}(0, 1)$ and $\lambda_{jk} \sim \Gamma(1, 1)$, but in our experiments we will perform MAP estimation over δ_{jk} and λ_{jk} .

Denoting observations collectively $\mathbf{Y} := \{\mathbf{y}_i^{(j)}\}$, standard VAE methodology would lower bound the log marginal likelihood as follows $\log p(\mathbf{Y}) \geq$

$$\geq \sum_{i=1}^N \mathbb{E}_{q_\phi(\mathbf{z}_i | \mathbf{y}_i)} [\log p_\theta(\mathbf{y}_i | \mathbf{z}_i)] - \text{KL}(q_\phi(\mathbf{z}_i | \mathbf{y}_i) || p(\mathbf{z}_i))$$

¹To our knowledge, existing work has only considered mixtures for clustering observations rather than features

However, unlike for VAE, in our model $\log p_\theta(\mathbf{y}_i|\mathbf{z}_i)$ is intractable due to the need to integrate over \mathbf{w} and $\boldsymbol{\pi}$.

3.3.1 Non-collapsed inference

One approach would be to introduce an approximate posterior $q(\mathbf{w}, \boldsymbol{\pi}) = \prod_{j=1}^P q(\mathbf{w}^j)q(\boldsymbol{\pi})$. For example, one can choose variational families $q(\mathbf{w}^j) = \text{Categorical}(\phi_{j,1}, \dots, \phi_{j,K})$ and $q(\boldsymbol{\pi}) = \text{Dirichlet}(\boldsymbol{\psi})$ where ϕ and $\boldsymbol{\psi}$ are variational parameters. We note that reparameterisation trick for the Dirichlet can be implemented as in (Figurnov et al., 2018). Assuming mean-field across $\{\mathbf{z}_i\}, \{\mathbf{w}^j\}, \boldsymbol{\pi}$, we obtain the ELBO

$$\sum_{i=1}^N \mathbb{E}_{q_\phi(\mathbf{z}_i|\mathbf{y}_i)} \mathbb{E}_{q(\mathbf{w})} \log p_\theta(\mathbf{y}_i|\mathbf{z}_i, \mathbf{w}) - \text{KL}(q_\phi(\mathbf{z}_i|\mathbf{y}_i)||p(\mathbf{z}_i)) - \text{KL}\left(\prod_{j=1}^P q(\mathbf{w}^j)q(\boldsymbol{\pi}) \parallel \prod_{j=1}^P p(\mathbf{w}^j|\boldsymbol{\pi})p(\boldsymbol{\pi})\right)$$

We refer to this approach as *non-collapsed* inference, and later demonstrate how this inference scheme can get stuck in local modes, resulting in undesirable behaviour.

3.3.2 Collapsed inference

Given the success of earlier work on collapsed inference schemes, e.g. collapsed Gibbs samplers (Neal, 2000) or collapsed VB schemes (Teh et al., 2007; Kurihara et al., 2007; Hensman et al., 2012, 2015), we hope to improve our previously described non-collapsed inference by marginalising out $\boldsymbol{\pi}$. If we adapted the approach of Teh et al. (2007), further approximations for evaluating the ELBO would be needed. Instead, we follow the approach of Hensman et al. (2012) where marginalisation is applied *after* (rather than before) making a variational approximation.

Now, following this strategy, we lower bound the term $\log p_\theta(\mathbf{y}|\mathbf{z})$, by first introducing an approximate posterior

$$q(\mathbf{w}) = \prod_{j=1}^P q(\mathbf{w}^j) = \prod_{j=1}^P \text{Categorical}(\phi_{j,1}, \dots, \phi_{j,K})$$

and then collapsing $\boldsymbol{\pi}$ from the bound without any further approximations. The detailed derivation can be found in Supplementary. As a result, we obtain the ELBO

$$\mathcal{L} = \sum_{i=1}^N \mathbb{E}_{q_\phi(\mathbf{z}_i|\mathbf{y}_i)} \mathbb{E}_{q(\mathbf{w})} \log p_\theta(\mathbf{y}_i|\mathbf{z}_i, \mathbf{w}) \quad (5)$$

$$+ \log \int \exp(\mathbb{E}_{q(\mathbf{w})} p(\mathbf{w}|\boldsymbol{\pi})) p(\boldsymbol{\pi}) d\boldsymbol{\pi} \quad (6)$$

$$- \mathbb{E}_{q(\mathbf{w})} \log q(\mathbf{w}) \quad (7)$$

$$- \sum_{i=1}^N \text{KL}(q_\phi(\mathbf{z}_i|\mathbf{y}_i)||p(\mathbf{z}_i)) \quad (8)$$

For the first term (5), we use a combination of reparameterisation and analytic marginalisation, the former for sampling $\mathbf{z}_i \sim q_\phi(\mathbf{z}_i|\mathbf{y}_i)$ and the latter for $q(\mathbf{w})$. So we evaluate

$$\sum_{i=1}^N \mathbb{E}_{q(\mathbf{z}_i|\mathbf{y}_i)} \sum_{j=1}^P \sum_{k=1}^K \phi_{j,k} \log \mathcal{N}(y_i^{(j)}|\lambda_{j,k} f_{\text{basis}}^{(k)}(\mathbf{z}_i + \delta_{jk}), \sigma_j^2).$$

The second term (6) can be derived analytically,

$$\log \frac{\Gamma(\sum_k \alpha_k)}{\prod_k \Gamma(\alpha_k)} \int \prod_{k=1}^K \pi_k^{n_k + \alpha_k - 1} d\boldsymbol{\pi}$$

where $n_k := \sum_{j=1}^P \phi_{j,k}$. This integral has a closed form, resulting in the analytic term

$$\log \frac{\Gamma(\sum_k \alpha_k) \prod_k \Gamma(n_k + \alpha_k)}{\prod_k \Gamma(\alpha_k) \Gamma(\sum_k n_k + \alpha_k)}.$$

Also the third and fourth term have analytic forms.

Finally, we perform MAP inference over λ_{jk} and δ_{jk} (but we note that variational inference could also be adapted), encouraging δ_{jk} values to be centered around 0 and λ_{jk} values centered around 1. So the final optimisation objective will incorporate additional penalties

$$\mathcal{L} + \sum_{j=1}^P \sum_{k=1}^K \log \mathcal{N}(\delta_{jk} | 0, 1) + \log \Gamma(\lambda_{jk} | 1, 1).$$

Adaptation of ELBO for large data sets: For large high-dimensional data sets, we note the lower bound \mathcal{L} will be dominated by the data log-likelihood term (5). As a result, our clustering prior will implicitly become less important for increasingly large N and P . While the property that the likelihood will dominate the prior in the large data regime is inherent to Bayesian models, it may not always be desirable, especially for mis-specified models. In practice, this can be alleviated via downweighing the likelihood or upweighing the prior. For example, β -VAE (Higgins et al., 2017) scales the KL-term by $\beta > 0$, which is closely connected to the ELBO on an alternative formulation involving an annealed prior $\propto p(\mathbf{z})^\beta$ (Hoffman et al., 2017; Mathieu et al., 2019). Similarly to β -VAE, we propose a modified objective \mathcal{L}_β (as outlined in Supplementary) where $\beta > 1$ lets us increase the relative importance of the sparse clustering prior for large-scale applications.

Scalability: As a result, we have derived an inference scheme which can be easily implemented as a VAE extension, thus making it possible to leverage the advantages of deep learning frameworks. From the computational efficiency perspective, the ELBO for the standard VAE involves summing over (a subset of) observations and all features, whereas our ELBO includes an additional sum over cluster components $k = 1, \dots, K$. Thus the additional scaling w.r.t. K is linear (but we note that this can be efficiently implemented by introducing an extra dimension to tensors).

4 Related work

To our knowledge, the problem of *clustering features* has not been previously addressed in the context of VAEs. In the regression setting, various probabilistic methods for curve clustering have been proposed before. E.g. Chudova et al. (2003) propose a method for translation-invariant curve clustering, assuming fixed inputs on an equi-spaced time grid. Gaffney and Smyth (2005) derive an EM-algorithm for curve clustering and alignment, assuming that curves are parameterised using either polynomial regression or cubic splines. Palla et al. (2012) propose a Bayesian non-parametric *linear* latent variable model. Thus, it could be seen as a special case of BasisVAE with a *linear* decoder and infinite K . However, the linearity assumption is too restrictive for the applications we consider in this paper. Campbell and Yau (2018) propose a specific translation-invariant framework for single-cell genomics, however it captures only two types of parametric mappings: sigmoidal and transient patterns. This is much more restrictive compared to BasisVAE that can *learn* the dynamics in an unsupervised way.

5 Experiments

It is challenging to evaluate the performance of clustering methods on real applications, as the ground truth cluster labels are rarely known (and even if labels exist, there may always be more fine-grained subgroups). For this reason, we first evaluate the performance of BasisVAE in a controlled setting, before considering the semi-synthetic clustering of motion capture data, and finally demonstrating the utility of our approach for the analysis of single-cell gene expression data.

Our PyTorch implementation of BasisVAE is available in <https://github.com/kasparmartens/BasisVAE>.

5.1 Toy data

We first consider the synthetic data shown in Figure 1 with 5 distinct groups of features (with 10 features each, $P = 5 \times 10$) as a function of latent $\mathbf{z} \in \mathbb{R}$. The data generative process was designed in order to mimic the patterns we expect to see in real gene expression data: first three groups of features exhibit a transient trend (where gene expression spikes for a certain range of \mathbf{z} values), whereas the fourth cluster exhibits gene up-regulation and fifth one down-regulation. The BasisVAE correctly infers five clusters (Figure 1B), whereas its translation-invariant version has merged the first three clusters together (Figure 1C).

Note that these results were obtained under our collapsed inference scheme. This would not have been

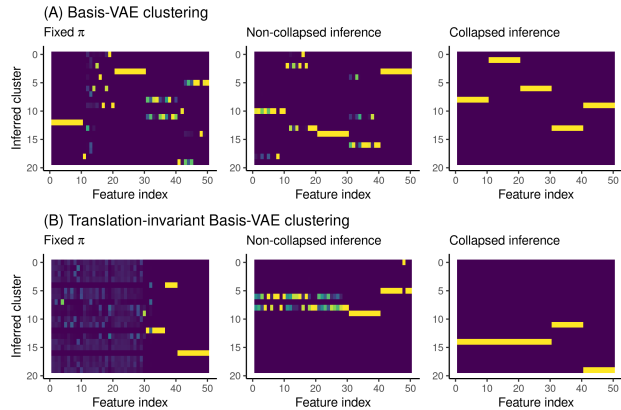


Figure 4: Comparison of inference techniques for $q(\mathbf{w}, \boldsymbol{\pi})$, for both (A) BasisVAE and (B) its translation-invariant version, using the synthetic data as in Fig 1 (consisting of 5 groups with 10 features each). The heatmaps show posterior probabilities for cluster allocations, i.e. elements of $q(\mathbf{w})$. Both “fixed $\boldsymbol{\pi}$ ” and “non-collapsed” inference (left and middle panels) result in suboptimal and non-sparse cluster assignments, whereas collapsed inference correctly identifies the underlying ((A) $K = 5$ and (B) $K = 3$) clusters.




Figure 5: A large number of basis functions is required to accurately capture the set of shifted mappings (cluster allocations shown in colour for BasisVAE using (A) $K=3$, (B) $K=5$, and (C) $K=12$), whereas (D) translation-invariant BasisVAE can represent this with a single basis function.

possible with simpler approaches such as the “fixed $\boldsymbol{\pi}$ ” and “non-collapsed” inference, as shown in Figure 4, which result in cluster allocations that are less sparse than the one inferred by “collapsed” inference (all comparisons were carried out with K fixed to $K = 20$ and $\alpha = 0.1$). “Fixed $\boldsymbol{\pi}$ ” has failed likely due to the reason that $\pi_k = 1/K$ does not encourage sparse solutions, whereas “non-collapsed” inference is more prone to getting stuck in local modes due to the dependence between \mathbf{w} and $\boldsymbol{\pi}$. Supp Fig 1 quantifies clustering performance using the V-measure (Rosenberg and Hirschberg, 2007) over 10 random restarts.

Various widely used methods for clustering gene expression data rely on the feature-by-feature similarity matrix. As shown in Fig. 1D, these do not accurately capture the underlying cluster structure, and do not provide a way to naturally incorporate scale and translation invariance. Also the linear BasisVAE, which can

be seen as a particular implementation of (Palla et al., 2012), does not provide the correct clustering either.

Next, we demonstrate the capabilities of BasisVAE to capture translation invariance in a setting where we observe features spanning a range of δ values (Fig. 5). The non-translation-invariant BasisVAE clustering behaves intuitively: for small $K \in \{3, 5, 12\}$ (panels A-C), features with similar δ values are clustered together (cluster allocations shown in colour). This may be sufficient for certain applications, but there is no way to quantify how different clusters relate to one another. This capability is present in the translation-invariant BasisVAE, which identifies a single cluster (Fig 5D) and relates all features within the cluster to one another explicitly via inferred δ .

5.2 Motion capture data

Next we investigate the behaviour of our model in a semi-controlled setting. We use CMU Motion Capture database which consists of short segments of various activities over a repeated number of trials. Note that even if we focus on the same activity performed by the same subject, the measurements (the locations of joints) are not *synchronised* between repeated trials. We note that this alignment problem has a particular structure which is inherently simpler than the structure in the genomics data in the next section. First, unlike for the genomics application, for motion capture data time is observed. Second, we know that we need to align measurements only *between* trials, but not within a trial (if measurements for one joint are shifted by δ , it is likely the case for other joints as well).

Here, we focus on the 10 repeated trials of the activity “golf swing”. For the purpose of demonstrating the capabilities of BasisVAE, we treat 10 trials of 59-dimensional measurements as a high-dimensional data set with $P = 590$ dimensions without *a priori* known structure. Our goal is three-fold: First, we hope to uncover clusters that are shared *across* trials, e.g. we expect the “right hand” movements to follow the same dynamics over all trials. Second, we hope to explicitly align the movements across trials. Third, we would like to find *shared* dynamics between the marker positions, e.g. two nearby joints may exhibit similar movement.

We have visualised results in Fig 6 for a subset of features (in columns) and all 10 trials (in rows), showing the inferred mappings (lines) and the respective cluster allocations (in colour). The features can be grouped into 1) right-hand, 2) head-and-neck, and 3) lower-and-upper-back related movement. Note that for many features, all 10 trials have been assigned to the same cluster, whereas sometimes there are subtle differences between the trials (e.g. for head and neck, the “bump”

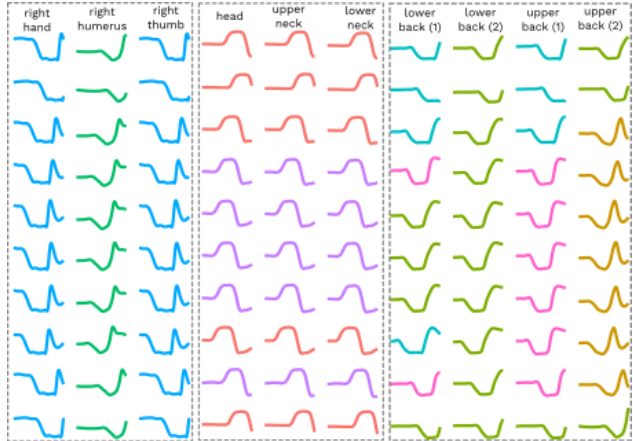


Figure 6: BasisVAE on motion capture data. The inferred mappings together with the clustering (shown in colour), shown over 10 repeated trials (in rows) and selected features (in columns). Features have been retrospectively grouped into 1) right-hand, 2) head-and-neck, and 3) lower-and-upper-back related features.




Figure 7: BasisVAE lets us explicitly align features in motion capture data, demonstrated on features “right hand” and “lower back (2)” (δ values shown in colour).

for some trials lasts slightly longer than for others). Our model provides interpretability that is not present in alternative methods. We can align features between trials (inferred δ values shown in Figure 7), and identify groups of similar features. E.g. “right hand” and “right thumb” have been assigned the same cluster, whereas “right humerus” exhibits a different dynamic.

5.3 Single-cell Spermatogenesis data

Recently, there has been a lot of interest in single-cell genomics, where gene expression is measured on the level of individual cells. This results in large data sets, both in the number of samples (i.e. cells) as well as the number of features (i.e. genes). The ability to quantify individual cells has led to a number of studies where the aim is to characterise *cellular differentiation* – the process by which a cell changes from one cell type to another. As the technology lets us measure every cell only once, we do not have access to the true underlying time, how far along the differentiation trajectory every cell is. Those quantities are referred to as *pseudotimes* (Trapnell et al., 2014), and are treated

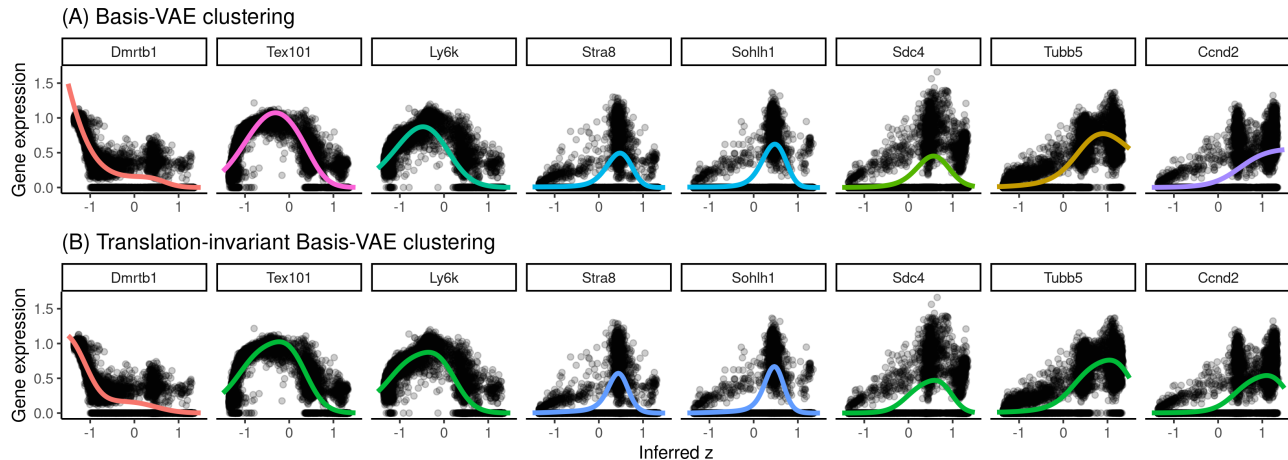


Figure 8: BasisVAE has discovered structure in the large-scale single-cell spermatogenesis data. For selected genes (whose peak expression along the z -axis varies from left to right, i.e. from *Dmrtb1* to *Ccnd2*), we have shown how (A) BasisVAE has identified a number of clusters with distinct gene expression behaviour (clusters shown in colour), whereas (B) its translation-invariant version has grouped genes with similar profiles together (clusters shown in colour), thus aiding further interpretability.

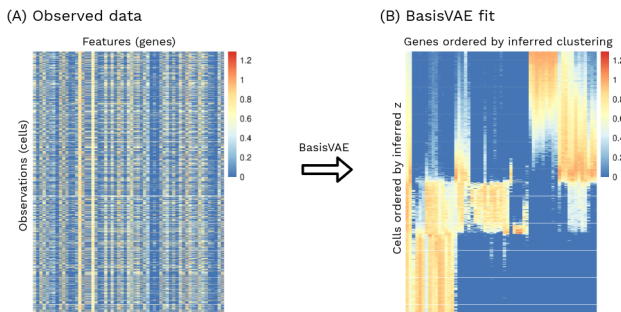


Figure 9: BasisVAE has uncovered structure in spermatogenesis data. (A) Heatmap of the observed data for a subset of genes. (B) Same data with re-ordered rows (ordered by z) and columns (ordered by clusters).

as one-dimensional latent variables \mathbf{z} .

The main scientific interest in these studies lies in getting insights into the biological processes that drive this latent dimension, e.g. identifying groups of genes whose expression changes with \mathbf{z} , and characterising those changes. In this section, we demonstrate the utility of BasisVAE for such purposes. Here, we re-analyse a publicly available data set from a recently published study [Ernst et al. \(2019\)](#), where the authors used single-cell RNA-Sequencing to profile mouse spermatogenesis. Due to the data being non-negative-valued and exhibiting zero-inflation, here we follow a common choice (e.g. as in [Lopez et al., 2018](#)) and we use the zero-inflated negative binomial likelihood in the decoder.

We pick a subset of samples ($N = 8509$) and $P = 7000$ most variable genes, and fit BasisVAE without and with translation invariance (we used $K = 50$ clusters). Supp Fig 2 illustrates how the inferred \mathbf{z} relates to the first two principal components (PC1, PC2) of the

data. Manually interpreting individual genes in such a large-scale setting is infeasible, however the inferred BasisVAE cluster allocations make interpretation easier. In Figure 8 we display some of the genes that were highlighted in the analysis of [Ernst et al. \(2019\)](#) together with our inferred mappings and cluster allocations. Note that in the scale-invariant version, genes exhibiting the same shape but different shift were allocated different clusters, whereas e.g. genes *Stra8* and *Sohlh1* were identified as exhibiting the same behaviour. This is different from the translation-invariant model, where e.g. genes *Tex101*, *Ly6k*, *Sdc4*, *Tubb5*, *Ccnd2* were all assigned to the same cluster, so not only do we know that they have the same shape, but we can also interpret their relative shift by inspecting the δ values.

While the inferred clustering structure (e.g. where the behaviour of gene might be delayed relative to another, but exhibiting the same shape) does not imply any biological functional relatedness, BasisVAE methodology can add to the exploratory toolkit for investigators to explore and rank such potential relationships between genes explicitly rather than retrospectively through *ad hoc* approaches that are common in genomics.

6 Discussion

We have proposed a joint model for dimensionality reduction and clustering of features, implemented in the VAE framework. Specifically, we have modified the decoder to incorporate a hierarchical Bayesian clustering prior, and demonstrated how collapsed variational inference can identify sparse solutions when over-specifying K . Furthermore, translation invariance lets us handle scenarios where not all features are *aligned*.

Acknowledgements

KM was supported by a UK Engineering and Physical Sciences Research Council Doctoral Studentship. 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

- Blei, D. M. and Jordan, M. I. (2006). Variational inference for Dirichlet process mixtures. *Bayesian analysis*, 1(1):121–143.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877.
- Campbell, K. R. and Yau, C. (2018). A descriptive marker gene approach to single-cell pseudotime inference. *Bioinformatics*, 35(1):28–35.
- Chudova, D., Gaffney, S., Mjolsness, E., and Smyth, P. (2003). Translation-invariant mixture models for curve clustering. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 79–88. ACM.
- Dilokthanakul, N., Mediano, P. A., Garnelo, M., Lee, M. C., Salimbeni, H., Arulkumaran, K., and Shananhan, M. (2016). Deep unsupervised clustering with gaussian mixture variational autoencoders. *arXiv preprint arXiv:1611.02648*.
- Ernst, C., Eling, N., Martinez-Jimenez, C. P., Marioni, J. C., and Odom, D. T. (2019). Staged developmental mapping and X chromosome transcriptional dynamics during mouse spermatogenesis. *Nature communications*, 10(1):1251.
- Figurnov, M., Mohamed, S., and Mnih, A. (2018). Implicit Reparameterization Gradients. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 441–452. Curran Associates, Inc.
- Gaffney, S. J. and Smyth, P. (2005). Joint probabilistic curve clustering and alignment. In *Advances in neural information processing systems*, pages 473–480.
- Hensman, J., Rattray, M., and Lawrence, N. D. (2012). Fast variational inference in the conjugate exponential family. In *Advances in neural information processing systems*, pages 2888–2896.
- Hensman, J., Rattray, M., and Lawrence, N. D. (2015). Fast Nonparametric Clustering of Structured Time-Series. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):383–393.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. (2017). beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. *International Conference on Learning Representations*.
- Hoffman, M. D., Riquelme, C., and Johnson, M. J. (2017). The beta-VAE’s Implicit Prior. In *Workshop on Bayesian Deep Learning, NIPS*, pages 1–5.
- Kingma, D. P. and Welling, M. (2014). Auto-encoding variational bayes. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Kurihara, K., Welling, M., and Teh, Y. W. (2007). Collapsed Variational Dirichlet Process Mixture Models. In *IJCAI*, volume 7, pages 2796–2801.
- Lopez, R., Regier, J., Cole, M. B., Jordan, M. I., and Yosef, N. (2018). Deep generative modeling for single-cell transcriptomics. *Nature methods*, 15(12):1053.
- Mathieu, E., Rainforth, T., Siddharth, N., and Teh, Y. W. (2019). Disentangling Disentanglement in Variational Autoencoders. In *International Conference on Machine Learning*, pages 4402–4412.
- Nalisnick, E., Hertel, L., and Smyth, P. (2016). Approximate inference for deep latent gaussian mixtures. In *NIPS Workshop on Bayesian Deep Learning*.
- Neal, R. M. (2000). Markov Chain Sampling Methods for Dirichlet Process Mixture Models. *Journal of Computational and Graphical Statistics*, 9(2):249–265.
- Palla, K., Ghahramani, Z., and Knowles, D. A. (2012). A nonparametric variable clustering model. In *Advances in Neural Information Processing Systems*, pages 2987–2995.
- Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*.
- Rosenberg, A. and Hirschberg, J. (2007). V-measure: A conditional entropy-based external cluster evaluation measure. In *Proceedings of the 2007 joint conference on empirical methods in natural language processing and computational natural language learning (EMNLP-CoNLL)*, pages 410–420.
- Teh, Y. W., Newman, D., and Welling, M. (2007). A collapsed variational Bayesian inference algorithm for latent Dirichlet allocation. In *Advances in neural information processing systems*, pages 1353–1360.
- Trapnell, C., Cacchiarelli, D., Grimsby, J., Pokharel, P., Li, S., Morse, M., Lennon, N. J., Livak, K. J., Mikkelsen, T. S., and Rinn, J. L. (2014). The dynamics and regulators of cell fate decisions are revealed by pseudotemporal ordering of single cells. *Nature biotechnology*, 32(4):381.