

GENERATIVE RESTRICTED KERNEL MACHINES

Anonymous authors

Paper under double-blind review

ABSTRACT

We introduce a novel framework for generative models based on Restricted Kernel Machines (RKMs) with multi-view generation and uncorrelated feature learning capabilities, called Gen-RKM. To incorporate multi-view generation, this mechanism uses a shared representation of data from various views. The mechanism is flexible to incorporate both kernel-based, (deep) neural network and convolutional based models within the same setting. To update the parameters of the network, we propose a novel training procedure which jointly learns the features and shared representation. Experiments demonstrate the potential of the framework through qualitative evaluation of generated samples.

1 INTRODUCTION

In the past decade, interest in generative models has grown tremendously, finding applications in multiple fields such as, generated art, on-demand video, image denoising (Vincent et al., 2010), exploration in reinforcement learning (Florensa et al., 2018), collaborative filtering (Salakhutdinov et al., 2007), inpainting (Yeh et al., 2017) and many more.

Some examples of graphical models based on a probabilistic framework with latent variables are Variational Auto-Encoders (Kingma & Welling, 2014) and Restricted Boltzmann Machines (RBMs) (Smolensky, 1986; Salakhutdinov & Hinton, 2009). More recently proposed models are based on adversarial training such as Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) and its many variants. Furthermore, auto-regressive models such as Pixel Recurrent Neural Networks (PixelRNNs) (Van Den Oord et al., 2016) model the conditional distribution of every individual pixel given previous pixels. All these approaches have their own advantages and disadvantages. For example, RBMs perform both learning and Bayesian inference in graphical models with latent variables. However, such probabilistic models must be properly normalized, which requires evaluating intractable integrals over the space of all possible variable configurations (Salakhutdinov & Hinton, 2009). Currently GANs are considered as the state-of-the-art for generative modeling tasks, producing high-quality images but are more difficult to train due to unstable training dynamics, unless more sophisticated variants are applied.

Many datasets are comprised of different representations of the data, or views. Views can correspond to different modalities such as sounds, images, videos, sequences of previous frames, etc. Although each view could individually be used for learning tasks, exploiting information from all views together could improve the learning quality (Pu et al., 2016; Liu & Tuzel, 2016; Chen & Denoyer, 2017). Also, it is among the goals of the latent variable modelling to model the description of data in terms of *uncorrelated or independent* components. Some classical examples are Independent Component Analysis; Hidden Markov models (Rabiner & Juang, 1986); Probabilistic Principal Component Analysis (PCA) (Tipping & Bishop, 1999); Gaussian-Process Latent variable model (Lawrence, 2005) and factor analysis. Hence, when learning a latent space in generative models, it becomes interesting to find a disentangled representation. Disentangled variables are generally considered to contain interpretable information and reflect separate factors of variation in the data (e.g. lighting conditions, style, colors, etc.). While the definition of disentanglement is not precise, many believe a representation with statistically independent variables is a good starting point (Schmidhuber, 1992; Ridgeway, 2016). Such representations extract information into a compact form which makes it possible to generate samples with specific characteristics (Chen et al., 2018; Bouchacourt et al., 2018; Tran et al., 2017; Chen et al., 2016). Additionally, these representations have been found to generalize better and be more robust against adversarial attacks (Alemi et al., 2017).

In this work, we propose an alternative generative mechanism based on the framework of Restricted Kernel Machines (RKMs) (Suykens, 2017), called Generative RKM (Gen-RKM). RKMs yield a representation of kernel methods with visible and hidden units establishing links between Kernel PCA, Least-Squares Support Vector Machines (LS-SVM) (Suykens et al., 2002) and RBMs. This framework has a similar energy form as RBMs, though there is a non-probabilistic training procedure where the eigenvalue decomposition plays the role of normalization. Recently, Houthuys & Suykens (2018) used this framework to develop tensor-based multi-view classification models and Schreurs & Suykens (2018) showed how kernel PCA fits into this framework.

Contributions. We make the following contributions: 1) A novel multi-view generative model based on the RKM framework where multiple views of the data can be generated simultaneously. 2) Two methods are proposed for computing the pre-image of the feature vectors: with the feature map explicitly known or unknown. We show that the mechanism is flexible to incorporate both kernel-based, (deep) convolutional neural network based models within the same setting. 3) When working with explicit feature maps, we propose a training algorithm that jointly performs the feature-selection and learns the common-subspace representation in the same procedure. 4) Experiments demonstrate that the model is capable of generating good quality images of natural objects. Further experiments on multi-view datasets exhibit the potential of the model. Thanks to the use of kernel PCA, the learned latent variables are uncorrelated. This resembles a disentangled representation, which makes it possible to generate data with specific characteristics.

This paper is organized as follows. In Section 2, we discuss the Gen-RKM training and generation mechanism when multiple data sources are available. In Section 3, we explain how the model incorporates both kernel methods and neural networks through the use of implicit and explicit feature maps respectively. When the feature maps are defined by neural networks, the Gen-RKM algorithm is explained in Section 4. In Section 5, we show experimental results of our model applied on various public datasets. Section 6 concludes the paper along with directions towards the future work.

2 GENERATIVE RESTRICTED KERNEL MACHINES FRAMEWORK

The proposed Gen-RKM framework consists of two phases: a training phase and a generation phase which occurs one after another.

2.1 TRAINING

Similar to Energy-Based Models (EBMs, see LeCun et al. (2004) for details), the RKM objective function captures dependencies between variables by associating a scalar energy to each configuration of the variables. Learning consists of finding an energy function in which the observed configurations of the variables are given lower energies than unobserved ones. Note that the schematic representation, as shown in Fig. 1 is similar to Discriminative RBMs (Larochelle & Bengio, 2008) and the objective function \mathcal{J}_t (defined below) has an energy form similar to RBMs with additional regularization terms.

We assume a dataset $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$, with $\mathbf{x}_i \in \mathbb{R}^d$, $\mathbf{y}_i \in \mathbb{R}^p$ comprising of N data points. Here \mathbf{y}_i may represent an additional view of \mathbf{x}_i , e.g., an additional image from a different angle, the caption of an image or a class label. Starting from the RKM interpretation of Kernel PCA, which gives an upper bound on the equality constrained L_2 Kernel PCA objective function (Suykens, 2017), and applying the feature-map $\phi_1 : \mathbb{R}^d \mapsto \mathbb{R}^{d_f}$ and $\phi_2 : \mathbb{R}^p \mapsto \mathbb{R}^{p_f}$ to the input data points, the training objective function \mathcal{J}_t for generative RKM is given by¹:

$$\mathcal{J}_t = \sum_{i=1}^N \left(-\phi_1(\mathbf{x}_i)^\top \mathbf{U} \mathbf{h}_i - \phi_2(\mathbf{y}_i)^\top \mathbf{V} \mathbf{h}_i + \frac{\lambda}{2} \mathbf{h}_i^\top \mathbf{h}_i \right) + \frac{\eta_1}{2} \text{Tr}(\mathbf{U}^\top \mathbf{U}) + \frac{\eta_2}{2} \text{Tr}(\mathbf{V}^\top \mathbf{V}) \quad (1)$$

where $\mathbf{U} \in \mathbb{R}^{d_f \times s}$ and $\mathbf{V} \in \mathbb{R}^{p_f \times s}$ are the unknown interaction matrices, and $\mathbf{h}_i \in \mathbb{R}^s$ are the latent variables modeling a common subspace \mathcal{H} between the two input spaces \mathcal{X} and \mathcal{Y} (see Fig. 1). The derivation of this objective function is given in the Appendix A.1. Given $\eta_1 > 0$ and $\eta_2 > 0$

¹For convenience, it is assumed that all the feature vectors are centered in the feature space \mathcal{F} using $\tilde{\phi}(\mathbf{x}) := \phi(\mathbf{x}) - \frac{1}{N} \sum_{i=1}^N \phi(\mathbf{x}_i)$. Otherwise, a centered kernel matrix could be obtained using Eq. 17 (Appendix A.4).




Figure 1: Gen-RKM schematic representation modeling a common subspace \mathcal{H} between two data sources \mathcal{X} and \mathcal{Y} . The ϕ_1, ϕ_2 are the feature maps (\mathcal{F}_x and \mathcal{F}_y represent the feature-spaces) corresponding to the two data sources. While ψ_1, ψ_2 represent the pre-image maps. The interconnection matrices \mathbf{U}, \mathbf{V} model dependencies between latent variables and the mapped data sources.

as regularization parameters, the stationary points of \mathcal{J}_t are given by:

$$\begin{cases} \frac{\partial \mathcal{J}_t}{\partial \mathbf{h}_i} = 0 \implies \lambda \mathbf{h}_i = \mathbf{U}^\top \phi_1(\mathbf{x}_i) + \mathbf{V}^\top \phi_2(\mathbf{y}_i), \quad \forall i = 1, \dots, N \\ \frac{\partial \mathcal{J}_t}{\partial \mathbf{U}} = 0 \implies \mathbf{U} = \frac{1}{\eta_1} \sum_{i=1}^N \phi_1(\mathbf{x}_i) \mathbf{h}_i^\top \\ \frac{\partial \mathcal{J}_t}{\partial \mathbf{V}} = 0 \implies \mathbf{V} = \frac{1}{\eta_2} \sum_{i=1}^N \phi_2(\mathbf{y}_i) \mathbf{h}_i^\top \end{cases} \quad (2)$$

Substituting \mathbf{U} and \mathbf{V} in the first equation above, denoting $\mathbf{\Lambda} = \text{diag}\{\lambda_1, \dots, \lambda_s\} \in \mathbb{R}^{s \times s}$ with $s \leq N$, yields the following eigenvalue problem:

$$\left[\frac{1}{\eta_1} \mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{K}_2 \right] \mathbf{H}^\top = \mathbf{H}^\top \mathbf{\Lambda}, \quad (3)$$

where $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_N] \in \mathbb{R}^{s \times N}$ with $s \leq N$ is the number of selected principal components and $\mathbf{K}_1, \mathbf{K}_2 \in \mathbb{R}^{N \times N}$ are the kernel matrices corresponding to data sources². Based on Mercer's theorem (Mercer, 1909), positive-definite kernel functions $k_1 : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$, $k_2 : \mathbb{R}^p \times \mathbb{R}^p \mapsto \mathbb{R}$ can be defined such that $k_1(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi_1(\mathbf{x}_i), \phi_1(\mathbf{x}_j) \rangle$, and $k_2(\mathbf{y}_i, \mathbf{y}_j) = \langle \phi_2(\mathbf{y}_i), \phi_2(\mathbf{y}_j) \rangle$, $\forall i, j = 1, \dots, N$ forms the elements of corresponding kernel matrices. The feature maps ϕ_1 and ϕ_2 , mapping the input data to the high-dimensional feature space (possibly infinite) are implicitly defined by kernel functions. Typical examples of such kernels are given by the Gaussian RBF kernel $k(\mathbf{x}_i, \mathbf{x}_j) = e^{-\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 / (2\sigma^2)}$ or the Laplace kernel $k(\mathbf{x}_i, \mathbf{x}_j) = e^{-\|\mathbf{x}_i - \mathbf{x}_j\|_2 / \sigma}$ just to name a few (Scholkopf & Smola, 2001). However, one can also define explicit feature maps, still preserving the positive-definiteness of the kernel function by construction (Suykens et al., 2002).

2.2 GENERATION

In this section, we derive the equations for the generative mechanism. RKMs resembling energy-based models, the inference consists in clamping the value of observed variables and finding configurations of the remaining variables that minimizes the energy (LeCun et al., 2004). Given the learned interconnection matrices \mathbf{U} and \mathbf{V} , and a given latent variable \mathbf{h}^* , consider the following objective function:

$$\mathcal{J}_g = -\phi_1(\mathbf{x}^*)^\top \mathbf{U} \mathbf{h}^* - \phi_2(\mathbf{y}^*)^\top \mathbf{V} \mathbf{h}^* + \frac{1}{2} \phi_1(\mathbf{x}^*)^\top \phi_1(\mathbf{x}^*) + \frac{1}{2} \phi_2(\mathbf{y}^*)^\top \phi_2(\mathbf{y}^*), \quad (4)$$

with an additional regularization term on data sources. The given latent variable \mathbf{h}^* can be the corresponding hidden variable of a training point, a newly sampled hidden unit or a specifically determined one. Above cases correspond to generating the reconstructed visible unit, generating a random new visible unit or exploring the latent space by carefully selecting hidden units respectively.

²While in the above section we have assumed that only two data sources (namely \mathcal{X} and \mathcal{Y}) are available for learning, the above procedure could be extended to multiple data-sources. For the M views or data-sources, this yields the training problem: $\left[\sum_{\ell=1}^M \frac{1}{\eta_\ell} \mathbf{K}_\ell \right] \mathbf{H}^\top = \mathbf{H}^\top \mathbf{\Lambda}$.

Here \mathcal{J}_g denotes the objective function for generation. The stationary points of \mathcal{J}_g are characterized by:

$$\begin{cases} \frac{\partial \mathcal{J}_g}{\partial \phi_1(\mathbf{x}^*)} = 0 \implies \phi_1(\mathbf{x}^*) = \mathbf{U}\mathbf{h}^*, \\ \frac{\partial \mathcal{J}_g}{\partial \phi_2(\mathbf{y}^*)} = 0 \implies \phi_2(\mathbf{y}^*) = \mathbf{V}\mathbf{h}^*. \end{cases} \quad (5)$$

Using \mathbf{U} and \mathbf{V} from Eq. 2, we obtain the generated feature vectors:

$$\phi_1(\mathbf{x}^*) = \left(\frac{1}{\eta_1} \sum_{i=1}^N \phi_1(\mathbf{x}_i) \mathbf{h}_i^\top \right) \mathbf{h}^*, \quad \phi_2(\mathbf{y}^*) = \left(\frac{1}{\eta_2} \sum_{i=1}^N \phi_2(\mathbf{y}_i) \mathbf{h}_i^\top \right) \mathbf{h}^*. \quad (6)$$

To obtain the generated data, one now needs to compute the inverse images of the feature maps $\phi_1(\cdot)$ and $\phi_2(\cdot)$ in the respective input spaces, i.e., solve the *pre-image problem*. We seek to find the functions $\psi_1: \mathbb{R}^{d_f} \mapsto \mathbb{R}^d$ and $\psi_2: \mathbb{R}^{p_f} \mapsto \mathbb{R}^p$ corresponding to the two data-sources, such that $(\psi_1 \circ \phi_1)(\mathbf{x}^*) \approx \mathbf{x}^*$ and $(\psi_2 \circ \phi_2)(\mathbf{y}^*) \approx \mathbf{y}^*$, where $\phi_1(\mathbf{x}^*)$ and $\phi_2(\mathbf{y}^*)$ are calculated using Eq. 6.

When using kernel methods, explicit feature maps are not necessarily known. Commonly used kernels such as the radial-basis function and polynomial kernels map the input data to a very high dimensional feature space. Hence finding the pre-image, in general, is known to be an ill-conditioned problem (Mika et al., 1999). However, various approximation techniques have been proposed (Bui et al., 2019; Kwok & Tsang, 2003; Honeine & Richard, 2011; Weston et al., 2004) which could be used to obtain the approximate pre-image $\hat{\mathbf{x}}$ of $\phi_1(\mathbf{x}^*)$. In section 3.1, we employ one such technique to demonstrate the applicability in our model, and consequently generate the multi-view data. One could also define explicit pre-image maps. In section 3.2, we define parametric pre-image maps and learn the parameters by minimizing the appropriately defined reconstruction errors. The next section describes the above two pre-image methods for both cases, i.e., when the feature map is explicitly known or unknown, in greater detail.

3 IMPLICIT & EXPLICIT FEATURE MAP

3.1 IMPLICIT FEATURE MAP

As noted in the previous section, since \mathbf{x}^* may not exist, we find an approximation $\hat{\mathbf{x}}$. A possible technique is shown by Schreurs & Suykens (2018). Left multiplying Eq. 6 by $\phi_1(\mathbf{x}_i^*)^\top$ and $\phi_2(\mathbf{y}_i^*)^\top$, $\forall i = 1, \dots, N$, we obtain:

$$\mathbf{k}_{\mathbf{x}^*} = \frac{1}{\eta_1} \mathbf{K}_1 \mathbf{H}^\top \mathbf{h}^*, \quad \mathbf{k}_{\mathbf{y}^*} = \frac{1}{\eta_2} \mathbf{K}_2 \mathbf{H}^\top \mathbf{h}^*, \quad (7)$$

where, $\mathbf{k}_{\mathbf{x}^*} = [k(\mathbf{x}_1, \mathbf{x}^*), \dots, k(\mathbf{x}_N, \mathbf{x}^*)]^\top$ represents the *similarities* between $\phi_1(\mathbf{x}^*)$ and training data points in the feature space, and $\mathbf{K}_1 \in \mathbb{R}^{N \times N}$ represents the centered kernel matrix of \mathcal{X} . Similar conventions follow for \mathcal{Y} respectively. Using the *kernel-smoother* method (Hastie et al., 2001), the pre-images are given by:

$$\hat{\mathbf{x}} = \psi_1(\phi_1(\mathbf{x}^*)) = \frac{\sum_{j=1}^{n_r} \tilde{k}_1(\mathbf{x}_j, \mathbf{x}^*) \mathbf{x}_j}{\sum_{j=1}^{n_r} \tilde{k}_1(\mathbf{x}_j, \mathbf{x}^*)}, \quad \hat{\mathbf{y}} = \psi_2(\phi_2(\mathbf{y}^*)) = \frac{\sum_{j=1}^{n_r} \tilde{k}_2(\mathbf{y}_j, \mathbf{y}^*) \mathbf{y}_j}{\sum_{j=1}^{n_r} \tilde{k}_2(\mathbf{y}_j, \mathbf{y}^*)}, \quad (8)$$

where $\tilde{k}_1(\mathbf{x}_i, \mathbf{x}^*)$ and $\tilde{k}_2(\mathbf{y}_i, \mathbf{y}^*)$ are the scaled similarities (see Eq. 8) between 0 and 1 and n_r the number of closest points based on the similarity defined by kernels \tilde{k}_1 and \tilde{k}_2 .

3.2 EXPLICIT FEATURE MAP

While using an explicit feature map, Mercer’s theorem still holds due to the positive semi-definiteness of the kernel function by construction, thereby allowing the derivation of Eq. 3. In the experiments, we use a set of (convolutional) neural networks as the feature maps $\phi_\theta(\cdot)$. Another (transposed convolutional) neural network is used for the pre-image map $\psi_\zeta(\cdot)$ (Dumoulin & Visin, 2016). The network parameters $\{\theta, \zeta\}$ are learned by minimizing the reconstruction errors defined by $\mathcal{L}_1(\mathbf{x}_i^*, \psi_{1_{\zeta_1}}(\phi_{1_{\theta_1}}(\mathbf{x}_i^*)))$ and $\mathcal{L}_2(\mathbf{y}_i^*, \psi_{2_{\zeta_2}}(\phi_{2_{\theta_2}}(\mathbf{y}_i^*)))$. In our experiments, we use the mean-squared errors $\mathcal{L}_1(\mathbf{x}_i^*, \psi_{1_{\zeta_1}}(\phi_{1_{\theta_1}}(\mathbf{x}_i^*))) = \frac{1}{N} \|\mathbf{x}_i^* - \psi_{1_{\zeta_1}}(\phi_{1_{\theta_1}}(\mathbf{x}_i^*))\|_2^2$

and $\mathcal{L}_2(\mathbf{y}_i^*, \psi_{2\zeta_2}(\phi_{2\theta_2}(\mathbf{y}_i^*))) = \frac{1}{N} \|\mathbf{y}_i^* - \psi_{2\zeta_2}(\phi_{2\theta_2}(\mathbf{y}_i^*))\|_2^2$, however, in principle, one can use any other loss appropriate to the dataset. Here $\phi_{1\theta_1}(\mathbf{x}_i^*)$ and $\phi_{2\theta_2}(\mathbf{y}_i^*)$ are computed from Eq. 6, i.e., the generated points in feature space from the subspace \mathcal{H} .

Adding the loss function directly into the objective function \mathcal{J}_t is not suitable for minimization. Instead, we use the stabilized objective function defined as $\mathcal{J}_{stab} = \mathcal{J}_t + \frac{c_{stab}}{2} \mathcal{J}_t^2$, where $c_{stab} \in \mathbb{R}^+$ is the regularization constant (Suykens, 2017). This tends to push the objective function \mathcal{J}_t towards zero, which is also the case when substituting the solutions λ_i, \mathbf{h}_i back into \mathcal{J}_t (see Appendix A.3 for details). The combined training objective is given by:

$$\min_{\theta_1, \theta_2, \zeta_1, \zeta_2} \mathcal{J}_c = \mathcal{J}_{stab} + \frac{c_{acc}}{2N} \left(\sum_{i=1}^N [\mathcal{L}_1(\mathbf{x}_i^*, \psi_{1\zeta_1}(\phi_{1\theta_1}(\mathbf{x}_i^*))) + \mathcal{L}_2(\mathbf{y}_i^*, \psi_{2\zeta_2}(\phi_{2\theta_2}(\mathbf{y}_i^*))) \right], \quad (9)$$

where $c_{acc} \in \mathbb{R}^+$ is a regularization constant to control the stability with reconstruction accuracy. In this way, we integrate feature-selection and subspace learning within the same training procedure.

4 THE GEN-RKM ALGORITHM

Based on the previous analysis, we propose a novel algorithm, called the Gen-RKM algorithm, combining kernel learning and generative models. We show that this procedure is efficient to train and evaluate. It is also scalable to large datasets when using explicit feature maps. The training procedure simultaneously involves feature selection, common-subspace learning and inverse-map learning. This is achieved via an optimization procedure where one iteration involves an eigendecomposition of the kernel matrix which is composed of the features from various views (see Eq. 3). The latent variables are given by the eigenvectors, which are then passed via a pre-image map to reconstruct the sample. The reconstruction error together with the energy function represents the cost that needs to be minimized. Fig. 1 shows a schematic representation of the algorithm when two data sources are available.

Thanks to training in m mini-batches, this procedure is scalable to large datasets (sample size N) with training time scaling super-linearly with $T_m = c \frac{N^\gamma}{m^{\gamma-1}}$, instead of $T_k = cN^\gamma$, where $\gamma \approx 3$ for algorithms based on decomposition methods, with some proportionality constant c . The training time could be further reduced by computing the covariance matrix (size $(d_f + p_f) \times (d_f + p_f)$) instead of a kernel matrix (size $\frac{N}{m} \times \frac{N}{m}$), when the sum of the dimensions of the feature-spaces is less than the samples in mini-batch i.e. $d_f + p_f \leq \frac{N}{m}$. While using neural networks as feature maps, d_f and p_f correspond to the number of neurons in the output layer, which are chosen as hyperparameters by the practitioner. Eigendecomposition of this smaller covariance matrix would yield \mathbf{U} and \mathbf{V} as eigenvectors (see Eq. 10 and Appendix A.2 for detailed derivation), where computing the \mathbf{h}_i involves only matrix-multiplication which is readily parallelizable on modern GPUs:

$$\begin{bmatrix} \frac{1}{\eta_1} \Phi_{\mathbf{x}} \Phi_{\mathbf{x}}^\top & \frac{1}{\eta_1} \Phi_{\mathbf{x}} \Phi_{\mathbf{y}}^\top \\ \frac{1}{\eta_2} \Phi_{\mathbf{y}} \Phi_{\mathbf{x}}^\top & \frac{1}{\eta_2} \Phi_{\mathbf{y}} \Phi_{\mathbf{y}}^\top \end{bmatrix} \begin{bmatrix} \mathbf{U} \\ \mathbf{V} \end{bmatrix} = \begin{bmatrix} \mathbf{U} \\ \mathbf{V} \end{bmatrix} \Lambda, \quad \begin{aligned} \Phi_{\mathbf{x}} &:= [\phi_1(\mathbf{x}_1), \dots, \phi_1(\mathbf{x}_N)], \\ \Phi_{\mathbf{y}} &:= [\phi_2(\mathbf{y}_1), \dots, \phi_2(\mathbf{y}_N)]. \end{aligned} \quad (10)$$

5 EXPERIMENTS

To demonstrate the applicability of the proposed framework and algorithm, we trained the Gen-RKM model on a variety of datasets commonly used to evaluate generative models: MNIST (LeCun & Cortes, 2010), Fashion-MNIST (Xiao et al., 2017), CIFAR-10 (Krizhevsky, 2009), CelebA (Liu et al., 2015) and Dsprites (Matthey et al., 2017). The experiments were performed using both the implicit feature map defined by a Gaussian kernel and parametric explicit feature maps defined by deep neural networks, either Convolutional or fully connected. As explained in Section 2, in case of kernel methods, training only involves constructing the kernel matrix and solving the eigenvalue problem in Eq. 3. In principle, one could also use the latent variables directly for generation. However, in our experiments, we fit a Gaussian mixture model (GMM) with l components to the latent variables of the training set, and randomly sample a new point \mathbf{h}^* for generating views using a kernel smoother. In case of explicit feature maps, we define $\phi_{1\theta_1}$ and $\psi_{1\zeta_1}$ as convolution and transposed-convolution neural networks, respectively (Dumoulin & Visin, 2016); and $\phi_{2\theta_2}$ and $\psi_{1\zeta_2}$

Algorithm 1 Gen-RKM**Input:** $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$, η_1 , η_2 , feature map $\phi_j(\cdot)$ - explicit *or* implicit via kernels $k_j(\cdot, \cdot)$, for $j \in \{1, 2\}$ **Output:** Generated data \mathbf{x}^* , \mathbf{y}^*

```

1: procedure TRAIN
2:   if  $\phi_j(\cdot)$  = Implicit then
3:     Hyperparameters: kernel specific
4:     Solve Eq. 3
5:     Select  $s$  principal components
6:   else if  $\phi_j(\cdot)$  = Explicit then
7:     while not converged do
8:        $\{\mathbf{x}, \mathbf{y}\} \leftarrow \{\text{Get mini-batch}\}$ 
9:        $\phi_1(\mathbf{x}) \leftarrow \mathbf{x}$ ;  $\phi_2(\mathbf{y}) \leftarrow \mathbf{y}$ 
10:      do steps 4-5
11:       $\{\phi_1(\mathbf{x}), \phi_2(\mathbf{y})\} \leftarrow h$  (Eq. 6)
12:       $\{\mathbf{x}, \mathbf{y}\} \leftarrow \{\psi_1(\phi_1(\mathbf{x})), \psi_2(\phi_2(\mathbf{y}))\}$ 
13:       $\Delta\theta_1 \propto -\nabla_{\theta_1} \mathcal{J}_c$ ;  $\Delta\theta_2 \propto -\nabla_{\theta_2} \mathcal{J}_c$ 
14:       $\Delta\zeta_1 \propto -\nabla_{\zeta_1} \mathcal{J}_c$ ;  $\Delta\zeta_2 \propto -\nabla_{\zeta_2} \mathcal{J}_c$ 
15:    end while
16:   end if
17: end procedure

1: procedure GENERATION
2:   Select  $h^*$ 
3:   if  $\phi_j(\cdot)$  = Implicit then
4:     Hyperparameter:  $n_r$ 
5:     Compute  $\mathbf{k}_{\mathbf{x}^*}$ ,  $\mathbf{k}_{\mathbf{y}^*}$  (Eq. 7)
6:     Get  $\hat{\mathbf{x}}$ ,  $\hat{\mathbf{y}}$  (Eq. 8)
7:   else if  $\phi_j(\cdot)$  = Explicit then
8:     do steps 11-12
9:   end if
10: end procedure

```

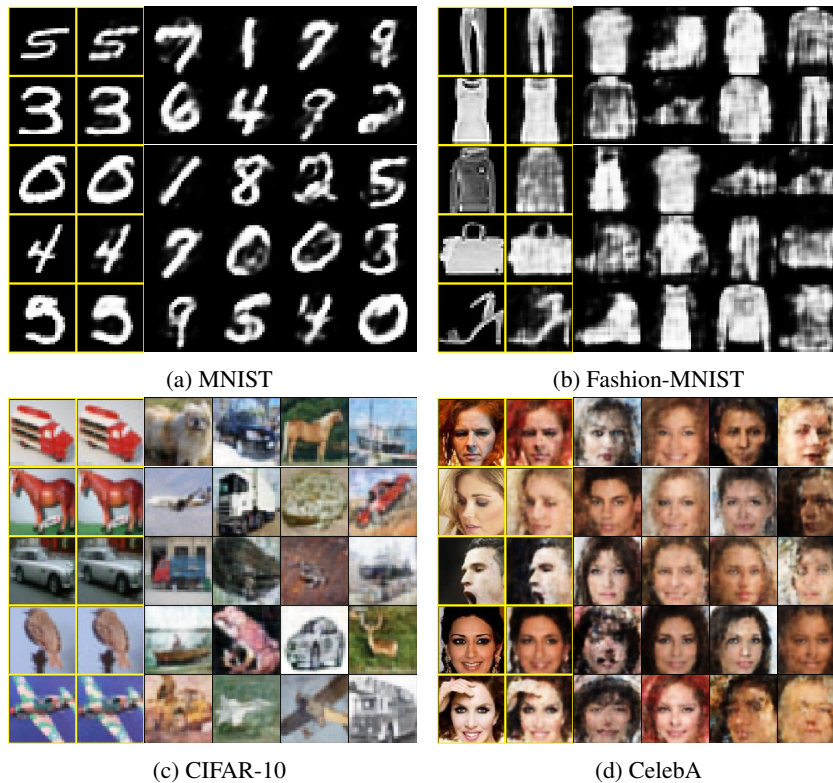


Figure 2: Generated samples from the model using CNN as explicit feature map in the kernel function. The yellow boxes in the first column show training examples and the adjacent boxes show the reconstructed samples. The other images (columns 3-6) are generated by random sampling from the fitted distribution over the learned latent variables.




Figure 3: Multi-view generation on CelebA dataset showing images and attributes.

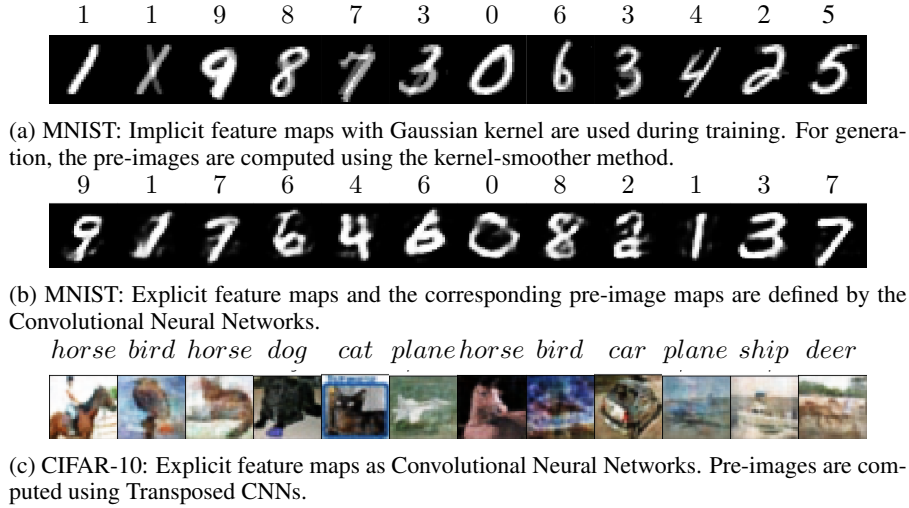


Figure 4: Multi-view Generation (images and labels) on various datasets using implicit and explicit feature maps.




Figure 5: Exploring the learned uncorrelated-features by traversing along the eigenvectors. The first column shows the scatter plot of latent variables using the top two principal components. The green lines within, show the traversal in the latent space and the related rows show the corresponding reconstructed images.

as fully-connected networks. The particular architecture details are outlined in Table 2 in the Appendix. The training procedure in case of explicitly defined maps consists of minimizing \mathcal{J}_c using the Adam optimizer (Kingma & Ba, 2014) to update the weights and biases. To speed-up learning, we subdivided the datasets into m mini-batches, and within each iteration of the optimizer, Eq. 3 is solved to update the value of \mathbf{H} . Information on the datasets and hyperparameters used for the experiments is given in Table 1 in the Appendix.

Generation: Figure 2a shows the generated images using a kernel smoother method. The first column in yellow-boxes shows the training samples and the second column on the right shows the reconstructed samples. The other images shown are generated by random sampling from a GMM

over the learned latent variables. Notice that the reconstructed samples are of better quality visually than the other images generated by random sampling. Figures 2b, 2c and 2d show the images generated when the convolutional neural network and transposed-convolutional neural network was used as the feature map and pre-image map respectively. To elucidate that the model has not merely memorized the training examples, we show the generated images via bilinear-interpolations of the latent variables in Appendix A.6.

Multi-view Generation: Figures 3 & 4 demonstrate the multi-view generative capabilities of the model. In these datasets, labels or attributes are seen as another view of the image that provides extra information. One-hot encoding of the labels was used to train the model. Figure 4a shows the generated images and labels when feature maps are only implicitly known i.e. through a Gaussian kernel. Figures 4b, 4c shows the same when using fully-connected networks as parametric functions to encode and decode labels. We can see that both the generated image and the generated label matches in most cases, albeit not all. Up to our knowledge, no universal evaluation metric exists to assess such characteristic of multi-view generation. Though one can use classifiers to crudely assess the matching, however, depending on the type of classifier and the way it was trained, the results would vary among researchers.

Targeted Generation: Since the components of the latent variables are the eigenvectors of the kernel matrix (see Eq. 3), one can exploit the orthogonality for targeted generation. Such targeted generation capabilities could be useful in critical applications where the data needs to be generated based on some prior-knowledge or with specific attributes. We explore the uncorrelated features learned by the models on the Dsprites and celebA dataset (See Fig. 5). In our experiments, the Dsprites training dataset comprised of 32×32 positions of oval and heart-shaped objects. The number of principal components chosen were 2 and the goal was to find-out whether traversing along the eigenvectors, corresponds to traversing the generated image in one particular direction while preserving the shape of the object. Rows 1 and 2 of Fig. 5 show the reconstructed images of an oval while moving along first and second principal component respectively. Notice that the first and second components correspond to the y and x positions respectively. Rows 3 and 4 show the same for hearts. On the celebA dataset, we train the Gen-RKM with 15 components. Rows 5 and 6 shows the reconstructed images while traversing along the principal components. When moving along the first component from left-to-right, the hair-color of the women transforms, while preserving the face structure. Whereas traversal along the second component, transforms a man to woman while preserving the orientation. When the number of principal components were 2 while training, the brightness and background light-source corresponds to the two largest variances in the dataset. Also notice that, the reconstructed images are more blurry due to the selection of less number of components to model \mathcal{H} .

6 CONCLUSION AND FUTURE WORK

The paper proposes a novel framework, called Gen-RKM, for generative models based on RKMs with extensions to multi-view generation and learning uncorrelated representations. This allows for a mechanism where the feature map can be defined using kernel functions or (deep) neural network based methods. When using kernel functions, the training consists of only solving an eigenvalue problem. In the case of a (convolutional) neural network based explicit feature map, we used (transposed) networks as the pre-image functions. Consequently, a training procedure was proposed which involves joint feature-selection and subspace learning. Thanks to training in mini-batches and capability of working with covariance matrices, the training is scalable to large datasets. Experiments on benchmark datasets illustrate the merit of the proposed framework. Furthermore, a targeted generation mechanism is demonstrated which uses the uncorrelated features modelled by the orthogonal eigenvectors. Extensions of this work consists of adapting the model to more advanced multi-view datasets involving speech, images and texts; further analysis on other feature maps, pre-image methods, loss-functions and uncorrelated feature learning. Finally, this paper has demonstrated the applicability of the Gen-RKM framework, suggesting new research directions to be worth exploring.

REFERENCES

- Alex Alemi, Ian Fischer, Josh Dillon, and Kevin Murphy. Deep variational information bottleneck. In *ICLR*, 2017.
- Diane Bouchacourt, Ryota Tomioka, and Sebastian Nowozin. Multi-level variational autoencoder: Learning disentangled representations from grouped observations. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Anh Tuan Bui, Joon-Ku Im, Daniel W. Apley, and George C. Runger. Projection-Free Kernel Principal Component Analysis for Denoising. *Neurocomputing*, 2019. ISSN 0925-2312.
- Mickaël Chen and Ludovic Denoyer. Multi-view generative adversarial networks. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 175–188. Springer, 2017.
- Tian Qi Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems*, pp. 2610–2620, 2018.
- Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pp. 2172–2180, 2016.
- Vincent Dumoulin and Francesco Visin. A guide to convolution arithmetic for deep learning. *arXiv preprint arXiv:1603.07285*, 2016.
- Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. Automatic Goal Generation for Reinforcement Learning Agents. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1515–1528, Stockholmströmssan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pp. 2672–2680, 2014.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer New York Inc., New York, NY, USA, 2001.
- Paul Honeine and Cedric Richard. Preimage Problem in Kernel-Based Machine Learning. *IEEE Signal Processing Magazine*, 28(2):77–88, March 2011. ISSN 1053-5888.
- Lynn Houthuys and Johan A K Suykens. Tensor-based Restricted Kernel Machines for Multi-View Classification. In *27th International Conference on Artificial Neural Networks ICANN, Rhodes, Greece*, volume 11140, pp. 205–215, 2018.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- James T. Kwok and Ivor Wai-Hung Tsang. The pre-image problem in kernel methods. *IEEE Transactions on Neural Networks*, 15:1517–1525, 2003.
- Hugo Larochelle and Yoshua Bengio. Classification using discriminative restricted Boltzmann machines. In *Proceedings of the 25th International Conference on Machine Learning - ICML '08*, pp. 536–543, Helsinki, Finland, 2008. ACM Press. ISBN 978-1-60558-205-4.

- Neil Lawrence. Probabilistic non-linear principal component analysis with gaussian process latent variable models. *JMLR*, 6:1783–1816, December 2005. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=1046920.1194904>.
- Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010. URL <http://yann.lecun.com/exdb/mnist/>.
- Yann LeCun, Fu Jie Huang, and Leon Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 2, pp. II–97–104 Vol.2, 2004.
- Ming-Yu Liu and Oncel Tuzel. Coupled Generative Adversarial Networks. In *Advances in Neural Information Processing Systems 29*, pp. 469–477. Curran Associates, Inc., 2016.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Loic Matthey, Irina Higgins, Demis Hassabis, and Alexander Lerchner. dsprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/>, 2017.
- James Mercer. Functions of Positive and Negative Type, and Their Connection the Theory of Integral Equations. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 209(441-458):415–446, January 1909.
- Sebastian Mika, Bernhard Schölkopf, Alex Smola, Klaus-Robert Müller, Matthias Scholz, and Gunnar Rätsch. Kernel PCA and De-noising in Feature Spaces. In *Proceedings of the 1998 Conference on Advances in Neural Information Processing Systems II*, pp. 536–542. MIT Press, 1999.
- Yunchen Pu, Zhe Gan, Ricardo Henao, Xin Yuan, Chunyuan Li, Andrew Stevens, and Lawrence Carin. Variational Autoencoder for Deep Learning of Images, Labels and Captions. NIPS’16, pp. 2360–2368. Curran Associates Inc., USA, 2016. ISBN 978-1-5108-3881-9.
- Lawrence R Rabiner and Biing-Hwang Juang. An introduction to hidden markov models. *IEEE ASSP magazine*, 3(1):4–16, 1986.
- Karl Ridgeway. A survey of inductive biases for factorial representation-learning. *CoRR*, abs/1612.05299, 2016.
- Ralph Tyrrell Rockafellar. *Conjugate Duality and Optimization*. SIAM, 1974.
- Ruslan Salakhutdinov and Geoffrey Hinton. Deep Boltzmann Machines. *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics*, Volume 5 of JMLR, 2009.
- Ruslan Salakhutdinov, Andriy Mnih, and Geoffrey Hinton. Restricted Boltzmann machines for collaborative filtering. In *ICML ’07*, pp. 791–798, Corvallis, Oregon, 2007. ACM Press.
- Jürgen Schmidhuber. Learning factorial codes by predictability minimization. *Neural Computation*, 4(6):863–879, 1992.
- Bernhard Scholkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001. ISBN 0262194759.
- Joachim Schreurs and Johan A. K. Suykens. Generative Kernel PCA. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, pp. 129–134, 2018.
- Paul Smolensky. Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1. chapter Information Processing in Dynamical Systems: Foundations of Harmony Theory, pp. 194–281. MIT Press, Cambridge, MA, USA, 1986. ISBN 0-262-68053-X.
- Johan A. K. Suykens. Deep Restricted Kernel Machines Using Conjugate Feature Duality. *Neural Computation*, 29(8):2123–2163, August 2017. ISSN 0899-7667, 1530-888X.

- Johan A. K. Suykens, Tony Van Gestel, Jos De Brabanter, Bart De Moor, and Joos Vandewalle. *Least Squares Support Vector Machines*. World Scientific, River Edge, NJ, January 2002. ISBN 978-981-238-151-4.
- Michael E. Tipping and Chris M. Bishop. Probabilistic principal component analysis. *Journal Of The Royal Statistical Society, series B*, 61(3):611–622, 1999.
- Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning GAN for pose-invariant face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1415–1424, 2017.
- Aaron Van Den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel Recurrent Neural Networks. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16*, pp. 1747–1756. JMLR.org, 2016.
- Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol. Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion. *Journal of Machine Learning Research*, 11:3371–3408, 2010.
- Jason Weston, Bernhard Schölkopf, and Gökhan H. Bakir. Learning to Find Pre-Images. In *NIPS 16*, pp. 449–456. 2004.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms, 2017.
- Raymond A. Yeh, Chen Chen, Teck Yian Lim, Alexander G. Schwing, Mark Hasegawa-Johnson, and Minh N. Do. Semantic image inpainting with deep generative models. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

A APPENDIX

A.1 DERIVATION OF GEN-RKM OBJECTIVE FUNCTION

Given $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$, where $\mathbf{x}_i \in \mathbb{R}^d$, $\mathbf{y}_i \in \mathbb{R}^p$ and feature-map $\phi_1 : \mathbb{R}^d \mapsto \mathbb{R}^{d_f}$ and $\phi_2 : \mathbb{R}^p \mapsto \mathbb{R}^{p_f}$, the Least-Squares Support Vector Machine (LS-SVM) formulation of Kernel PCA (Suykens et al., 2002) for the two data sources can be written as:

$$\begin{aligned} \min_{\mathbf{U}, \mathbf{V}, \mathbf{e}_i} \quad & \frac{\eta_1}{2} \text{Tr}(\mathbf{U}^\top \mathbf{U}) + \frac{\eta_2}{2} \text{Tr}(\mathbf{V}^\top \mathbf{V}) - \frac{1}{2\lambda} \sum_{i=1}^N \mathbf{e}_i^\top \mathbf{e}_i \\ \text{s.t.} \quad & \mathbf{e}_i = \mathbf{U}^\top \phi_1(\mathbf{x}_i) + \mathbf{V}^\top \phi_2(\mathbf{y}_i) \quad \forall i = 1, \dots, N, \end{aligned} \quad (11)$$

where $\mathbf{U} \in \mathbb{R}^{d \times s}$ and $\mathbf{V} \in \mathbb{R}^{p \times s}$ are the interconnection matrices.

Using the notion of *conjugate feature duality* introduced in (Suykens, 2017), the error variables \mathbf{e}_i are conjugated to latent variables \mathbf{h}_i using:

$$\frac{1}{2\lambda} \mathbf{e}^\top \mathbf{e} + \frac{\lambda}{2} \mathbf{h}^\top \mathbf{h} \geq \mathbf{e}^\top \mathbf{h}, \quad \forall \mathbf{e}, \mathbf{h} \in \mathbb{R}^s \quad (12)$$

which is also known as the Fenchel-Young inequality for the case of quadratic functions (Rockafellar, 1974). By eliminating the variables \mathbf{e}_i from Eq. 11 and using Eq. 12, we obtain the Gen-RKM training objective function:

$$\mathcal{J}_t = \sum_{i=1}^N \left(-\phi_1(\mathbf{x}_i)^\top \mathbf{U} \mathbf{h}_i - \phi_2(\mathbf{y}_i)^\top \mathbf{V} \mathbf{h}_i + \frac{\lambda}{2} \mathbf{h}_i^\top \mathbf{h}_i \right) + \frac{\eta_1}{2} \text{Tr}(\mathbf{U}^\top \mathbf{U}) + \frac{\eta_2}{2} \text{Tr}(\mathbf{V}^\top \mathbf{V}). \quad (13)$$

A.2 KERNEL PCA IN THE PRIMAL

From Eq. 2, eliminating the variables \mathbf{h}_i yields the following:

$$\begin{aligned} \frac{1}{\eta_1} \left[\sum_{i=1}^N \phi_1(\mathbf{x}_i) \phi_1(\mathbf{x}_i)^\top \mathbf{U} + \sum_{i=1}^N \phi_1(\mathbf{x}_i) \phi_2(\mathbf{y}_i)^\top \mathbf{V} \right] &= \lambda \mathbf{U}, \\ \frac{1}{\eta_2} \left[\sum_{i=1}^N \phi_2(\mathbf{y}_i) \phi_1(\mathbf{x}_i)^\top \mathbf{U} + \sum_{i=1}^N \phi_2(\mathbf{y}_i) \phi_2(\mathbf{y}_i)^\top \mathbf{V} \right] &= \lambda \mathbf{V}. \end{aligned} \quad (14)$$

Denote $\Phi_{\mathbf{x}} := [\phi_1(\mathbf{x}_1), \dots, \phi_1(\mathbf{x}_N)]$, $\Phi_{\mathbf{y}} := [\phi_2(\mathbf{y}_1), \dots, \phi_2(\mathbf{y}_N)]$ and $\Lambda = \text{diag}\{\lambda_1, \dots, \lambda_s\} \in \mathbb{R}^{s \times s}$ with $s \leq N$. Now, composing the above equations in matrix form, we get the following eigen-decomposition problem:

$$\begin{bmatrix} \frac{1}{\eta_1} \Phi_{\mathbf{x}} \Phi_{\mathbf{x}}^\top & \frac{1}{\eta_1} \Phi_{\mathbf{x}} \Phi_{\mathbf{y}}^\top \\ \frac{1}{\eta_2} \Phi_{\mathbf{y}} \Phi_{\mathbf{x}}^\top & \frac{1}{\eta_2} \Phi_{\mathbf{y}} \Phi_{\mathbf{y}}^\top \end{bmatrix} \begin{bmatrix} \mathbf{U} \\ \mathbf{V} \end{bmatrix} = \begin{bmatrix} \mathbf{U} \\ \mathbf{V} \end{bmatrix} \Lambda. \quad (15)$$

Here the size of the covariance matrix is $(d_f + p_f) \times (d_f + p_f)$. The latent variables \mathbf{h}_i can be computed using Eq. 2, which simply involves matrix multiplications.

A.3 STABILIZING THE OBJECTIVE FUNCTION

Proposition 1. *All stationary solutions for \mathbf{H}, Λ in Eq. 3 of \mathcal{J}_t lead to $\mathcal{J}_t = 0$.*

Proof. Let λ_i, \mathbf{h}_i are given by Eq. 3. Using Eq. 2 to substitute \mathbf{V} and \mathbf{U} in Eq. 1 yields:

$$\begin{aligned} \mathcal{J}_t(\mathbf{V}, \mathbf{U}, \Lambda, \mathbf{H}) &= \sum_{i=1}^N -\frac{\lambda}{2} \mathbf{h}_i^\top \mathbf{h}_i + \frac{\eta_1}{2} \text{Tr} \left(\frac{1}{\eta_1^2} \sum_{i=1}^N \mathbf{h}_i \phi_1(\mathbf{x}_i)^\top \sum_{j=1}^N \phi_1(\mathbf{x}_j) \mathbf{h}_j^\top \right) \\ &+ \frac{\eta_2}{2} \text{Tr} \left(\frac{1}{\eta_2^2} \sum_{i=1}^N \mathbf{h}_i \phi_2(\mathbf{y}_i)^\top \sum_{j=1}^N \phi_2(\mathbf{y}_j) \mathbf{h}_j^\top \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^N -\frac{\lambda}{2} \mathbf{h}_i^\top \mathbf{h}_i + \frac{\eta_1}{2} \text{Tr} \left(\frac{1}{\eta_1^2} \mathbf{H} \mathbf{K}_1 \mathbf{H}^\top \right) + \frac{\eta_2}{2} \text{Tr} \left(\frac{1}{\eta_2^2} \mathbf{H} \mathbf{K}_2 \mathbf{H}^\top \right) \\
&= \sum_{i=1}^N -\frac{\lambda}{2} \mathbf{h}_i^\top \mathbf{h}_i + \frac{1}{2} \text{Tr} \left(\mathbf{H} \left[\frac{1}{\eta_1} \mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{K}_2 \right] \mathbf{H}^\top \right).
\end{aligned}$$

From Eq. 3, we get:

$$\mathcal{J}_t(\mathbf{V}, \mathbf{U}, \boldsymbol{\Lambda}, \mathbf{H}) = \sum_{i=1}^N -\frac{\lambda}{2} \mathbf{h}_i^\top \mathbf{h}_i + \frac{1}{2} \text{Tr}(\mathbf{H} \mathbf{H}^\top \boldsymbol{\Lambda}) = \sum_{i=1}^N -\frac{\lambda}{2} \mathbf{h}_i^\top \mathbf{h}_i + \frac{\lambda}{2} \sum_{i=1}^N \mathbf{h}_i^\top \mathbf{h}_i = 0.$$

□

Proposition 2. Let $J(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}$ be a smooth function, for all $\mathbf{x} \in \mathbb{R}^N$ and for $c \in \mathbb{R}_{>0}$, define $\bar{J}(\mathbf{x}) := J(\mathbf{x}) + \frac{c}{2} J(\mathbf{x})^2$. Assuming $(1 + cJ(\mathbf{x})) \neq 0$, then \mathbf{x}^* is the stationary points of $\bar{J}(\mathbf{x})$ iff \mathbf{x}^* is the stationary point for $J(\mathbf{x})$.

Proof. Let \mathbf{x}^* be a stationary point of $J(\mathbf{x})$, meaning that $\nabla J(\mathbf{x}^*) = 0$. The stationary points for $\bar{J}(\mathbf{x})$ can be obtained from:

$$\frac{d\bar{J}}{d\mathbf{x}} = (\nabla J(\mathbf{x}) + cJ(\mathbf{x})\nabla J(\mathbf{x})) = (1 + cJ(\mathbf{x})) \nabla J(\mathbf{x}). \quad (16)$$

It is easy to see from Eq. 2 that if $\mathbf{x} = \mathbf{x}^*$, $\nabla J(\mathbf{x}^*) = 0$, we have that $\left. \frac{d\bar{J}}{d\mathbf{x}} \right|_{\mathbf{x}^*} = 0$, meaning that all the stationary points of $J(\mathbf{x})$ are stationary points of $\bar{J}(\mathbf{x})$.

To show the other way, let \mathbf{x}^* be stationary point of $\bar{J}(\mathbf{x})$ i.e. $\nabla \bar{J}(\mathbf{x}^*) = 0$. Assuming $(1 + cJ(\mathbf{x}^*)) \neq 0$, then from Eq. 16 for all $c \in \mathbb{R}_{>0}$, we have

$$(1 + cJ(\mathbf{x}^*)) \nabla J(\mathbf{x}^*) = 0,$$

implying that $\nabla J(\mathbf{x}^*) = 0$. □

Based on the above propositions, we stabilize our original objective function Eq. 1 to keep it bounded and hence is suitable for minimization with Gradient-descent methods. Without the reconstruction errors, the stabilized objective function is

$$\min_{\mathbf{U}, \mathbf{V}, \mathbf{h}_i} \mathcal{J}_t + \frac{c}{2} \mathcal{J}_t^2.$$

Denoting $\bar{J} = \mathcal{J}_t + \frac{c_{stab}}{2} \mathcal{J}_t^2$. Since the derivatives of \mathcal{J}_t are given by Eq. 2, the stationary points of \bar{J} are:

$$\begin{cases} \frac{\partial \bar{J}}{\partial \mathbf{V}} = (1 + c_{stab} \mathcal{J}_t) \left(-\sum_{i=1}^N \phi_1(\mathbf{x}_i) \mathbf{h}_i^\top + \eta_1 \mathbf{V} \right) = 0 & \implies \mathbf{V} = \frac{1}{\eta_1} \sum_{i=1}^N \phi_1(\mathbf{x}_i) \mathbf{h}_i^\top, \\ \frac{\partial \bar{J}}{\partial \mathbf{U}} = (1 + c_{stab} \mathcal{J}_t) \left(-\sum_{i=1}^N \phi_2(\mathbf{y}_i) \mathbf{h}_i^\top + \eta_2 \mathbf{U} \right) = 0 & \implies \mathbf{U} = \frac{1}{\eta_2} \sum_{i=1}^N \phi_2(\mathbf{y}_i) \mathbf{h}_i^\top, \\ \frac{\partial \bar{J}}{\partial \mathbf{h}_i} = (1 + c_{stab} \mathcal{J}_t) \left(-\mathbf{V}^\top \phi_1(\mathbf{x}_i) - \mathbf{U}^\top \phi_2(\mathbf{y}_i) + \lambda \mathbf{h}_i \right) = 0 & \implies \lambda \mathbf{h}_i = \mathbf{V}^\top \phi_1(\mathbf{x}_i) \\ & \quad + \mathbf{U}^\top \phi_2(\mathbf{y}_i), \end{cases}$$

assuming $1 + c_{stab} \mathcal{J}_t \neq 0$. Elimination of \mathbf{V} and \mathbf{U} yields $\left[\frac{1}{\eta_1} \mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{K}_2 \right] \mathbf{H}^\top = \mathbf{H}^\top \boldsymbol{\Lambda}$, which is indeed the same solution for $c_{stab} = 0$ in Eq. 1 and Eq. 3.

A.4 CENTERING OF KERNEL MATRIX

Centering of the kernel matrix is done by the following equation:

$$\mathbf{K}_c = \mathbf{K} - N^{-1} \mathbf{1} \mathbf{1}^\top \mathbf{K} - N^{-1} \mathbf{K} \mathbf{1} \mathbf{1}^\top + N^{-2} \mathbf{1} \mathbf{1}^\top \mathbf{K} \mathbf{1} \mathbf{1}^\top, \quad (17)$$

where $\mathbf{1}$ denotes an N -dimensional vector of ones and \mathbf{K} is either \mathbf{K}_1 or \mathbf{K}_2 .

A.5 ARCHITECTURE DETAILS

See Table 1 and 2 for details on model architectures, datasets and hyperparameters used in this paper. The PyTorch library in Python was used as the programming language with a 8GB NVIDIA QUADRO P4000 GPU.

Table 1: Datasets and hyperparameters used for the experiments.

Dataset	N	d	N_{subset}	s	m	σ	n_r	l
MNIST	60000	28×28	5000	500	50	1.3	4	10
Fashion-MNIST	60000	28×28	500	100	5	/	/	10
CIFAR-10	60000	$32 \times 32 \times 3$	500	500	5	/	/	10
CelebA	202599	$128 \times 128 \times 3$	500	15	5	/	/	20
Dsprites	737280	64×64	1024	2	5	/	/	/

Table 2: Details of model architectures used in the paper. All convolutions and transposed-convolutions are with stride 2 and padding 1. Unless stated otherwise, the layers have Parametric-RELU ($\alpha = 0.2$) activation function, except the output layers of the pre-image maps which has sigmoid activation function.

Dataset	Optimizer (Adam)	Architecture		
		\mathcal{X}	\mathcal{Y}	
MNIST	1e-3	Input	28x28x1	10 (One-hot encoding)
		Feature-map (fm)	Conv 32x4x4; Conv 64x4x4; FC 128 (Linear)	FC 15, 20 (Linear)
		Pre-image map	reverse of fm	reverse of fm
		Latent space dim.		500
Fashion-MNIST	1e-3	Input	28x28x1	10 (One-hot encoding)
		Feature-map	Conv 32x4x4; 64x4x4; FC 128 (Linear)	FC 15, 20
		Pre-image map (fm)	reverse of fm	reverse of fm
		Latent space dim.		100
CIFAR-10	1e-3	Input	32x32x3	10 (One-hot encoding)
		Feature-map (fm)	Conv 64x4x4; Conv 128x4x4; FC 128 (Linear)	FC 15, 20
		Pre-image map	reverse of fm	reverse of fm
		Latent space dim.		500
CelebA	1e-4	Input	64x64x3	-
		Feature-map (fm)	Conv 32x4x4; Conv 64x4x4; Conv 128x4x4; Conv 256x4x4 ; FC 128 (Linear)	-
		Pre-image map	reverse of fm	-
		Latent space dim.		15
Dsprites	1e-4	Input	64x64x1	-
		Feature-map (fm)	Conv 20x4x4; Conv 40x4x4; Conv 80x4x4; FC 300 (Linear)	-
		Pre-image map	reverse of fm	-
		Latent space dim.		2

A.6 FURTHER EMPIRICAL RESULTS

In the following section, we show generated images by bilinear-interpolation of latent vectors. Given four vectors $\mathbf{h}_1, \mathbf{h}_2, \mathbf{h}_3$ and \mathbf{h}_4 (reconstructed images from these vectors are shown at the edges of Figs. 6, 7), the interpolated vector \mathbf{h}^* is given by:

$$\mathbf{h}^* = (1 - \alpha)(1 - \gamma)\mathbf{h}_1 + \alpha(1 - \gamma)\mathbf{h}_2 + \gamma(1 - \alpha)\mathbf{h}_3 + \gamma\alpha\mathbf{h}_4, \quad 0 \leq \alpha, \gamma \leq 1.$$

This \mathbf{h}^* is then used in step 8 of the generation procedure of Gen-RKM algorithm (see Algorithm 1) to compute \mathbf{x}^* .




Figure 6: Reconstructed images by bilinear-interpolation in latent space. The model was trained using the Binary Cross-entropy reconstruction loss.



Figure 7: CelebA: Reconstructed images by bilinear-interpolation in latent space.