
Correlating Variational Autoencoders Natively For Multi-View Imputation

Ella S. C. Orme

Department of Mathematics
Imperial College London
ella.orme18@imperial.ac.uk

Marina Evangelou

Department of Mathematics
Imperial College London
m.evangelou@imperial.ac.uk

Ulrich Paquet

African Institute for Mathematical Sciences, South Africa
ulrich@aims.ac.za

Abstract

Multi-view data from the same source often exhibit correlation. This is mirrored in correlation between the latent spaces of separate variational autoencoders (VAEs) trained on each data-view. A multi-view VAE approach is proposed that incorporates a joint prior with a non-zero correlation structure between the latent spaces of the VAEs. By enforcing such correlation structure, more strongly correlated latent spaces are uncovered. Using conditional distributions to move between these latent spaces, missing views can be imputed and used for downstream analysis. Learning this correlation structure involves maintaining validity of the prior distribution, as well as a successful parameterization that allows end-to-end learning.

1 Introduction

Data from multiple sources describing the same subjects arises in a wealth of settings. This can be clinical information of patients alongside genetic information and scan data. Datasets consisting of multiple views are referred to as multi-view or multi-modal data. There are instances where not all views are always available for every realisation. For example, a patient may miss an appointment or machinery may falter, resulting in no reading for a specific view. Missing data results in smaller usable datasets and reduced statistical power with many methods only applicable to full datasets [1] and can result in reduced performance in downstream analysis [2]. This manuscript presents a multi-view imputation approach, where its aim is to impute the realisations of a missing view by incorporating the information learnt from the other view.

The proposed multi-view imputation method, named Joint Prior Variational Autoencoder (JPVAE), is based on variational autoencoders (VAEs) [3], with the views connected solely through a joint prior on the VAEs' latent embeddings. Standard autoencoders seek to encode a latent representation of data, and from this encoding reconstruct the original data via a decoder. Multi-view approaches allow the latent representation of a missing view to be obtained, from which the reconstruction can be used as an imputation of the missing view. Several multi-view imputation approaches exist in the literature based on autoencoders [4]. However, as variational autoencoders learn a continuous embedding, they provide better interpolation of the latent space than standard autoencoders, making them a more suitable approach for an imputation method. The proposed joint prior in JPVAE incorporates a non-zero correlation structure that is found to increase the observed correlation between views in the latent spaces. This allows for successful movement between latent spaces, improving the ability to impute missing views.

Various approaches have been proposed extending VAEs to the multi view case, including those by [5–8], with missing data imputation included as a feature of these methods. Most recently, proposed methods differ via their method of approximating a joint posterior. Daunhawer et al. [9] discuss the undesirable upper bound these methods put on a lower limit of the log-likelihood used as the objective function in VAEs. In contrast to existing approaches, which largely assume some joint posterior and/or shared latent space, our work has the novelty that is based on a joint prior between the latent variables. To the best of our knowledge, this is the first attempt made to correlate the latent spaces of VAEs natively through a joint prior.

JPVAE is most similar in structure to the private version of Deep Variational Canonical Correlation Analysis (VCCA-private), a multi-view VAE approach by Wang et al. [10]. However their model contains both private and shared variables in the latent space and doesn't have a suitable approach for imputing missing views. JPVAE can also be seen as an application of the ideas from self-supervised learning techniques such as Barlow Twins [11], where the embeddings from noisy versions of the same input are driven to be highly correlated via a loss function.

Figure 1 illustrates an application of the proposed approach JPVAE, on imputing view 2 of the data that corresponds to the bottom half of MNIST digits using the top half of the digits data (view 1). Through the proposed approach, reconstructions of missing views can be obtained. This is achieved through the conditional distribution between the two latent variables, as illustrated in Figure 2. By incorporating a joint prior with a non-zero cross-correlation structure, we observe better quality results in imputing view 2. Further realisations of image imputation, including imputation of view 1 from view 2, can be found in Appendix A.5.

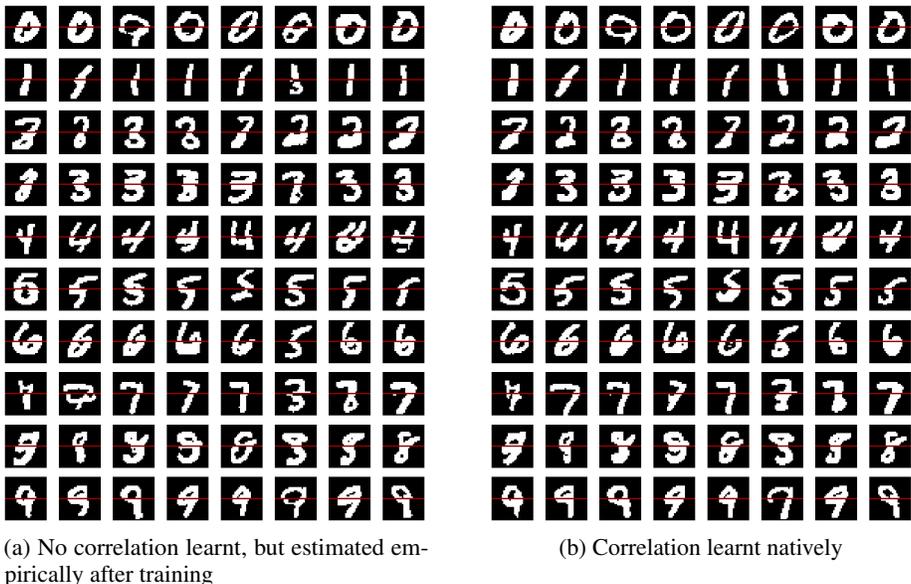


Figure 1: Imputation of the bottom half of MNIST digits (view 2 of the data) using the top half of the image (view 1) on a JPVAE model trained with (a) independent priors (completely separate VAEs) and (b) a joint prior with learnt correlation structure between latent spaces. The cross entropy loss between true bottom half of image and imputation is 109.1 in (a) and 93.04 in (b). A classifier trained on the concatenation of the reconstruction of view 1 and imputation of view 2 achieves average testing accuracy of 79.92% (1.0) in (a) and 87.45% (0.27) in (b). See Appendix A.5 for further details.

1.1 Contributions

A novel multi-view imputation approach based on variational autoencoders, named JPVAE, is proposed that in contrast to existing work assumes a joint prior between the latent variables.¹ As the multiple views are correlated, the generated latent spaces are found also to be correlated. JPVAE

¹The code for implementation of JPVAE and the numerical experiments contained in this manuscript can be found at <https://github.com/eso28599/JPVAE>.

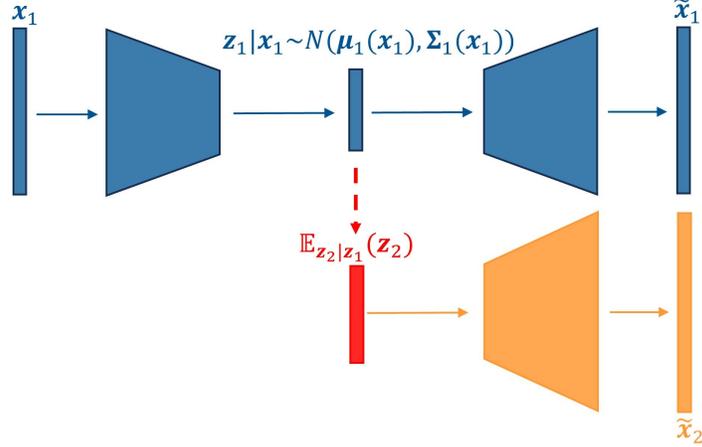


Figure 2: Workflow to obtain reconstruction \tilde{x}_1 and imputation $\tilde{x}_{2|1}$ from input x_1 only, on pre-trained JPVAE network. Trapeziums represent encoders/decoders and rectangles represent vectors. The structures for view 1 and 2 are shown in blue and orange respectively. The expectation step is shown in red.

takes advantage of this correlation and whilst marginals for each VAE are assumed to follow a standard normal distribution, a joint prior is assumed with a non-zero cross correlation. The imposed correlation structure is learnt natively by forcing a stronger correlation in the latent space. Section 2 presents the proposed method and related theorems that showcase the validity of the proposed method. As part of our work, we present and prove theorems that enable the parameterization of differentiable positive definite matrices that allow end-to-end learning. Through our conducted experiments presented in Section 3, learning the correlation structure leads to improved imputation ability, lower reconstruction loss and better predictive likelihood. Imputed views from JPVAE can be used for downstream tasks, with improved classification accuracy demonstrated. Lastly, better VAE models are learnt with JPVAE preventing posterior collapse, a common problem observed with VAEs [12].

1.2 Notation

Throughout this paper matrices are denoted by capital letters and column vectors are denoted by lower case letters, both emboldened. $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$ with vector $\mathbf{x}_i \in \mathbb{R}^c$ represents an entire dataset in $\mathbb{R}^{n \times c}$, with \mathbf{x}_i an individual realisation. A diagonal matrix with vector \mathbf{a} along the diagonal is represented by $\text{diag}(\mathbf{a})$. A block diagonal matrix with matrices \mathbf{A}_i along the diagonal is represented by $\text{bdiag}(\mathbf{A}_i)$. Vertical concatenation of column vectors \mathbf{a} and \mathbf{b} is denoted by $(\mathbf{a}; \mathbf{b})$ i.e. $(\mathbf{a}; \mathbf{b}) = (\mathbf{a}^T; \mathbf{b}^T)^T$. The eigenvalues and singular values of matrix \mathbf{A} are denoted by $\lambda_i(\mathbf{A})$ and $\sigma_i(\mathbf{A})$ respectively, indexed by subscript i in order of decreasing magnitude. \mathbf{I} denotes an identity matrix of relevant dimension.

2 Multi-view variational autoencoder with a joint prior

This section presents the novel multi-view VAE approach JPVAE where each view has a separate associated VAE. Before discussing the proposed approach in detail, the main concepts of VAEs are introduced. Subsequently the single-view VAE is extended to the multi-view setting and the joint prior is presented. The different components of JPVAE are presented including theoretical work on matrix parameterization which enables end-to-end learning.

2.1 Variational autoencoder

A standard VAE seeks to encode data \mathbf{X} in a probabilistic latent space and decode from this space to reconstructed data $\tilde{\mathbf{X}}$ [3]. The goal is for $\tilde{\mathbf{X}}$ to be as close as possible to \mathbf{X} . As the latent space is a probabilistic embedding rather than one obtained by a deterministic mapping, it has the advantage that

it can be explored fully, including the sampling of new realisations from the same distribution as $\tilde{\mathbf{X}}$. These realisations are assumed to be taken from the same underlying distribution as \mathbf{X} , $p_\theta(\cdot) = p(\cdot|\theta)$ where θ is the set of parameters defining the distribution.

VAEs are a type of variational Bayesian method that seek to find a lower bound for this marginal probability, $p_\theta(\mathbf{X})$, through a Bayesian framework. The likelihood of $\mathbf{x} \in \mathbb{R}^c$ given latent variable $\mathbf{z} \in \mathbb{R}^d$ is denoted by $p_\theta(\cdot|\mathbf{z})$ and the prior for the latent variables, usually taken to be a standard normal, is denoted by $p_\theta(\mathbf{z})$. The encoder and decoder are neural networks that seek to learn the posterior distribution of \mathbf{z} given \mathbf{x} , $p_\theta(\cdot|\mathbf{x})$, and the likelihood of \mathbf{x} given \mathbf{z} respectively, given the assumed prior over the latent space. In order to make the posterior learnable, an *approximate* posterior distribution is used that usually is a multivariate normal distribution. The approximation $p_\theta(\cdot|\mathbf{x}) \approx q_\phi(\cdot|\mathbf{x})$ is made where ϕ denotes the parameters of the probabilistic encoder, $q_\phi(\cdot|\mathbf{x})$. The encoder maps an input \mathbf{x} to a mean vector $\boldsymbol{\mu}(\mathbf{x}) \in \mathbb{R}^d$ and to the log of the vector of variances, $\boldsymbol{\sigma}^2(\mathbf{x}) \in \mathbb{R}^d$. A sample is drawn from the approximate posterior distribution and this is used as an input to the neural network which acts as the decoder, D_θ . As ϕ appears within the distribution of the latent variables, derivatives cannot simply be taken inside the expectation term. Differentiability of the loss function is required for the loss-driven parameter update steps and therefore a reparameterization trick needs to be implemented [3].

Instead of drawing \mathbf{z} directly, $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \mathbf{I})$ is drawn and $\mathbf{z} = \boldsymbol{\mu}(\mathbf{x}) + \text{diag}(\boldsymbol{\sigma}^2(\mathbf{x}))\boldsymbol{\epsilon}$ is determined. The distribution over which we take expectation is now independent of the parameters for which we take derivatives, and derivative update steps may now be implemented. The neural network decoder then seeks to reconstruct \mathbf{x} from the latent variable \mathbf{z} .

A maximum likelihood principle is then applied to the marginal likelihood of the data $p_\theta(\mathbf{X})$ to obtain estimates for the parameters within the encoders and decoders. As the likelihood is intractable a lower bound on the log-likelihood, known as the Evidence Lower Bound (ELBO) is instead maximised, given by:

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_{\mathbf{z} \sim q_\phi(\cdot|\mathbf{x})} [\ln(p_\theta(\mathbf{x}|\mathbf{z}))] - D_{KL}(q_\phi(\cdot|\mathbf{x})||p_\theta(\cdot)) \quad (1)$$

where $D_{KL}(r||s)$ denotes the Kullback-Leibler (KL) divergence between distributions r and s [13]. $\mathbb{E}_{\mathbf{z} \sim q_\phi(\cdot|\mathbf{x})}(\cdot)$ denotes the expectation with respect to the conditional distribution of \mathbf{z} given \mathbf{x} . Maximising the ELBO is equivalent to finding a balance between an approximate posterior that is close to the prior and putting weight on the latent variable space that maximises the likelihood of the data given these same variables, $\ln(p_\theta(\mathbf{x}|\mathbf{z}))$. Without the KL term, the delta function is returned as the approximate posterior and the autoencoder is recovered.

With this VAE formulation the KL term often becomes very small or ‘vanishes’ leading to the posterior equalling the prior (posterior collapse). This is referred to as KL vanishing and leads to a decoder that is largely independent of the latent variables. See [14] for further discussion of this problem. To combat this, we implement KL annealing – a procedure where a weight β is introduced on the KL term and gradually increased, typically from 0, as learning occurs [15]. The objective function becomes:

$$\mathcal{L}_\beta(\theta, \phi) = \mathbb{E}_{\mathbf{z} \sim q_\phi(\cdot|\mathbf{x})} [\ln(p_\theta(\mathbf{x}|\mathbf{z}))] - \beta D_{KL}(q_\phi(\cdot|\mathbf{x})||p_\theta(\cdot)). \quad (2)$$

It is now assumed two data views with shared row dimension are present and represented by $\mathbf{X} = \{\mathbf{X}_1, \mathbf{X}_2\}$ with $\mathbf{X}_i \in \mathbb{R}^{N \times c_i}$. For brevity, the sample subscript is dropped and $\mathbf{x}_i \in \mathbb{R}^{c_i}$ refers to an individual realisation from view i . Here N is the number of rows in both views and c_i is the number of columns in view $i = 1, 2$. This shared row dimension implies the views are paired, with row j representing the same individual/sample in both views. This allows for meaningful correlation in the latent space. The notation introduced in this section for single view VAEs is used in the following sections where multi-view VAEs are presented. Subscripts are used to denote the relevant views e.g. the encoder and decoder in view i are represented by $E_{i\phi_i}$ and $D_{i\theta_i}$ respectively.

2.2 Joint prior variational autoencoder

As the different views in a multi-view dataset are from a common source, correlation exists between the views. This translates to correlation between the latent spaces of each view from independently trained VAEs. JPVAE takes advantage of this correlation, enforcing the relationship between the two views via a joint prior on the latent variables, as illustrated in Figure 3. Enforcing the proposed correlation structure between the latent spaces ensures we can move from the original space where

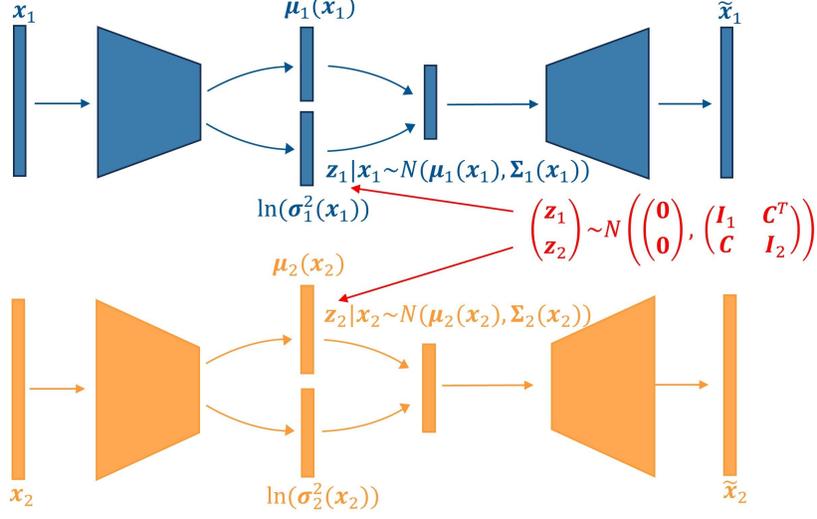


Figure 3: Workflow to obtain reconstructions \tilde{x}_1 and \tilde{x}_2 from realisations x_1 and x_2 using a learnt JPVAE structure. Trapeziums present the encoders and decoders of the two views. Vectors are presented by the rectangles. The structures for view 1 and 2 are shown in blue and orange respectively. The prior on the latent variables is shown in red.

data are highly correlated in a non-linear fashion, to a space where the correlation is linear, and back to the reconstructed space that has a non-linear correlation. The marginal prior on the latent variables corresponding to each view is a standard normal, as with the traditional VAE. However, a cross-covariance matrix C is assumed between the latent variables z_1 and z_2 . Having separate VAE structures for each view assumes conditional independence in both directions: (a) given the data the latent variables are independent, and (b) given the latent variables the data are independent. This allows the unique features of each view to be encoded and decoded.

As the latent spaces are linearly correlated, it is possible to move between them via the conditional distribution, as illustrated in Figure 2. This allows for imputation of missing views, obtaining a reconstruction of x_2 solely from realisation x_1 (and vice versa). If a joint prior is not assumed, separate VAEs are trained on each view and there is no correlation enforced between latent spaces. Whilst some correlation exists between the latent spaces, it is not as strong, and therefore may not produce as accurate reproductions, as illustrated in the numerical experiments of Section 3.

2.3 Objective function

By assuming that latent variables for each view are independent given the data, the approximated posterior distributed can be expressed as:

$$q_\phi(z|\mathbf{x}) = q_{(\phi_1, \phi_2)}((z_1, z_2)|(\mathbf{x}_1, \mathbf{x}_2)) = q_{\phi_1}(z_1|\mathbf{x}_1)q_{\phi_2}(z_2|\mathbf{x}_2). \quad (3)$$

The individual approximate posteriors $q_{\phi_i}(\cdot|\mathbf{x}_i)$ are multivariate Gaussians with mean and covariance matrix determined by the output of $E_{i\phi_i}$. For input \mathbf{x}_i , these are $\mu_i(\mathbf{x}_i)$ and $\text{diag}(\sigma_i^2(\mathbf{x}_i))$ respectively. As the posteriors are assumed to be independent, the joint distribution is multivariate Gaussian with mean $(\mu_1(\mathbf{x}_1); \mu_2(\mathbf{x}_2))$. The covariance matrix is represented by $\text{bdiag}(\Sigma_i(\mathbf{x}_i))$ with $\Sigma_i(\mathbf{x}_i) = \text{diag}(\sigma_i^2(\mathbf{x}_i))$. Similarly, by assuming the data is independent given the latent variables, the likelihood function can be expressed as:

$$p_\theta(\mathbf{x}|\mathbf{z}) = p_{(\theta_1, \theta_2)}((\mathbf{x}_1, \mathbf{x}_2)|(\mathbf{z}_1, \mathbf{z}_2)) = p_{\theta_1}(\mathbf{x}_1|\mathbf{z}_1)p_{\theta_2}(\mathbf{x}_2|\mathbf{z}_2). \quad (4)$$

This is equivalent to having separate encoders and decoders for each view but with a joint prior. Assuming independence means the expectation term in the ELBO can be separated into terms corresponding to the separate views:

$$\mathbb{E}_{\mathbf{z} \sim q_\phi(\cdot|\mathbf{x})} [\text{In}p_\theta(\mathbf{x}|\mathbf{z})] = \mathbb{E}_{\mathbf{z}_1 \sim q_{\phi_1}(\cdot|\mathbf{x}_1)} [\text{In}p_{\theta_1}(\mathbf{x}_1|\mathbf{z}_1)] + \mathbb{E}_{\mathbf{z}_2 \sim q_{\phi_2}(\cdot|\mathbf{x}_2)} [\text{In}p_{\theta_2}(\mathbf{x}_2|\mathbf{z}_2)].$$

The objective function therefore becomes:

$$\begin{aligned} \mathcal{L}_\beta(\theta_1, \phi_1, \theta_2, \phi_2) = & \mathbb{E}_{\mathbf{z}_1 \sim q_{\phi_1}(\cdot|\mathbf{x})} [\ln p_{\theta_1}(\mathbf{x}_1|\mathbf{z}_1)] + \mathbb{E}_{\mathbf{z}_2 \sim q_{\phi_2}(\cdot|\mathbf{x})} [\ln p_{\theta_2}(\mathbf{x}_2|\mathbf{z}_2)] \\ & - \beta D_{KL}(q_{\phi_1}(\cdot|\mathbf{x}_1)q_{\phi_2}(\cdot|\mathbf{x}_2) \parallel p_\theta(\cdot)). \end{aligned} \quad (5)$$

2.4 Joint prior

There are assumed to be n_i latent variables in the VAE for view i , represented by the random vector $\mathbf{z}_i \in \mathbb{R}^{n_i}$. Without loss of generality it is assumed that $n_1 \leq n_2$. The joint prior distribution of the random vector $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2) \in \mathbb{R}^{n_1+n_2}$ is assumed to be a multivariate normal with mean $\mathbf{0}$ and covariance matrix, Σ_C :

$$\Sigma_C = \begin{pmatrix} \mathbf{I}_1 & \mathbf{C}^T \\ \mathbf{C} & \mathbf{I}_2 \end{pmatrix} \quad (6)$$

where $\mathbf{C} \in \mathbb{R}^{n_2 \times n_1}$ is the cross-covariance matrix encapsulating the relationship between the two latent spaces. \mathbf{I}_i is the $n_i \times n_i$ identity matrix. As Σ_C is a covariance matrix, it needs to satisfy two conditions: symmetry and positive semi-definiteness. With the defined structure it is clearly symmetric and so it is sufficient to require Σ_C to be a positive semi-definite matrix to ensure it is a well-defined covariance matrix. For a covariance matrix with structure as defined in Eq. 6, it is useful to note that Σ_C is positive definite if and only if $\mathbf{I}_1 - \mathbf{C}^T\mathbf{C}$ (or $\mathbf{I}_2 - \mathbf{C}\mathbf{C}^T$) is positive definite [16]. Additionally, as (a) a real symmetric matrix is positive (semi) definite if and only if all of its eigenvalues are positive (non-negative) [17, p. 51] and (b) a matrix is invertible if and only if all of its eigenvalues are non-zero, covariance matrix Σ_C (and $\mathbf{I}_1 - \mathbf{C}^T\mathbf{C}/\mathbf{I}_2 - \mathbf{C}\mathbf{C}^T$) is positive definite if and only if it is invertible.

2.5 Kullback–Leibler divergence term

As discussed earlier, for calculating the ELBO for the proposed prior, the KL divergence between $q_{\phi_1}(\cdot|\mathbf{x}_1)q_{\phi_2}(\cdot|\mathbf{x}_2)$ and the prior on $(\mathbf{z}_1, \mathbf{z}_2)$, p_C , is required. The following general result on KL divergence between multivariate Gaussians is applied in this setting.

Between two k -dimensional multivariate Gaussians r and s with respective means $\boldsymbol{\mu}_r$ and $\boldsymbol{\mu}_s$ and respective covariance matrices Σ_r and Σ_s , the KL divergence is given by [18]:

$$D_{KL}(s||r) = \frac{1}{2} \left[\ln \frac{|\Sigma_r|}{|\Sigma_s|} - k + (\boldsymbol{\mu}_s - \boldsymbol{\mu}_r)^T \Sigma_r^{-1} (\boldsymbol{\mu}_s - \boldsymbol{\mu}_r) + \text{tr} \{ \Sigma_r^{-1} \Sigma_s \} \right]. \quad (7)$$

As this assumes the existence of Σ_r^{-1} and Σ_s^{-1} , both covariance matrices must be positive definite.

By utilising the result of Eq. 7 with the prior over $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2)$ set as $r = p_C$, and the approximate posterior distribution as $s = q_{\phi_1}(\cdot|\mathbf{x}_1)q_{\phi_2}(\cdot|\mathbf{x}_2)$, the KL term of Eq. 2 is obtained. Explicitly, $p_C = N(\mathbf{0}, \Sigma_C)$ and $q_{\phi_1}(\cdot|\mathbf{x}_1)q_{\phi_2}(\cdot|\mathbf{x}_2) = N((\boldsymbol{\mu}_1(\mathbf{x}_1), \boldsymbol{\mu}_2(\mathbf{x}_2)), \Sigma_q)$ with $\Sigma_q = \text{bdiag}(\Sigma_i)$. For the utilisation of Eq. 7, it is assumed that Σ_C is positive definite and the variances of the approximate posterior are positive.

Using the block matrix inverse result from [19, p. 18] (stated in Appendix A.1.1) the inverse of Σ_C is given by:

$$\Sigma_C^{-1} = \begin{pmatrix} \mathbf{D}_1 & -\mathbf{D}_1\mathbf{C}^T \\ -\mathbf{D}_2\mathbf{C} & \mathbf{D}_2 \end{pmatrix} \quad (8)$$

where $\mathbf{D}_1 = (\mathbf{I}_1 - \mathbf{C}^T\mathbf{C})^{-1}$ and $\mathbf{D}_2 = (\mathbf{I}_2 - \mathbf{C}\mathbf{C}^T)^{-1}$. These inverses are guaranteed to exist, given the assumption of positive definiteness on Σ_C . Using Eq. 8 and properties of the trace, we obtain:

$$\text{tr} \{ \Sigma_C^{-1} \Sigma_q \} = \text{tr} \left\{ \begin{pmatrix} \mathbf{D}_1 \Sigma_1 & -\mathbf{D}_1 \mathbf{C}^T \Sigma_2 \\ -\mathbf{D}_2 \mathbf{C} \Sigma_1 & \mathbf{D}_2 \Sigma_2 \end{pmatrix} \right\} = \text{tr} \{ \mathbf{D}_1 \Sigma_1 \} + \text{tr} \{ \mathbf{D}_2 \Sigma_2 \}.$$

Additionally, as Σ_C is a block matrix, $|\Sigma_C| = |\mathbf{I}_1 - \mathbf{C}^T\mathbf{C}| = 1/|\mathbf{D}_1| = 1/|\mathbf{D}_2|$ [20, p. 114]. Similarly, as Σ_q is a diagonal matrix we have $|\Sigma_q| = |\Sigma_1||\Sigma_2|$. Writing $\boldsymbol{\mu}_i = \boldsymbol{\mu}_i(\mathbf{x}_i)$ and noting

that $\mu_p = 0$, we obtain:

$$\begin{aligned}
D_{KL}(q_{\phi_1}(\cdot|\mathbf{x}_1)q_{\phi_2}(\cdot|\mathbf{x}_2)||p_C) &= \frac{1}{2} [\boldsymbol{\mu}_1^T \mathbf{D}_1 \boldsymbol{\mu}_1 - \ln |\boldsymbol{\Sigma}_1| - n_1 + \text{tr} \{\mathbf{D}_1 \boldsymbol{\Sigma}_1\}] \\
&+ \frac{1}{2} [\boldsymbol{\mu}_2^T \mathbf{D}_2 \boldsymbol{\mu}_2 - \ln |\boldsymbol{\Sigma}_2| - n_2 + \text{tr} \{\mathbf{D}_2 \boldsymbol{\Sigma}_2\}] \\
&- \frac{1}{2} [\ln |\mathbf{D}_1| + \boldsymbol{\mu}_1^T \mathbf{D}_1 \mathbf{C}^T \boldsymbol{\mu}_2 + \boldsymbol{\mu}_2^T \mathbf{D}_2 \mathbf{C} \boldsymbol{\mu}_1].
\end{aligned} \tag{9}$$

2.6 Matrix parameterization

The matrix \mathbf{C} is unknown and therefore must be either chosen apriori or be optimised. In our work, \mathbf{C} is updated along with all other parameters at each update step, using the ELBO function as the loss function. By doing this, two challenges were faced and addressed. Firstly, \mathbf{C} must be updated in such a way that $\boldsymbol{\Sigma}_C$ is a valid covariance matrix. Secondly, a differentiable parameterization of \mathbf{C} is needed to allow for end-to-end learning. Conditions for validity of the update step are outlined in the following theorems.

Theorem 2.1. $\boldsymbol{\Sigma}_C$ defined as in Eq. 6 is positive semi-definite if and only if all singular values of \mathbf{C} are bounded by 1.

Proof. Due to previous remarks, it is equivalent to show that $\mathbf{I}_1 - \mathbf{C}^T \mathbf{C}$ is positive semi-definite if and only if the singular values of \mathbf{C} are bounded by 1. Further, this is equivalent to the showing that eigenvalues of $\mathbf{I}_1 - \mathbf{C}^T \mathbf{C}$ are non-negative if and only if the eigenvalues of \mathbf{C} are bounded by 1.

The eigenvalues of $\mathbf{C}^T \mathbf{C}$ are given by $\{\sigma_k^2(\mathbf{C})\}_{k=1}^{n_1}$ [21] meaning the eigenvalues of $-\mathbf{C}^T \mathbf{C}$ are given by $\{-\sigma_k^2(\mathbf{C})\}_{k=1}^{n_1}$. Applying Weyl's Theorem [19, p. 242] (stated in Appendix A.1.2) to $\mathbf{I}_1 - \mathbf{C}^T \mathbf{C}$ implies:

$$\begin{aligned}
\lambda_1(\mathbf{I}_1) + \lambda_{n_1+1-k}(-\mathbf{C}^T \mathbf{C}) &\leq \lambda_k(\mathbf{I}_1 - \mathbf{C}^T \mathbf{C}) \leq \lambda_{n_1}(\mathbf{I}_1) + \lambda_{n_1+1-k}(-\mathbf{C}^T \mathbf{C}) \\
\implies 1 &\leq \lambda_k(\mathbf{I}_1 - \mathbf{C}^T \mathbf{C}) - \lambda_{n_1+1-k}(-\mathbf{C}^T \mathbf{C}) \leq 1
\end{aligned} \tag{10}$$

which combined with the relations above gives $\lambda_k(\mathbf{I}_1 - \mathbf{C}^T \mathbf{C}) = 1 + \lambda_{n_1+1-k}(-\mathbf{C}^T \mathbf{C}) = 1 - \sigma_{n_1+1-k}^2(\mathbf{C})$. Therefore all eigenvalues of $\mathbf{I}_1 - \mathbf{C}^T \mathbf{C}$ are non-negative if and only if all singular values of \mathbf{C} are bounded by 1. \square

If the inequality on $\sigma_1(\mathbf{C})$ is replaced by a strict inequality then $\boldsymbol{\Sigma}_C$ is guaranteed to be positive definite. This is clear as all eigenvalues of $\mathbf{I}_1 - \mathbf{C}^T \mathbf{C}$ are now positive which ensures positive definiteness. This guarantees the applicability of Eq. 9. A further restriction, outlined in Theorem 2.2, can be made which enables a different parameterization of \mathbf{C} which assumes a scaled orthogonal relationship between the views.

Theorem 2.2. If $\mathbf{C} = \alpha \tilde{\mathbf{C}} \in \mathbb{R}^{n_2 \times n_1}$ where $|\alpha| \leq 1$ and $\tilde{\mathbf{C}}$ is a semi-orthogonal matrix ($\tilde{\mathbf{C}}^T \tilde{\mathbf{C}} = \mathbf{I}_1$) then $\boldsymbol{\Sigma}_C$ is positive semi-definite.

Proof. Given the assumed structure on \mathbf{C} , we have the following:

$$\mathbf{I}_1 - \mathbf{C}^T \mathbf{C} = \mathbf{I}_1 - \alpha^2 \tilde{\mathbf{C}}^T \tilde{\mathbf{C}} = (1 - \alpha^2) \mathbf{I}_1. \tag{11}$$

As \mathbf{I}_1 is positive semi definite, if $|\alpha| \leq 1$ then, so too is $\mathbf{I}_1 - \mathbf{C}^T \mathbf{C}$. \square

Alternatively, this could be seen as a corollary to Theorem 2.1 - due to the semi-orthogonality of $\tilde{\mathbf{C}}$, all singular values of $\tilde{\mathbf{C}}$ are equal to 1 which means $\sigma_i(\mathbf{C}) = \alpha$ with $|\alpha| \leq 1$. Again, if the condition on α is replaced with a strict inequality ($|\alpha| < 1$) this guarantees applicability of Eq. 9. A simplification of Eq. 9 assuming this orthogonality condition can be found in Appendix A.2.

To ensure applicability of Eq. 9, we can either require (a) $\mathbf{C} = \alpha \tilde{\mathbf{C}}$ with $\tilde{\mathbf{C}}$ an orthogonal matrix and $|\alpha| < 1$ or (b) $\sigma_1(\mathbf{C}) < 1$. The latter requires a matrix factorisation of \mathbf{C} which includes singular values. The most obvious approach is to use a singular value decomposition (SVD) of

C i.e. $C = USV^T$ where U and V are orthogonal matrices and S is the matrix of singular values i.e. $S = \text{diag}(\sigma_k(C))$. These singular values can be parameterized by $\sigma_k(C) = (1 - \exp(-\sigma_k))/(1 + \exp(-\sigma_k))$ for $\sigma_k \in \mathbb{R}$. To fully parameterize C with the singular value constraint, a parameterization of orthogonal matrices U and V is needed. Similarly, a parameterization of orthogonal \tilde{C} is needed for option (a). The following assumes $n_1 = n_2 = n$.

As discussed by Shepard et al. [22], there are four popular parameterizations of orthogonal matrices. All methods in their preliminary form parameterize at most a subset of the orthogonal matrices (at maximum those with determinant +1 or -1). These parameterizations must therefore be extended to map to the entire orthogonal matrix space. As it is the only one-to-one mapping, the rational Cayley transform has been chosen for use within JPVAE. In its original form, the Cayley transform tells us that for all orthogonal matrices with no eigenvalues equal to 1, there exists a unique skew-symmetric matrix $A \in \mathbb{R}^{m \times m}$ such that $O = (I + A)(I - A)^{-1}$ [22]. To extend the parameterization to the full space of orthogonal matrices, an extra matrix consisting of 1s and -1s along the diagonal is needed [23, 24]. Let $J = \text{diag}([\mathbf{1}_{m-r}; -\mathbf{1}_r])$ where $r = \text{floor}(m/(1 + \exp(-s)))$ for $s \in \mathbb{R}$. Then $O = J(I + A)(I - A)^{-1}$ is a mapping onto the space of orthogonal matrices. \tilde{C} is therefore parameterized in full by $(n - 1)^2/2 + 1$ parameters: $(\{a_{ij} \in \mathbb{R} : i < j \text{ for } i, j \in \{1, \dots, n\}\}, \{s \in \mathbb{R}\})$.

2.7 Imputation

Once the model has been learned on the training data, missing views can be imputed. Assume data is available for view j , but not view i , given x_j we can impute the missing value of x_i . Data x_j is fed into the encoder for view j , $E_{j\phi}$, and latent variables are sampled giving $z_j = a$. An estimate of z_i can be obtained using the conditional distribution of z_i given z_j . This is fed through decoder i , $D_{i\theta}$, to produce an estimate of x_i given x_j , $\tilde{x}_{i|j}$.

The joint marginal of z_1 and z_2 is assumed to be a multivariate normal with mean $[\mu_1; \mu_2]$ and covariance matrix Σ . The maximum likelihood estimator for the mean and covariance are obtained and denoted by $[\hat{\mu}_1; \hat{\mu}_2]$ and $\hat{\Sigma}$ respectively. As any subset of variables from a multivariate normal conditioned on a known second subset of variables also follows a multivariate normal distribution, the conditional distribution can be found explicitly. For $i \neq j$, the distribution of z_i given $z_j = a$ is observed is:

$$z_i | z_j = a \sim N\left(\hat{\mu}_i + \hat{\Sigma}_{ij} \hat{\Sigma}_{jj}^{-1} (a - \hat{\mu}_j), \hat{\Sigma}_{ii} - \hat{\Sigma}_{ij} \hat{\Sigma}_{jj}^{-1} \hat{\Sigma}_{ji}\right) \quad (12)$$

where $\hat{\Sigma}_{kl} \in \mathbb{R}^{n_k \times n_l}$ is the submatrix of $\hat{\Sigma}$ corresponding to the variables associated with view k and view l . The conditional mean $\mathbb{E}(z_i | z_j = a) = \hat{\mu}_i + \hat{\Sigma}_{ij} \hat{\Sigma}_{jj}^{-1} (a - \hat{\mu}_j)$ is then used as an estimate of the latent variables in latent space i and fed into $D_{i\theta}$ to obtain $\tilde{x}_{i|j}$, the imputed value of x_i given x_j .

3 Numerical experiments

Through a series of experiments the performance of JPVAE is explored for both imputation purposes as well as for downstream analyses like classification. As JPVAE enables imputation of missing views ($\tilde{X}_{i|j}$), this ability is investigated alongside reconstruction of views (\tilde{X}_i).

A multi-view dataset was created from the binary version of the popular MNIST dataset, which consists of handwritten digits from 0 to 9 [25]. For each image, the top half was taken as view 1 and the bottom half was taken as view 2. This dataset of halved images is referred to as hvdMNIST. The hvdMNIST dataset has the desirable property of having a strong correlation between views. The dataset contains 50,000 training images and 10,000 test images. Experiments are repeated with 5 different random seeds and the average and standard deviation reported. Details on the model architecture and training details can be found in Appendices A.3 and A.4 respectively.

Two variants of JPVAE are explored, where in each one the validity of Σ_C is enforced via different ways. The first variant imposes a bound on singular values (such that $\sigma_1(C) < 1$). The second variant enforces a scaled orthogonality where $CC^T = C^T C = \alpha^2 I$, with the value of α set as 0.95. The two variants of JPVAE are compared with the $C = 0$ case, which corresponds to completely disjoint VAEs for each view.

Without explicitly learning a correlation structure, correlation is present between the two generated latent spaces. By incorporating a joint prior with a non-zero cross-correlation as presented in Eq. 6

into the loss function, JPVAE increases the correlation between views in the latent space as shown in Figure 4. Table 1 illustrates the improvement in the ability to reconstruct view 2 given view 1 and vice versa when correlation structure is learnt. Enforcing the orthogonality restriction improves the performance of JPVAE compared with simply applying the singular value bound.

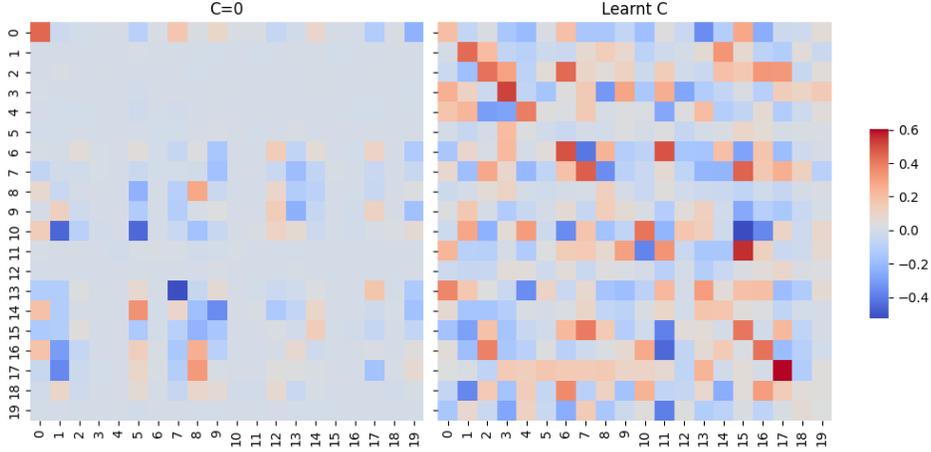


Figure 4: Empirical cross-correlation between view 1 and view 2 in the latent spaces. The left plot represents empirical cross-correlation for $C = \mathbf{0}$ and the right shows the same for C learnt with the orthogonality restriction imposed. The Frobenius norm of these matrices are 1.703 and 3.448 respectively.

Table 1: Average reconstruction losses across the entire dataset. Best results in bold, standard deviation in brackets.

	Reconstruction		Imputation	
	\tilde{X}_1	\tilde{X}_2	$\tilde{X}_{1 2}$	$\tilde{X}_{2 1}$
$C = \mathbf{0}$	24.64 (0.37)	25.56 (0.24)	114.1 (2.3)	127.5 (3.2)
$\sigma_1(C) < 1$	24.08 (0.44)	25.02 (0.21)	106.6 (1.4)	117.4 (4.1)
$CC^T = C^T C = \alpha^2 I$	23.41 (0.29)	23.98 (0.31)	97.25 (1.9)	106.6 (1.9)

The improved performance of the method is enabled by the learnt correlation that prevents posterior collapse. Following the work by [26], the phenomenon of posterior collapse is indicated by the percentage of active units (AU). The activity of unit u in the latent space is measured by $A_u = \text{Cov}_{\mathbf{x}}(\mathbb{E}_{u \sim q(u|\mathbf{x})}[u])$. A unit is considered active, *i.e.* to not have suffered from posterior collapse, if $A_u \geq 10^{-2}$. A higher percentage of AUs are preserved when C is learnt, with the best case scenario observed with the orthogonality constraint (Table 2). Wang et al. [12] proved that posterior collapse is equivalent to latent variable non-identifiability. This indicates that by enforcing the orthogonality restriction, we make the latent variable space identifiable.

Table 2: The percentage (%) of active units (AU) across latent spaces (z_1, z_2), out of a total 40. Best results in bold, standard deviation in brackets.

	AU
$C = \mathbf{0}$	61 (1.4)
$\sigma_1(C) < 1$	66 (1.4)
$CC^T = C^T C = \alpha^2 I$	98.5 (2.2)

The imputed views, $\tilde{X}_{1|2}$ and $\tilde{X}_{2|1}$, can be used for downstream tasks. If not all views are available for an individual, this allows techniques requiring all views to be applied. As an illustration of the performance of the imputed data in downstream tasks, a basic multi-layer perceptron classifier was

trained and tested on different combinations of reconstructed and imputed views. The performance of a classifier trained on the reconstructed data indicates that when C is learnt, the relevant signal remains present in both the reconstructed and the imputed data. The performance on imputed data is greatly improved by learning C , and in particular by enforcing the orthogonality constraint (Table 3; results with standard deviation can be found in Figure 9.).

Table 3: Results for (Y, Z) represent classification accuracy % for model trained on the training split of Y and tested on the test split of Z . Accuracies for (X_1, X_1) and (X_2, X_2) with standard deviation in brackets are 93.59% (0.25) and 90.83% (0.23) respectively. Best results in bold. For clarity, results with standard deviation reported can be found in Figure 9.

	View 1			View 2		
	$(\tilde{X}_1, \tilde{X}_1)$	$(\tilde{X}_1, \tilde{X}_{1 2})$	$(\tilde{X}_{1 2}, \tilde{X}_{1 2})$	$(\tilde{X}_2, \tilde{X}_2)$	$(\tilde{X}_2, \tilde{X}_{2 1})$	$(\tilde{X}_{2 1}, \tilde{X}_{2 1})$
$C = \mathbf{0}$	89.07	47.21	77.22	85.88	51.83	78.81
$\sigma_1(C) < 1$	89.67	64.50	80.22	86.46	67.06	82.38
$CC^T = C^T C = \alpha^2 I$	90.31	77.22	83.54	87.80	76.55	86.30

Not only does learning correlation structure improve the ability to impute data, the reconstruction loss and classification accuracy for reconstructed data x_i is improved compared with those scores seen when training each view on separate VAEs (Table 1). This may be due to the increased use of the latent space, as evidenced by the higher percentage of active units. Whilst a classifier trained on the imputed data from $C = \mathbf{0}$ demonstrates the retention of signal, the low accuracy seen for a classifier trained on reconstructed data and high reconstruction loss indicates that it does not retain the specific signature of the input. For example, taking the top half of a digit ‘2’ as the input, it may correctly reconstruct the bottom half of a realisation of the digit ‘2’ but not correctly reconstruct the specific realisation (as seen in Figure 1a).

Using the joint prior we see that the view with stronger signal (view 1) is bolstering the classification of the view with weaker signal (view 2). Whilst classification accuracy on x_1 is higher than that on x_2 , the accuracy on the imputed data $\tilde{X}_{2|1}$ sees a smaller drop, and is greater than that of $\tilde{X}_{1|2}$. Notably, the accuracy on imputed data $(\tilde{X}_{2|1}, \tilde{X}_{2|1})$ is comparable to that on $(\tilde{X}_2, \tilde{X}_2)$ whilst the same for view 1 experiences a drop in performance.

4 Conclusions

A novel multi-view VAE approach has been proposed that natively strengthens the correlation between latent spaces via a joint prior. This is the first time that a connection between multi-view VAEs is made through a joint prior rather than a joint posterior, as has previously been implemented in the literature. Theoretical guarantees and parameterizations are presented that allow for end-to-end learning. By simultaneously preventing posterior collapse, JPVAE returns superior models and demonstrates a promising ability to impute missing data suitable for downstream tasks.

References

- [1] Javier E. Flores, Daniel M. Claborne, Zachary D. Weller, Bobbie-Jo M. Webb-Robertson, Katrina M. Waters, and Lisa M. Bramer. Missing data in multi-omics integration: Recent advances through artificial intelligence. *Frontiers in Artificial Intelligence*, 6, 2023. ISSN 2624-8212. doi: 10.3389/frai.2023.1098308. URL <https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2023.1098308>.
- [2] Dongdong Lin, Jigang Zhang, Jingyao Li, Chao Xu, Hong-Wen deng, and Yu-Ping Wang. An integrative imputation method based on multi-omics datasets. *BMC Bioinformatics*, 17, 06 2016. doi: 10.1186/s12859-016-1122-6.
- [3] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations*, 2014. URL <http://arxiv.org/abs/1312.6114>.
- [4] Meng Song, Jonathan Greenbaum, Joseph Luttrell, Weihua Zhou, Chong Wu, Hui Shen, Ping Gong, Chaoyang Zhang, and Hong-Wen Deng. A review of integrative imputation for multi-omics datasets. *Frontiers in Genetics*, 11, 2020.
- [5] Thomas Sutter, Imant Daunhawer, and Julia Vogt. Multimodal generative learning utilizing jensen-shannon-divergence. *Advances in neural information processing systems*, 33:6100–6110, 2020.
- [6] Yuge Shi, N. Siddharth, Brooks Paige, and Philip H. S. Torr. Variational mixture-of-experts autoencoders for multi-modal deep generative models. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019.
- [7] Mike Wu and Noah Goodman. Multimodal generative models for scalable weakly-supervised learning. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, page 5580–5590, 2018.
- [8] Noah Cohen Kalafut, Xiang Huang, and Daifeng Wang. Joint variational autoencoders for multimodal imputation and embedding. *Nature Machine Intelligence*, 5:631–642, 2023.
- [9] Imant Daunhawer, Thomas M. Sutter, Kieran Chin-Cheong, Emanuele Palumbo, and Julia E Vogt. On the limitations of multimodal VAEs. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=w-CPUXXrAj>.
- [10] Weiran Wang, Xinchun Yan, Honglak Lee, and Karen Livescu. Deep variational canonical correlation analysis. *arXiv preprint arXiv:1610.03454*, 2016.
- [11] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *38th International Conference on Machine Learning, ICML 2021*, pages 12310–12320. ML Research Press, 2021.
- [12] Yixin Wang, David Blei, and John P Cunningham. Posterior collapse and latent variable non-identifiability. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 5443–5455, 2021.
- [13] S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79 – 86, 1951. doi: 10.1214/aoms/1177729694.
- [14] Hao Fu, Chunyuan Li, Xiaodong Liu, Jianfeng Gao, Asli Celikyilmaz, and Lawrence Carin. Cyclical annealing schedule: A simple approach to mitigating KL vanishing. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 240–250. Association for Computational Linguistics, 2019. doi: 10.18653/v1/N19-1021. URL <https://aclanthology.org/N19-1021>.
- [15] Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew M. Dai, Rafal Józefowicz, and Samy Bengio. Generating sentences from a continuous space. In *Conference on Computational Natural Language Learning*, 2015.

- [16] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*, chapter Mathematical background, page 650. Cambridge University Press, 2004.
- [17] K. B. Petersen and M. S. Pedersen. The matrix cookbook, Nov 2012. URL <http://www2.compute.dtu.dk/pubdb/pubs/3274-full.html>. Version 20121115.
- [18] L. Pardo. *Statistical Inference Based on Divergence Measures*. Statistics: A Series of Textbooks and Monographs. CRC Press, 2018. ISBN 9781420034813.
- [19] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 2 edition, 2012.
- [20] K.M. Abadir and J.R. Magnus. *Matrix Algebra*. Econometric Exercises. Cambridge University Press, 2005. ISBN 9781139443647.
- [21] James W. Demmel. *Applied Numerical Linear Algebra*, chapter Linear Least Squares Problems, pages 110–111. Society for Industrial and Applied Mathematics, 1997. doi: 10.1137/1.9781611971446.ch3.
- [22] Ron Shepard, Scott R. Brozell, and Gergely Gidofalvi. The representation and parametrization of orthogonal matrices. *The Journal of Physical Chemistry A*, 119(28):7924–7939, 07 2015.
- [23] W.L. Ferrar. *Algebra: a Textbook of Determinants, Matrices and Algebraic Forms*. Oxford University Press, 1960. URL <https://archive.org/details/algebra032104mbp/page/n3/mode/2up>.
- [24] A.I. Khuri and I.J. Good. The parameterization of orthogonal matrices: a review mainly for statisticians. *South African Statistical Journal*, 23, 1989.
- [25] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.
- [26] Yuri Burda, Roger Baker Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. *ICLR*, 2015.
- [27] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. URL <https://arxiv.org/abs/1412.6980>.

A Appendix / supplemental material

The Appendix contains additional results and figures to supplement the main body of text.

Theoretical results used within the manuscript are presented in Section A.1. This is followed by a simplification of the KL term when orthogonality constraints are assumed which is outlined in Section A.2. The model architecture and training details implemented in the numerical experiments are detailed in Sections A.4 and A.3 respectively. Lastly, additional results from the conducted numerical experiments can be found in Section A.5. All results throughout this manuscript are given to 4 significant figures, with standard deviations reported to 2 significant figures.

A.1 Theoretical results

Theoretical results utilised within the paper are now presented.

A.1.1 Block matrix inverse

The inverse of a block matrix, which is needed in Section 2.5 to determine the inverse of Σ_C in order to derive the KL term, is now outlined.

For a matrix \mathbf{A} with block matrix structure

$$\begin{pmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{pmatrix} \tag{13}$$

the following result holds [19, p. 18]:

Lemma A.1. *Assuming all relevant inverses exist, the inverse of \mathbf{A} as defined in Eq. 13 is given by:*

$$\begin{pmatrix} (\mathbf{A}_{11} - \mathbf{A}_{12}\mathbf{A}_{22}^{-1}\mathbf{A}_{21})^{-1} & \mathbf{A}_{11}^{-1}\mathbf{A}_{12}(\mathbf{A}_{21}\mathbf{A}_{11}^{-1}\mathbf{A}_{12} - \mathbf{A}_{22})^{-1} \\ \mathbf{A}_{22}^{-1}\mathbf{A}_{21}(\mathbf{A}_{12}\mathbf{A}_{22}^{-1}\mathbf{A}_{21} - \mathbf{A}_{11})^{-1} & (\mathbf{A}_{22} - \mathbf{A}_{21}\mathbf{A}_{11}^{-1}\mathbf{A}_{12})^{-1} \end{pmatrix}. \quad (14)$$

Applying this with $\mathbf{A}_{ii} = \mathbf{I}_i$, $\mathbf{A}_{12} = \mathbf{C}^T$ and $\mathbf{A}_{21} = \mathbf{C}$ gives the result in the text.

A.1.2 Weyl's inequality

The following inequality is used within the proof of Theorem 2.1 and concerns the sequence of eigenvalues of matrices. In contrast to the rest of the paper, the eigenvalues are indexed in non-increasing order, not non-increasing order of magnitude. For matrix \mathbf{M} these are denoted by $\{\hat{\lambda}_j(\mathbf{M})\}_{j=1}^n$.

Weyl's inequality [19, p. 242, Corollary 4.3.15] says:

Lemma A.2. *Let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$ be symmetric matrices. The following inequality holds*

$$\hat{\lambda}_j(\mathbf{A}) + \hat{\lambda}_1(\mathbf{B}) \leq \hat{\lambda}_j(\mathbf{A} + \mathbf{B}) \leq \hat{\lambda}_j(\mathbf{A}) + \hat{\lambda}_n(\mathbf{B}), \quad j = 1, \dots, n \quad (15)$$

Notice that $\lambda_k(\mathbf{I}_1) = \hat{\lambda}_j(\mathbf{I}_1) = 1$ for all j, k , and that as $-\mathbf{C}^T\mathbf{C}$ is negative semi-definite all eigenvalues are non-positive. This means $\hat{\lambda}_j(-\mathbf{C}^T\mathbf{C}) = \lambda_{n-j}(-\mathbf{C}^T\mathbf{C})$. Applying this lemma with $\mathbf{A} = -\mathbf{C}^T\mathbf{C}$ and $\mathbf{B} = \mathbf{I}_1$ gives Eq. 10.

A.2 KL term simplification

Under the orthogonality assumption on \mathbf{C} , the KL term derived in Eq. 9 can be simplified. Specifically, if $\mathbf{C}\mathbf{C}^T = \mathbf{C}^T\mathbf{C} = \alpha^2\mathbf{I}$ for some $\alpha \in (0, 1)$ then $\mathbf{D}_1 = \mathbf{D}_2 = \gamma\mathbf{I}$ where $\mathbf{I} = \mathbf{I}_1 = \mathbf{I}_2$ and $\gamma = 1/(1 - \alpha^2)$. Therefore Eq. 9 reduces to:

$$\begin{aligned} D_{KL}(q_{\phi_1}(\cdot|\mathbf{x}_1)q_{\phi_2}(\cdot|\mathbf{x}_1)||p_{\mathbf{C}}) &= \frac{1}{2} [\gamma\boldsymbol{\mu}_1^T\boldsymbol{\mu}_1 - \ln|\boldsymbol{\Sigma}_1| - n + \gamma \text{tr}\{\boldsymbol{\Sigma}_1\}] \\ &+ \frac{1}{2} [\gamma\boldsymbol{\mu}_2^T\boldsymbol{\mu}_2 - \ln|\boldsymbol{\Sigma}_2| - n + \gamma \text{tr}\{\boldsymbol{\Sigma}_2\}] \\ &- \frac{1}{2} [n \ln|\gamma| + 2\gamma\boldsymbol{\mu}_1^T\mathbf{C}\boldsymbol{\mu}_2]. \end{aligned} \quad (16)$$

A.3 Model architecture

All encoders and decoders for the JPVAE, as well as the neural network used for classification, consist of two dense layers with 512 units each. The latent distribution layers consists of two 20 unit dense layers which each parameterize the mean and the log of the variances of 20 normal random variables. The decoders output layer parameterizes the distribution of a Bernoulli random variable for each pixel. The classifier output layer parameterizes the categorical distribution with 10 outcomes. All activation functions were ReLu.

A.4 Training details

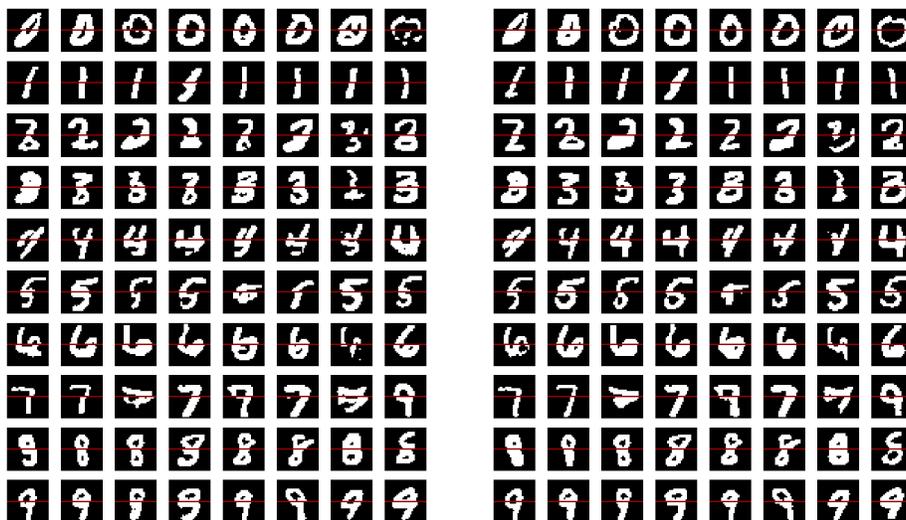
The JPVAE used an Adam optimiser [27] with learning rate 0.001, trained for 30 epochs with a batch size of 32 and used binary cross entropy as the reconstruction error. The cyclical KL annealing schedule as introduced by [14] is implemented, with $M = 30$.

The classifier is trained for 50 epochs with a batch size of 32, step size of 0.01, except for on original data where it is trained for 15 epochs to prevent overfitting. Cross entropy loss is used as the loss function.

A.5 Additional results

Additional results are presented in this section.

Figure 5 illustrates examples of reconstructing view 2 from view 1, with and without learning correlation natively, for a model trained with a different random seed to that displayed in Figure 1. Figures 6 and 7 illustrate examples of reconstructing view 1 from view 2, again with and without learning correlation natively, for models trained with different random seeds. As in Figure 1, the imputation of the missing view obtained when correlation of the latent spaces is learnt natively is of higher quality than when no correlation is learnt. To quantitatively evaluate this superior performance, a simple multi-layer perceptron classifier is trained the concatenation of the reconstructed views. It is then tested on a combination of the reconstructed and imputed views. Explicitly, for datasets where *e.g.* view 2 is imputed from view 1, the classifier is trained on (the training split of) $[\tilde{X}_1; \tilde{X}_2]$ and tested on (the testing split of) $[\tilde{X}_1; \tilde{X}_{2|1}]$. Results of this classification task for various test datasets are displayed in Figure 8. Results for the classification task described in the main body of the text, and presented in Table 3, are illustrated in Figure 9 with standard deviation incorporated via error bars.



(a) No correlation learnt, but estimated empirically after training

(b) Correlation learnt natively

Figure 5: Additional realisation of an imputation of the bottom half of MNIST digits (view 2 of the data) using the top half of the image (view 1) on a JPVAE model trained with (a) independent priors (completely separate VAEs) and (b) a joint prior with learnt correlation structure between latent spaces. The cross entropy loss between true bottom half of image and imputation is 111.9 in (a) and 101.5 in (b).



(a) No correlation learnt, but estimated empirically after training



(b) Correlation learnt natively

Figure 6: Imputation of the top half of MNIST digits (view 1 of the data) using the bottom half of the image (view 2) on a JPVAE model trained with (a) independent priors (completely separate VAEs) and (b) a joint prior with learnt correlation structure between latent spaces. The cross entropy loss between true top half of image and imputation is 117.8 in (a) and 100.2 in (b).



(a) No correlation learnt, but estimated empirically after training



(b) Correlation learnt natively

Figure 7: Additional realisation of an imputation of the top half of MNIST digits (view 1 of the data) using the bottom half of the image (view 2) on a JPVAE model trained with (a) independent priors (completely separate VAEs) and (b) a joint prior with learnt correlation structure between latent spaces. The cross entropy loss between true top half of image and imputation is 114.3 in (a) and 104.1 in (b).

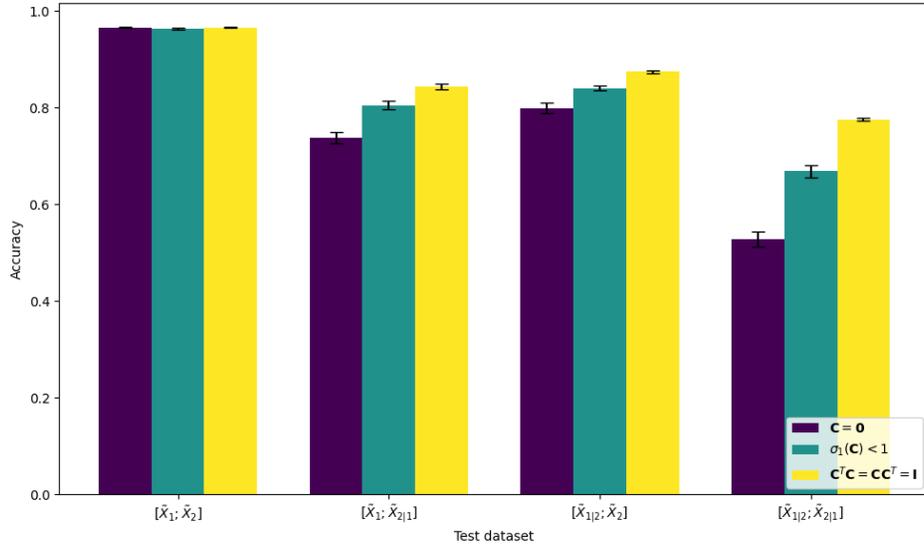


Figure 8: Results for $[Y_1; Y_2]$ represent classification accuracy % for model trained on the training split of $[\tilde{X}_1; \tilde{X}_2]$ and tested on the test split of $[Y_1; Y_2]$ (the column wise concatenation of Y_1 and Y_2). Accuracy for $[X_1; X_2]$ with standard deviation in brackets is 98.04% (0.074). Error bars present +/- one standard deviation.

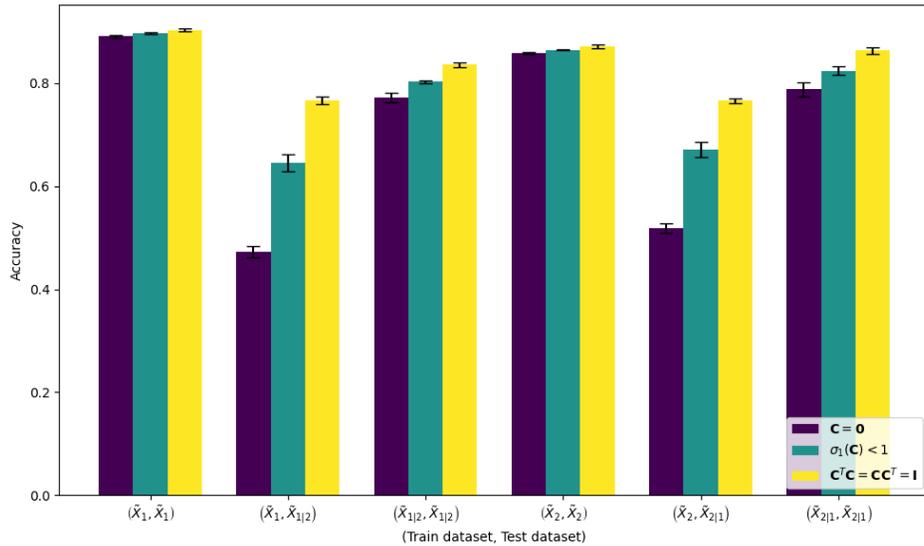


Figure 9: Results for (Y, Z) represent classification accuracy % for model trained on the training split of Y and tested on the test split of Z . Accuracies for (X_1, X_1) and (X_2, X_2) with standard deviation in brackets are 93.59% (0.25) and 90.83% (0.23) respectively. Error bars present +/- one standard deviation.