

Full Length Article

Deep Kernel Principal Component Analysis for multi-level feature learning

Francesco Tonin^{*}, Qinghua Tao^{*}, Panagiotis Patrinos, Johan A.K. Suykens

Department of Electrical Engineering, ESAT-STADIUS, KU Leuven, Kasteelpark Arenberg 10, B-3001 Leuven, Belgium

ARTICLE INFO

Keywords:

Kernel Principal Component Analysis
 Deep learning
 Generative models
 Manifold optimization

ABSTRACT

Principal Component Analysis (PCA) and its nonlinear extension Kernel PCA (KPCA) are widely used across science and industry for data analysis and dimensionality reduction. Modern deep learning tools have achieved great empirical success, but a framework for deep principal component analysis is still lacking. Here we develop a deep kernel PCA methodology (DKPCA) to extract multiple levels of the most informative components of the data. Our scheme can effectively identify new hierarchical variables, called deep principal components, capturing the main characteristics of high-dimensional data through a simple and interpretable numerical optimization. We couple the principal components of multiple KPCA levels, theoretically showing that DKPCA creates both forward and backward dependency across levels, which has not been explored in kernel methods and yet is crucial to extract more informative features. Various experimental evaluations on multiple data types show that DKPCA finds more efficient and disentangled representations with higher explained variance in fewer principal components, compared to the shallow KPCA. We demonstrate that our method allows for effective hierarchical data exploration, with the ability to separate the key generative factors of the input data both for large datasets and when few training samples are available. Overall, DKPCA can facilitate the extraction of useful patterns from high-dimensional data by learning more informative features organized in different levels, giving diversified aspects to explore the variation factors in the data, while maintaining a simple mathematical formulation.

1. Introduction

Principal Component Analysis (PCA) is a popular technique for dimensionality reduction (Jolliffe, 1986) and has been widely used in many fields (Lever et al., 2017). In fact, high-dimensional data are very common in data science when multiple variables are used to describe one sample; e.g., in biology, PCA has been applied to mass spectrometry, where thousands of proteins can be quantitatively profiled (Ringnér, 2008). PCA learns the most effective principal components to successfully reduce the dimensionality of the data while retaining most of the trends and patterns. This relies on the assumption that the given observations lie in a lower-dimensional linear subspace. Under this assumption, PCA seeks the best low-rank representation of the given data. PCA can be efficiently computed using the Singular Value Decomposition (SVD) and is optimal when data are corrupted by small Gaussian noises (Wright et al., 2009). Real-world data commonly show nonlinear relationships, so, for nonlinear problems, PCA can be extended to Kernel PCA (KPCA) (Schölkopf et al., 1998), which manages to simplify such complexity and high dimensionality to extract useful patterns in nonlinear subspaces. KPCA first maps the inputs to a high-dimensional feature space and then applies PCA to the mapped features either through nonlinear feature mappings in the primal or

equivalently kernel functions in the dual. In the Lagrange dual formulation of KPCA, the feature map does not need to be explicitly defined and positive-definite kernel functions are instead used by Mercer's theorem (Mercer, 1909).

In deep learning, dimensionality reduction and learning informative features are also widely studied through the latent space models, such as Variational Autoencoders (VAEs) (Kingma & Welling, 2014), which have become popular tools to extract latent features describing the factors of variation in the given training distribution. These models assume that there exists a prior distribution $p(z)$ over a small number of ground-truth factors of variation, such that an observation x is obtained by first sampling z from $p(z)$ and then sampling from a conditional distribution $p(x|z)$. In this setting, the goal is to find a representation of the data that learns the factors of variation in z independently, i.e., that disentangles the factors of variation. State-of-the-art models for disentangled feature learning include InfoGAN (Chen et al., 2016), Restricted Boltzmann machines (Hinton et al., 2006; Reed et al., 2014), β -VAE (Higgins et al., 2017) and its variants (Chen et al., 2018; Kim & Mnih, 2018). For instance, in β -VAE, $p(z) = \mathcal{N}(0, I)$ and the encoder $q(z|x)$ is matched to the prior $p(z)$ by minimizing the Kullback–Leibler divergence $D_{KL}(q(z|x)||p(z))$. Neural networks are used to model

^{*} Corresponding authors.

E-mail addresses: francesco.tonin@esat.kuleuven.be (F. Tonin), qinghua.tao@esat.kuleuven.be (Q. Tao), panos.patrinos@esat.kuleuven.be (P. Patrinos), johan.suykens@esat.kuleuven.be (J.A.K. Suykens).

<https://doi.org/10.1016/j.neunet.2023.11.045>

Received 22 February 2023; Received in revised form 4 October 2023; Accepted 21 November 2023

Available online 30 November 2023

0893-6080/© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

the generative model with probabilistic encoder $q(z|x)$ and decoder $p(x|z)$ (Kingma & Welling, 2014). A recent large-scale extensive experimental research has shown that the performance of VAE-based models varies greatly with random initialization, hyperparameters, and dataset, so reliable extraction of independent components describing the variation factors of data remains challenging (Locatello et al., 2019).

While (K)PCA has been widely used in science and industry, the modeling flexibility of using a single feature mapping or kernel function can be insufficient and it also cannot learn well-disentangled representations (Higgins et al., 2017). For such feature-learning tools, disentanglement of the variation factors (components) in the data is highly desirable (Bengio, 2009; Bengio et al., 2013) and it has been suggested that disentangled representations can benefit interpretation analysis, e.g., in the medical domain Holzinger et al. (2019), Sarhan et al. (2019). For instance, a model trained on gene expression data may learn components such as the cell type or the cell state. In addition, because (K)PCA is a shallow model employing a single feature mapping, it learns only one flat level of components. On the other hand, deep learning has achieved pervasive empirical success with great modeling flexibility (Goodfellow et al., 2016), but a framework combining deep architectures and principal component analysis remains lacking.

Deep kernel learning tackles multiple latent spaces for greater flexibility, more informative hierarchical investigation of the data, and kernel-based interpretations. There exist many works in deep kernel learning considering supervised learning (see Bohn et al. (2019) and references therein), but little investigation has been spared on the unsupervised settings, though a concatenation of operator-valued kernel layers was considered for data autoencoding in Laforgue et al. (2019). In Deng et al. (2019), it is proposed to conduct the shallow PCA to extract principal components, which are then applied to another KPCA, where each KPCA independently and sequentially optimizes its variance maximization. Importantly, when extending to deep architectures, Allen-Zhu and Li (2023) warn that simply doing a sequential kernel learning is not enough to achieve good accuracy due to the lack of backward feature correction, meaning that shallow layers need to use the information from deeper layers to boost their own learned representation. In Allen-Zhu and Li (2023), it is proved that hierarchical learning cannot be efficiently achieved without backward feature correction.

In this paper, we establish a novel Deep Kernel Principal Component Analysis (DKPCA) framework with the following main aspects.

- DKPCA presents multiple levels of principal components associated with the key properties of the data for more informative feature learning in multiple subspaces. The objective of each level is attained as an upper bound of a shallow KPCA¹ problem, and multiple levels are constructed by coupling the latent space of level $j - 1$ ($j \geq 2$) with the input space of level j , where the depth is given by the learned spaces directly relating to the principal components, as shown in Fig. 1. We derive that the optimization problem of our method explicitly formulates a set of nonlinear equations for each level resembling an eigenvalue problem of some matrix M_j , in contrast with black-box optimization in deep learning.
- Interestingly, M_j fuses the hidden features of previous and subsequent levels. This means that the proposed deep architecture introduces not only *forward* couplings between the levels, but also *backward* couplings, which by far has not been explored in kernel methods and yet is crucial for effective hierarchical representation learning according to the theoretical analysis in Allen-Zhu and Li (2023). As the levels are coupled together, we formulate a

multi-level constrained optimization problem with an eigenvalue problem at each level with hidden features as optimization variables, also facilitating deep approximation analysis of the given data.

- The solution of the proposed optimization process gives both the deep eigenvectors and the deep eigenvalues of the DKPCA: they correspond to the solution of the eigenvalue problem of each level. Within the considered deep architecture, we then construct a generative procedure for the DKPCA by defining both an out-of-sample encoding scheme and a decoding procedure, discussing connections with Autoencoders. The generative procedure generates new samples from multiple latent spaces in different levels, makes it possible to explore the role of the deep eigenvectors of each level through the latent space traversals, and gives diversified aspects to explore the variation factors of data. Our method can also be implemented with out-of-sample extensions which allow to efficiently tackle large-scale cases.

Extensive numerical experiments demonstrate the efficacy and advantages of the proposed DKPCA from different aspects and in different tasks on multiple data types. (1) DKPCA gives higher explained variance than shallow KPCA, indicating that more information is captured in fewer components. We also provide a strategy for practitioners to select the numbers of components and levels, which is in contrast with typical deep learning tools that use trial and error strategies in determining the network structure. (2) DKPCA effectively facilitates hierarchical data exploration, as the role of each principal component in each level can be investigated through the generation of new data. In images of 3D objects with different generative factors (i.e., colors, size, etc.), our deep method creates a learning hierarchy in the components in each level. Prevailing features are typically learned in the shallower levels, e.g., colors, while the deeper levels capture more subtle features, e.g., the specific object shape. (3) Quantitative performances are evaluated by comparing to state-of-the-art methods in disentangled feature learning (Chen et al., 2018; Higgins et al., 2017; Kim & Mnih, 2018) when few training samples are available, which is of particular interest in many real-world problems where data are difficult or expensive to collect. (4) We show that the more informative features extraction by DKPCA can be applied to multiple data types benefiting various downstream tasks in data science, such as regression and classification.

2. Background and related work

This section describes the shallow KPCA problem under the RKM framework through the Fenchel–Young inequality, and elaborates on the main novelties with respect to existing works in unsupervised deep kernel machines.

2.1. KPCA with restricted kernel machines

The RKM formulation of KPCA gives another expression of the Least-Squares Support Vector Machine (LS-SVM) KPCA problem (Suykens et al., 2003) with visible and hidden units similar to the energy of Restricted Boltzmann Machines (RBMs) (Bengio, 2009; Fischer & Igel, 2014; Hinton et al., 2006; Salakhutdinov, 2015). In this new formulation, contrary to RBMs, both the visible units and the hidden units can be continuous. To derive this formulation, consider training data $D = \{\mathbf{x}_i\}_{i=1}^N$ with $\mathbf{x}_i \in \mathbb{R}^d$, a feature map $\varphi : \mathbb{R}^d \mapsto \mathbb{R}^{d_r}$, and let s be the number of selected principal components. In the LS-SVM setting, the KPCA problem can be written as minimizing a regularization term and finding directions of maximum variance (Suykens et al., 2002):

$$\begin{aligned} & \underset{\mathbf{W}, \mathbf{e}_i}{\text{minimize}} && J_{\text{kPCA}} = \frac{\eta}{2} \text{Tr}(\mathbf{W}^T \mathbf{W}) - \frac{1}{2} \sum_{i=1}^N \mathbf{e}_i^T \Lambda^{-1} \mathbf{e}_i \\ & \text{subject to} && \mathbf{e}_i = \mathbf{W}^T \varphi(\mathbf{x}_i), \quad i = 1, \dots, N, \end{aligned} \quad (1)$$

¹ To differentiate DKPCA, we name the classical KPCA as shallow KPCA considering its one-level architecture.




Fig. 1. Topology of the RKM-based deep KPCA with n_L levels. An input vector \mathbf{x} is mapped to the feature space of the first level using a feature map φ_1 with hidden features $\mathbf{h}^{(1)}$ in the latent space of the first level. Subsequently, the input of level j , with feature map φ_j , are the hidden features of level $j - 1$.

where $\mathbf{W} \in \mathbb{R}^{d_F \times s}$ is the interconnection matrix, $\mathbf{e}_i \in \mathbb{R}^s$ are the score variables along the selected s projection directions, and $\mathbf{\Lambda} = \text{diag}\{\lambda_1, \dots, \lambda_s\} > 0, \eta > 0$ are regularization hyperparameters.

The RKM formulation of KPCA (Suykens, 2017) is given by an upper bound of $J_{k\text{pca}}$ obtained component-wise with the Fenchel–Young inequality $\frac{1}{2\lambda}e^2 + \frac{\lambda}{2}h^2 \geq eh, \forall e, h \in \mathbb{R}$ which introduces the hidden features \mathbf{h} and leads to the following objective with conjugate feature duality:

$$\bar{J}_{k\text{pca}} = - \sum_{i=1}^N \varphi(\mathbf{x}_i)^T \mathbf{W} \mathbf{h}_i + \frac{1}{2} \sum_{i=1}^N \mathbf{h}_i^T \mathbf{\Lambda} \mathbf{h}_i + \frac{\eta}{2} \text{Tr}(\mathbf{W}^T \mathbf{W}), \quad (2)$$

where $\mathbf{h}_i \in \mathbb{R}^s$ are the conjugated hidden features corresponding to each training sample \mathbf{x}_i ; in representation learning, \mathbf{h}_i is also known as the latent representation of \mathbf{x}_i consisting of s latent variables or of s hidden features. Note that the first term of (2) is similar to the energy of an RBM with connections between visible units \mathbf{x}_i in the input space and hidden units \mathbf{h}_i in the latent space. The stationary point conditions of $\bar{J}_{k\text{pca}}(\mathbf{W}, \mathbf{h}_i)$ are given by:

$$\begin{cases} \frac{\partial \bar{J}_{k\text{pca}}(\mathbf{W}, \mathbf{h}_i)}{\partial \mathbf{h}_i} = 0 \implies \mathbf{W}^T \varphi(\mathbf{x}_i) = \mathbf{\Lambda} \mathbf{h}_i, \forall i = 1, \dots, N \\ \frac{\partial \bar{J}_{k\text{pca}}(\mathbf{W}, \mathbf{h}_i)}{\partial \mathbf{W}} = 0 \implies \mathbf{W} = \frac{1}{\eta} \sum_{i=1}^N \varphi(\mathbf{x}_i) \mathbf{h}_i^T. \end{cases} \quad (3)$$

Eliminating \mathbf{W} and considering a positive definite kernel function $k : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$ with $k(\mathbf{x}_i, \mathbf{x}_j) = \varphi(\mathbf{x}_i)^T \varphi(\mathbf{x}_j)$, the stationary points of $\bar{J}_{k\text{pca}}(\mathbf{W}, \mathbf{h}_i)$ are given in the dual by the following eigenvalue problem

$$\frac{1}{\eta} \mathbf{K} \mathbf{H} = \mathbf{H} \mathbf{\Lambda}, \quad (4)$$

where $\mathbf{K} \in \mathbb{R}^{N \times N}$ denotes the kernel matrix induced by $k(\cdot, \cdot)$ and the matrix $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_N]^T$ incorporates the conjugate hidden features for all N data points. In (4), the hidden features \mathbf{H} conjugated along s projection directions now correspond to the first s eigenvectors, with the first s eigenvalues corresponding to the hyperparameter $\mathbf{\Lambda}$ in (2). Meanwhile, η becomes a scaling coefficient that does not change the solution space, and thus can be simply set as 1. Note that, in the conjugate feature duality of RKMs, the dual variables \mathbf{h} correspond to the latent variables playing the role of hidden features living in the latent space.

The dual problem (4) corresponds to the kernel PCA problem as defined in Schölkopf et al. (1998). While (4) is regularized by normalizing the eigenvectors to the unit ball in feature space, the primal problem (2) is explicitly regularized with coefficients $\mathbf{\Lambda}, \eta$ chosen at the hyperparameter selection level. Each eigenvalue/eigenvector pair corresponds to a principal component in KPCA. Therefore, for the first s principal components, one can solve the dual problem (4) by considering the s largest eigenvalues and their eigenvectors, which lead to $\bar{J}_{k\text{pca}} = 0$. Since $\bar{J}_{k\text{pca}}$ is unbounded below regarding its optimization in the primal, Suykens (2017) proposed to instead minimize a stabilized version to make the objective suitable for minimization, such that $\bar{J}_{k\text{pca}, \text{stab}} = \bar{J}_{k\text{pca}} + \frac{c_{\text{stab}}}{2} \bar{J}_{k\text{pca}}^2$, where $c_{\text{stab}} > 0$ is a hyperparameter. It can be shown that $\bar{J}_{k\text{pca}}$ and $\bar{J}_{k\text{pca}, \text{stab}}$ share the same stationary points (Pandey et al., 2021).

2.2. Related works

Suykens (2017) first introduced the RKM framework but only presented deep configurations for supervised learning. In this work, we study deep unsupervised learning with RKMs, which is a more challenging problem due to the objective in (2) being unbounded. Thus, one cannot straightforwardly solve such problem by a objective minimization, which is instead feasible in the supervised cases studied in Suykens (2017) thanks to the classification or regression losses. The training algorithm of Suykens (2017) only works with linear kernels in deep levels ($j \geq 2$), where the nonlinearity can be applied only to the first level. In our DKPCA, all levels are allowed to be nonlinear in the proposed new formulation. In Tonin et al. (2021), a two-level kernel architecture for unsupervised learning was proposed with orthogonality constraints on the latent variables within each level and across levels. It employs a straightforward numerical approach to solve an unconstrained optimization problem augmented with a penalty term for the constraints, where the backward couplings between the levels are omitted in the formulation, which can be instead realized by our DKPCA. Compared to previous related works in deep kernel machines (Suykens, 2017; Tonin et al., 2021), we derive a new interpretation of the deep KPCA problem as a multi-level eigenvalue problem with both forward and backward couplings, which has not been proposed yet in kernel methods. Differently from Suykens (2017), Tonin et al. (2021), building upon this new interpretation, we propose a residual minimization procedure augmented by orthogonality constraints of the hidden features within each level. We develop a new end-to-end training algorithm that jointly optimizes all KPCA levels and realizes the backward couplings in the learning process. Additionally, to more formally understand unsupervised deep kernel machines, we present novel analytical results that have not been addressed in previous works. We derive deep error approximation bounds through the Eckart-Young theorem applied to multi-level architectures, and we prove that the explained variance by DKPCA can be greater than that by shallow KPCA with appropriate regularization coefficients. In Deng et al. (2019), PCA is firstly conducted to extract principal components of the data and then further dimensionality reduction is sequentially applied to the extracted features from the previous (K)PCA layer. Their procedure only considers forward couplings, as each layer simply optimizes its variance maximization objective, which is independent of other layers. Thanks to the new formulation of our DKPCA, all levels can be jointly optimized, where both forward and backward couplings between the levels are considered. Notably, contrary to Deng et al. (2019), Suykens (2017), Tonin et al. (2021), our proposed DKPCA can also be used as a generative model. While reconstruction in linear shallow PCA can be realized by basis transformation, the generation in nonlinear kernel machines and especially in deep kernel machines with multiple nonlinearities is challenging. Our DKPCA makes the generation feasible through multiple latent spaces, facilitating hierarchical data exploration, through which the role of each principal component in each level can be investigated.

3. Deep kernel principal component analysis

In this section, we present the proposed DKPCA. We start by describing the model formulation of DKPCA. Next, we derive the optimization algorithm. Finally, the generative DKPCA model is introduced.

$$\left\{ \begin{array}{l} \text{Level 1:} \\ \text{Level } j: \\ \text{Level } n_{\text{levels}}: \end{array} \right. \begin{cases} \left[\frac{1}{\eta_1} \mathbf{K}_1 + \frac{1}{\eta_2} \mathcal{G}_1(\mathbf{H}_1, \mathbf{H}_2) \mathbf{H}_1^\top \right] \mathbf{H}_1 \\ \left[\frac{1}{\eta_j} \mathbf{K}_j(\mathbf{H}_{j-1}) + \frac{1}{\eta_{j+1}} \mathcal{G}_j(\mathbf{H}_j, \mathbf{H}_{j+1}) \mathbf{H}_j^\top \right] \mathbf{H}_j \\ \frac{1}{\eta_{n_{\text{levels}}}} \mathbf{K}_{n_{\text{levels}}}(\mathbf{H}_{n_{\text{levels}}-1}) \mathbf{H}_{n_{\text{levels}}} \end{cases} = \begin{cases} \mathbf{H}_1 \boldsymbol{\Lambda}_1, \\ \mathbf{H}_j \boldsymbol{\Lambda}_j, \quad \forall j = 2, \dots, n_{\text{levels}} - 1, \\ \mathbf{H}_{n_{\text{levels}}} \boldsymbol{\Lambda}_{n_{\text{levels}}}. \end{cases} \quad (6)$$

Box I.

3.1. DKPCA model formulation

We construct the objective function of DKPCA by joining the KPCA objectives of multiple levels in the Restricted Kernel Machine (RKM) framework (Suykens, 2017), which combines the flexibility of deep architectures and the interpretations rooted in kernel methods. DKPCA considers general cases consisting of n_{levels} ($n_{\text{levels}} \geq 2$) KPCA levels stacked in the corresponding latent spaces, i.e., the hidden features of level j are the input of level $j + 1$, inducing inter-level couplings, similar to the stacked Autoencoders (Bengio, 2009). Correspondingly, the objective for the proposed DKPCA is formulated in the primal model representation:

$$\begin{aligned} J = & - \sum_{i=1}^N \varphi_1(\mathbf{x}_i)^\top \mathbf{W}_1 \mathbf{h}_i^{(1)} + \frac{1}{2} \sum_{i=1}^N \mathbf{h}_i^{(1)\top} \boldsymbol{\Lambda}_1 \mathbf{h}_i^{(1)} + \frac{\eta_1}{2} \text{Tr}(\mathbf{W}_1^\top \mathbf{W}_1) \\ & + \sum_{j=2}^{n_{\text{levels}}} \left[- \sum_{i=1}^N \varphi_j(\mathbf{h}_i^{(j-1)})^\top \mathbf{W}_j \mathbf{h}_i^{(j)} + \frac{1}{2} \sum_{i=1}^N \mathbf{h}_i^{(j)\top} \boldsymbol{\Lambda}_j \mathbf{h}_i^{(j)} \right. \\ & \left. + \frac{\eta_j}{2} \text{Tr}(\mathbf{W}_j^\top \mathbf{W}_j) \right]. \quad (5) \end{aligned}$$

The feature map $\varphi_1 : \mathbb{R}^d \mapsto \mathbb{R}^{d_{r1}}$ of the first level takes the original data as the input, while $\varphi_j : \mathbb{R}^{s_j} \mapsto \mathbb{R}^{d_{rj}}$ is the feature map of level $j = 2, \dots, n_{\text{levels}}$ that takes the hidden features $\mathbf{h}_i^{(j-1)}$ of level $j - 1$ as the input, where $\mathbf{W}_j \in \mathbb{R}^{d_{rj} \times s_j}$ is the interconnection matrix of level j . Here, the matrix $\mathbf{H}_j = [\mathbf{h}_1^{(j)}, \dots, \mathbf{h}_N^{(j)}]^\top \in \mathbb{R}^{N \times s_j}$ incorporates the hidden features conjugated along s_j projection directions for all N data points, where s_j is the number of selected principal components by the j th level of our DKPCA. In the primal formulation, $\boldsymbol{\Lambda}_j = \text{diag}\{\lambda_1^{(j)}, \dots, \lambda_{s_j}^{(j)}\}$ and $\eta_j \neq 0$ both serve as the hyperparameters of level j . While $\eta > 0$ in the shallow KPCA case for variance maximization in Eq. (2), this constraint is not required in the deep objective (5), having complex inter-level couplings. Note that, in our DKPCA formulation, the visible units \mathbf{x}_i in the input space are conjugated with the multi-level hidden features $\mathbf{h}_i^{(j)}$ in the latent space of each level j , giving an energy function that resembles the deep Boltzmann machine (Salakhutdinov & Hinton, 2009). The DKPCA topology in its primal formulation is visualized in Fig. 1.

The projection directions of shallow (K)PCA are uncorrelated due to the orthogonality of different principal components as in (4). Similarly for DKPCA, we impose intra-level orthogonality on \mathbf{H}_j , i.e., $\mathbf{H}_j^\top \mathbf{H}_j = \mathbf{I}$. From the stationary points of (5), the formulation of DKPCA in the dual variables is in Box I: A graphical illustration of (6) is given in Fig. 2. The kernel matrices are obtained as follows: $\mathbf{K}_1 \in \mathbb{R}^{N \times N}$ is attained as $(\mathbf{K}_1)_{ik} = k_1(\mathbf{x}_i, \mathbf{x}_k)$ and $\mathbf{K}_j \in \mathbb{R}^{N \times N}$ as $(\mathbf{K}_j)_{ik} = k_j(\mathbf{h}_i^{(j-1)}, \mathbf{h}_k^{(j-1)})$, where $k_1 : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$ with $k_1(\mathbf{x}, \mathbf{y}) = \varphi_1(\mathbf{x})^\top \varphi_1(\mathbf{y})$ is the kernel function of the first level and $k_j : \mathbb{R}^{s_{j-1}} \times \mathbb{R}^{s_{j-1}} \mapsto \mathbb{R}$ with $k_j(\mathbf{x}, \mathbf{y}) = \varphi_j(\mathbf{x})^\top \varphi_j(\mathbf{y})$ is the kernel function of level $j = 2, \dots, n_{\text{levels}}$ by the kernel trick. Instead of first defining a feature map φ_j , one can simply choose a positive definite kernel k_j due to Mercer’s theorem (Mercer, 1909), guaranteeing the existence of a feature map φ_j such that $k_j(y, z) = \varphi_j(y)^\top \varphi_j(z)$. The stationary points of (5) can be found in Appendix A.1.

In (6), $\mathcal{G}_j(\mathbf{H}_j, \mathbf{H}_{j+1}) \in \mathbb{R}^{N \times s_j}$ are matrices jointly depending on the conjugated hidden features \mathbf{H}_j and \mathbf{H}_{j+1} . In particular, they are formulated as $\mathcal{G}_j(\mathbf{H}_j, \mathbf{H}_{j+1}) = \left(\mathbf{I}_N \odot \left(\mathbf{H}_{j+1} \mathbf{H}_{j+1}^\top \right) \right)^\top \mathbf{J}_{\mathbf{K}_{j+1}}(\mathbf{H}_j) \in \mathbb{R}^{N \times s_j}$, where the Khatri–Rao product between matrices $\mathbf{A} = [\mathbf{a}_1 \dots \mathbf{a}_n] \in \mathbb{R}^{m_1 \times n}$ and $\mathbf{B} = [\mathbf{b}_1 \dots \mathbf{b}_n] \in \mathbb{R}^{m_2 \times n}$ is $\mathbf{A} \odot \mathbf{B} = [\mathbf{a}_1 \otimes \mathbf{b}_1 \dots \mathbf{a}_n \otimes \mathbf{b}_n]$ with \otimes denoting the Kronecker product,

$$\begin{aligned} \mathbf{J}_{\mathbf{K}_{j+1}}(\mathbf{H}_j) & \triangleq \begin{bmatrix} \mathbf{J}_{\mathbf{K}_{j+1,1}}(\mathbf{H}_j) \\ \mathbf{J}_{\mathbf{K}_{j+1,2}}(\mathbf{H}_j) \\ \vdots \\ \mathbf{J}_{\mathbf{K}_{j+1,N}}(\mathbf{H}_j) \end{bmatrix} \in \mathbb{R}^{N^2 \times s_j} \quad \text{and} \\ \mathbf{J}_{\mathbf{K}_{j+1,i}}(\mathbf{H}_j) & \triangleq \begin{bmatrix} \nabla k_{j+1}(\mathbf{h}_i^{(j)}, \mathbf{h}_1^{(j)})^\top \\ \nabla k_{j+1}(\mathbf{h}_i^{(j)}, \mathbf{h}_2^{(j)})^\top \\ \vdots \\ \nabla k_{j+1}(\mathbf{h}_i^{(j)}, \mathbf{h}_N^{(j)})^\top \end{bmatrix} \in \mathbb{R}^{N \times s_j}. \quad (7) \end{aligned}$$

Below, two examples are illustrated on the derivations of \mathcal{G} when different kernel functions are chosen.

Example 3.1 (Linear Kernel). In the case of linear $k_j(\mathbf{z}, \mathbf{y}) = \mathbf{z}^\top \mathbf{y}$, we obtain $\mathbf{J}_{\mathbf{K}_{j,i}} = \mathbf{H}_{j-1}^\top$, so we further have $\mathcal{G}_{j-1} = \mathbf{H}_j \mathbf{H}_j^\top \mathbf{H}_{j-1}$, and $\mathbf{K}_j(\mathbf{H}_{j-1}) = \mathbf{H}_{j-1} \mathbf{H}_{j-1}^\top$. In two-level architectures with $k_2(\mathbf{z}, \mathbf{y}) = \mathbf{z}^\top \mathbf{y}$, \mathcal{G}_1 has a linear dependency on \mathbf{H}_1 as $\mathcal{G}_1 = \mathbf{H}_2 \mathbf{H}_2^\top \mathbf{H}_1$, where the eigendecomposition for the first level is written in the form of $\mathbf{M}_1(\mathbf{H}_2) \mathbf{H}_1 = \mathbf{H}_1 \boldsymbol{\Lambda}_1$ with $\mathbf{M}_1(\mathbf{H}_2)$ independent of \mathbf{H}_1 .

Example 3.2 (RBF Kernel). Consider $k_j(\mathbf{z}, \mathbf{y}) = \exp(-\|\mathbf{z} - \mathbf{y}\|_2^2 / (2\sigma^2))$. The partial derivative is

$$\frac{\partial k_j(\mathbf{h}_i^{(j-1)}, \mathbf{h}_k^{(j-1)})}{\partial \mathbf{h}_i^{(j-1)}} = -\frac{1}{\sigma^2} \left(\mathbf{h}_i^{(j-1)} - \mathbf{h}_k^{(j-1)} \right) k_j(\mathbf{h}_i^{(j-1)}, \mathbf{h}_k^{(j-1)}), \quad (8)$$

so $\mathbf{J}_{\mathbf{K}_{j,i}} = -2\gamma \text{diag}(\mathbf{K}_{j,i}) [\mathbf{h}_i^{(j-1)} \mathbf{1}^\top - \mathbf{H}_{j-1}]$, where $\mathbf{1}$ is a vector of all ones and $\mathbf{K}_{j,i}$ is the i th column of \mathbf{K}_j .

The derivations to the dual formulations show that $\boldsymbol{\Lambda}_j$ relates to the first s_j eigenvalues corresponding to the s_j eigenvectors \mathbf{H}_j in the optimization of DKPCA, indicating that all the pairs $(\boldsymbol{\Lambda}_j, \mathbf{H}_j)$, $j = 1, \dots, n_{\text{levels}}$ solving the dual problem constitute a pool of candidate solutions that lead to $\mathbf{J} = 0$ in the primal objective (5). Thus, the regularization hyperparameters $\boldsymbol{\Lambda}_j$ in the primal are automatically determined in the dual by the solutions of (6). Such obtained \mathbf{H}_j and $\boldsymbol{\Lambda}_j$ with $j = 1, \dots, n_{\text{levels}}$ are named as deep eigenvectors and deep eigenvalues in DKPCA, respectively. The dual problem of DKPCA in each level is interpreted as an eigenvalue problem, giving the conjugated hidden features (principal components) \mathbf{H}_j solved by the deep eigenvectors corresponding to level j . The existing (shallow) KPCA is a special case of DKPCA with $n_{\text{levels}} = 1$, where $\boldsymbol{\Lambda}_1$ degenerates to the first s_1 eigenvalues corresponding to the s_1 eigenvectors (principal components) \mathbf{H}_1 of the kernel matrix \mathbf{K}_1 .

3.2. Optimization algorithm

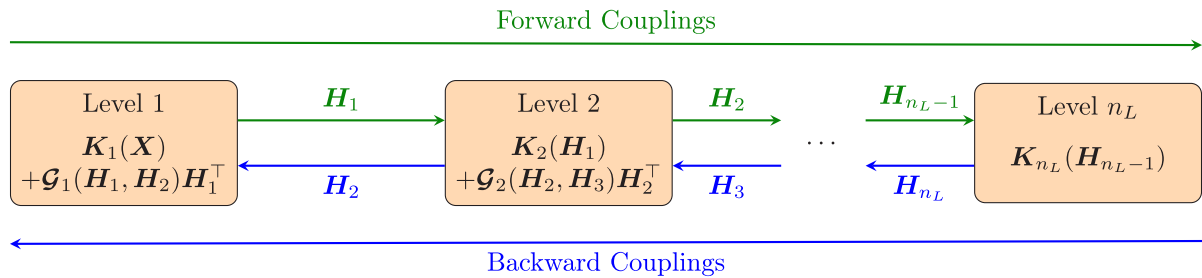


Fig. 2. Graphical illustration of the DKPCA dual problem (6) with n_L levels. Each arrow goes from the level that is characterized by the corresponding hidden features to the level where it is used as input. DKPCA introduces not only *forward* couplings (green arrows), but also *backward* couplings (blue arrows) between the levels. For simplicity, $n_j = 1$ in the diagram.

Algorithm 1 DKPCA Training using PGD. The stepsize α is selected via backtracking for each variable.

```

1: function DEEPKPCA( $\{x_i\}_{i=1}^N, \epsilon > 0$ )
2:   Compute  $K_1$  from  $\{x_i\}_{i=1}^N$ 
3:   Initialize  $\{H_1^1, \dots, H_{n_{\text{levels}}}^1, \Lambda_1^1, \dots, \Lambda_{n_{\text{levels}}}^1\}$ 
4:    $k \leftarrow 0$ 
5:   repeat
6:      $k \leftarrow k + 1$ 
7:     Compute  $K_j^k$  from  $H_{j-1}^k, \forall j = 2, \dots, n_{\text{levels}}$ 
8:     Compute  $G_j^k$  from  $H_j^k, H_{j+1}^k, \forall j = 1, \dots, n_{\text{levels}} - 1$ 
9:     Compute the residuals in (9)
10:     $H_j^{k+1} \leftarrow \Pi_{\text{St}(s_j, N)} \left( H_j^k - \alpha_k \nabla_{H_j} \tilde{J} \left( H_1^k, \dots, H_{n_{\text{levels}}}^k, \Lambda_1^k, \dots, \Lambda_{n_{\text{levels}}}^k \right) \right)$  ▷ Update for all levels
11:     $\Lambda_j^{k+1} \leftarrow \Lambda_j^k - \alpha_k \nabla_{\Lambda_j} \tilde{J} \left( H_1^k, \dots, H_{n_{\text{levels}}}^k, \Lambda_1^k, \dots, \Lambda_{n_{\text{levels}}}^k \right)$  ▷ Update for all levels
12:    until  $\|H_j^{k+1} - H_j^k\|_{\text{max}} / \alpha_k \leq \epsilon$  and  $\|\Lambda_j^{k+1} - \Lambda_j^k\|_{\text{max}} / \alpha_k \leq \epsilon$  ▷ Condition for all levels
13:    return  $H_1, \dots, H_{n_{\text{levels}}}, \Lambda_1, \dots, \Lambda_{n_{\text{levels}}}$ 
14: end function

```

For general positive definite kernels k_j , (6) is interpreted as a set of eigendecompositions with optimization variables H_j coupled with previous and subsequent layers. In the algorithmic aspect, we propose to train the DKPCA by residual minimization of (6), which considers the orthogonality constraints on intra-level hidden features and results in the following constrained optimization problem:

$$\begin{aligned}
 \underset{H_j, \Lambda_j}{\text{minimize}} \quad & \tilde{J} \triangleq \frac{1}{2} \left\| \frac{1}{\eta_1} K_1 H_1 + \frac{1}{\eta_2} G_1(H_1, H_2) - H_1 \Lambda_1 \right\|_F^2 + \\
 & \sum_{j=2}^{n_{\text{levels}}-1} \left\| \frac{1}{\eta_j} K_j(H_{j-1}) H_j + \frac{1}{\eta_{j+1}} G_j(H_j, H_{j+1}) - H_j \Lambda_j \right\|_F^2 + \\
 & \left\| \frac{1}{\eta_{n_{\text{levels}}}} K_{n_{\text{levels}}}(H_{n_{\text{levels}}-1}) H_{n_{\text{levels}}} - H_{n_{\text{levels}}} \Lambda_{n_{\text{levels}}} \right\|_F^2 \quad (9) \\
 \text{subject to} \quad & H_j^\top H_j = I_{s_j}, \quad \forall j = 1, \dots, n_{\text{levels}},
 \end{aligned}$$

where \tilde{J} denotes the optimization objective and the residual error is adopted as the Frobenius norm $\|\cdot\|_F$. During the training, the hidden features not only flow forward from the previous level, but also backward from the subsequent level, as H_j comes from the eigendecomposition depending on H_{j-1} and H_{j+1} in a level-wise fashion.

The constraint set for the hidden features of level j is the Stiefel manifold $\text{St}(s_j, N) = \{H_j \in \mathbb{R}^{N \times s_j} \mid H_j^\top H_j = I_{s_j}\}$. Optimization of (9) can be tackled by the Projected Gradient Descent (PGD) algorithm, where the iterates for H_j are specified by $H_j^{k+1} = \Pi_{\text{St}(s_j, N)} \left(H_j^k - \alpha_k \nabla_{H_j} \tilde{J} \left(H_1^k, \Lambda_1^k, \dots, H_{n_{\text{levels}}}^k, \Lambda_{n_{\text{levels}}}^k \right) \right)$, in the $(k + 1)$ th iteration, with $\Pi_{\text{St}(s_j, N)}$ being the Euclidean projection onto the Stiefel manifold, where α_k is the stepsize selected via backtracking and the projection in PGD can be computed by the compact SVD of H_j^k . This algorithm is detailed in Algorithm 1. In Algorithm 1, the deep principal components H_j and deep eigenvalues Λ_j are iteratively updated, during which the forward and backward couplings across levels are realized in the deep kernel machine.

In comparison, shallow (with one level) KPCA is solved by an eigendecomposition of the kernel matrix, and the existing deep KPCA (Deng

et al., 2019) extends its levels simply in a forward-wise manner by sequentially learning a KPCA in each level. Note that the optimization of DKPCA involves end-to-end training jointly optimizing all levels, which is a significantly different training scheme than the existing KPCA methods. In fact, (Deng et al., 2019) can be simply seen as a specific initialization of our DKPCA in (9), where the hidden features H_j of each level are initialized with the eigenvectors from sequential level-wise KPCA from Deng et al. (2019). Regarding computational analysis, shallow KPCA requires the computation of the first \bar{s} eigenvectors of the kernel matrix with complexity $\mathcal{O}(\bar{s}N^2)$, and the deep KPCA (Deng et al., 2019) consisting of n_{levels} levels requires sequential eigendecompositions with complexity $\mathcal{O}(n_{\text{levels}}\bar{s}N^2)$, with $\bar{s} = \sum_{j=1}^{n_{\text{levels}}} s_j$. The main complexity in Algorithm 1 is the projection step in PGD computing the compact SVD with complexity $\mathcal{O}(n_{\text{levels}}\bar{s}N^2)$. Optimization of (9) can also be addressed with more efficient algorithms than classical PGD. An alternative is the Riemannian Adam (Becigneul & Ganea, 2019), which can exploit the Stiefel manifold constraints of (9) without performing SVD, resulting in significantly less expensive iterations, as shown in Table B.3 in Appendix B.2.

3.3. Generative DKPCA

In linear PCA, performing reconstruction is straightforward by a linear basis transformation, while the nonlinear KPCA faces the well-known pre-image challenges in reconstructions (Mika et al., 1999). The proposed DKPCA employs multiple nonlinear feature maps and consists of multiple latent spaces, posing even greater challenges for the reconstruction. We propose a procedure for generative DKPCA from the sampled hidden features $h^{(j)}$ in latent spaces with parametric feature maps φ_j of each level, which also induces a positive definite kernel matrix (Pandey et al., 2021; Suykens et al., 2002). We also describe how the proposed generative model can facilitate the exploration of the role of the deep eigenvectors of each level.

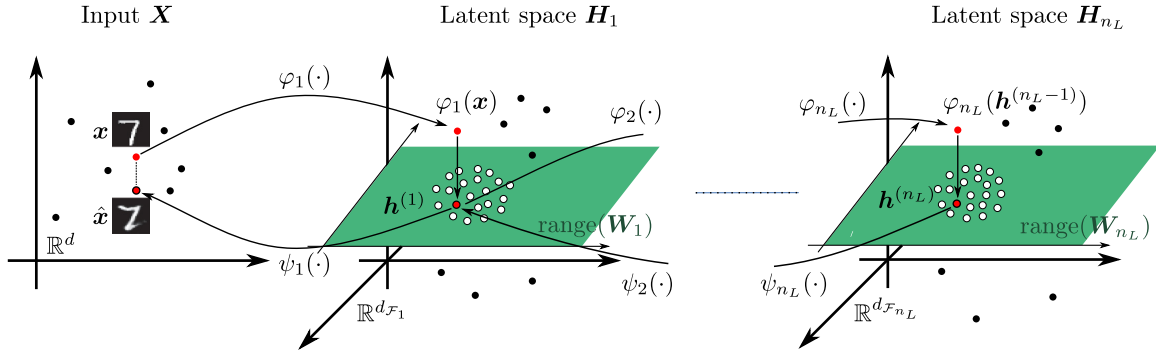


Fig. 3. Overview of generative DKPCA with n_L levels. Multiple latent spaces are considered with multi-level hidden features $h^{(j)}$, $\forall j = 1, \dots, n_L$. The feature maps φ_j are indicated with arrows going from left to right. The generative model employs the pre-image maps ψ_j , represented by the arrows going from right to left. The dashed line in input space represents the reconstruction error. The projecting vector in latent spaces indicates the projections in the corresponding s_j -dimensional latent subspace.

Given the learned $h^{(j)}$, we consider a generative objective introducing one term per level to the objective (5) for a point x : $\frac{1}{2} \varphi_1(x)^T \varphi_1(x)$ for the first level and $\frac{1}{2} \varphi_j(h^{(j-1)})^T \varphi_j(h^{(j-1)})$ for level $j = 2, \dots, n_{\text{levels}}$. By the characterization of the stationary points given in Appendix A.2, a new point \hat{x} is generated through the inverse maps of the multiple levels:

$$\hat{x} = \varphi_1^{-1}(W_1 a^{(2)}), \quad (10)$$

such that $a^{(j)} = \varphi_j^{-1}(W_j a^{(j+1)})$, $j = 2, \dots, n_{\text{levels}}$, $a^{(n_{\text{levels}}+1)} = h^{(n_{\text{levels}})}$, and where φ_j is invertible with the inverse map denoted as φ_j^{-1} . Note that (10) has a similar structure to the decoder of an Autoencoder architecture. This process is visualized in Fig. 3.

In practice, it is particularly useful to employ parametric feature maps, as they can learn to well map high-dimensional complex data from the unknown training distribution. For instance, in computer vision tasks one can define a convolutional neural network as the feature map φ_1 . A transposed convolutional network ψ_1 is used in the generation formula (10) to approximate the inverse map φ_1^{-1} such that $(\psi_1 \circ \varphi_1)(x) \approx x$. In such cases when the inverse map φ_1^{-1} is unknown explicitly in advance, one can employ a learnable pre-image map to approximate the inverse map, which resembles the decoder part in an Autoencoder architecture. Thus, we add the reconstruction error, e.g., $\mathcal{L}_i(x_i, \psi_1(\varphi_1(x_i))) = \|x_i - \psi_1(\varphi_1(x_i))\|^2$ to the optimization objective \tilde{J} in (9) for the learning of the inverse feature map ψ_1 . The full objective is thereby cast as

$$\tilde{J} + \gamma \sum_{i=1}^N \mathcal{L}_i(x_i, \psi_1(\varphi_1(x_i))), \quad (11)$$

where ψ_1 is the learnable pre-image map that approximates the inverse map φ_1^{-1} , \mathcal{L}_i is the reconstruction error of sample x_i , and $\gamma > 0$ balances the reconstruction error and the residuals minimization. Besides H_j and Λ_j in \tilde{J} , the network parameters of φ_1 and ψ_1 also need to be learned. In this optimization problem, an alternating update scheme is adopted: the Adam optimizer (Kingma & Ba, 2015) is used to update the parameters of φ_1 and ψ_1 , while keeping the deep eigenvectors and eigenvalues fixed; the hidden features H_j and the corresponding eigenvalues Λ_j are updated using the DKPCA training algorithm described in Section 3.2 with φ_1 and ψ_1 fixed.

In this case, the optimization to the proposed generative model includes both the latent variables $h_i^{(j)}$ in the dual and the explicit feature map φ_1 in the primal. This combination allows both the couplings of the levels in the latent variables of each level and deep powerful parametric feature maps better suited for more complex tasks. The deep architecture of DKPCA consists of feature maps over multiple levels, where depth is given both by multiple KPCA levels and by feature maps possibly consisting of multi-layered neural networks. This generative model resolves the pre-image problem in performing reconstruction and also enables to obtain new data corresponding to

any sampling in the multiple latent spaces. For reconstruction, given any input, its hidden features (principal components) in latent spaces are first computed and are then fed to the inverse feature maps for reconstruction in the original input space. For generation, given any sampling in latent spaces, their correspondingly generated samples in the input space can be obtained through the inverse feature maps using (10). This makes it viable to explore the role of the deep eigenvectors relating to the principal components of each level, i.e., the generation of newly sampled latent variables can be investigated by changing only one latent variable (principal component) at a time, performing the traversals over these latent variables.

Besides, DKPCA also pertains the out-of-sample extension, which allows to predict unseen input data without retraining. This property is of particular interest in large-scale case for such unsupervised settings (Suykens et al., 2002), as a subset of $M \ll N$ samples are used for the efficient training and the rest $N - M$ samples can be predicted through out-of-sample extensions, as detailed in Appendix A.2. In this way, the storage complexities for the kernel matrices and the hidden features matrices of level j decrease from $\mathcal{O}(N^2)$ and $\mathcal{O}(Ns_j)$ to $\mathcal{O}(M^2)$ and $\mathcal{O}(Ms_j)$, respectively. One approach to the subset selection is to take a random subsample of M data points for the training, which is capable of well balancing both efficiency and accuracy as evaluated in Fig. 8. One can also use more sophisticated selection schemes, such as the quadratic Renyi entropy (Girolami, 2002) or the leverage score sampling (Rudi et al., 2018). The optimal selection strategy is nevertheless data-dependent in practical applications (Espinoza et al., 2003; Fanuel et al., 2021).

4. Analytical findings

In this section, first we show that the optimization problem of our method explicitly formulates a set of nonlinear equations for each level resembling an eigenvalue problem of some matrix M_j fusing the principal components of previous and subsequent levels, i.e., DKPCA introduces not only *forward* couplings, but also *backward* couplings between the levels. Further, we illustrate that the additional levels act as a regularization on the first level. Then, we apply the Eckart-Young theorem to the deep kernel machine for approximation error bounds on the kernel matrix of the given data. Finally, we show conditions under which the explained variance of DKPCA is strictly greater than the one from KPCA.

4.1. Forward and backward couplings between levels

The equations in (6) give the level-wise eigendecomposition interpretation of DKPCA, in which the forward and backward couplings between levels are embodied. The first level resembles the eigendecomposition of the regularized kernel matrix of the given data $M_1 \triangleq \frac{1}{\eta_1} K_1 + \frac{1}{\eta_2} \mathcal{G}_1(H_1, H_2)H_1^T$; the last level is the eigendecomposition of




Fig. 4. Deep approximation bounds. Illustration of Lemma 4.1. Lower (green dashed line) and upper bound (red dashed line) for the deep approximation error (black solid line) of \mathbf{K}_1 on Synth 3, for varying η_2 when (a) $\eta_2 > 0$ and when (b) $\eta_2 < 0$, (c) s_1 , and (d) s_2 . In (a)(b), the full decomposition case is considered.

the symmetric matrix $\mathbf{M}_{n_{\text{levels}}} \triangleq \frac{1}{\eta_{n_{\text{levels}}}} \mathbf{K}_{n_{\text{levels}}}$; the intermediate levels $j = 2, \dots, n_{\text{levels}} - 1$ are related to the eigendecomposition of

$$\mathbf{M}_j(\mathbf{H}_{j-1}, \mathbf{H}_j, \mathbf{H}_{j+1}) \triangleq \frac{1}{\eta_j} \mathbf{K}_j(\mathbf{H}_{j-1}) + \frac{1}{\eta_{j+1}} \mathbf{G}_j(\mathbf{H}_j, \mathbf{H}_{j+1}) \mathbf{H}_j^T, \quad (12)$$

with deep eigenvectors \mathbf{H}_j and deep eigenvalues Λ_j . Fig. 2 visualizes this process.

The optimization of DKPCA discussed in Section 3.2 is interpreted as a set of n_{levels} eigendecomposition problems, each of which (\mathbf{H}_j) depends on the hidden features of both previous (\mathbf{H}_{j-1}) and subsequent (\mathbf{H}_{j+1}) levels. In this way, information not only flows forward but also backward in the learning process, as \mathbf{M}_j has dependency on both \mathbf{H}_{j-1} and \mathbf{H}_{j+1} . This is an important property, as previous theoretical works in deep learning such as (Allen-Zhu & Li, 2023) stressed that forward propagation alone in a level-wise fashion is not enough to learn efficient deep architectures, as the levels also need to be coupled in backward directions so that more abstract representation of subsequent levels can be utilized to improve the learning of the current level. With the forward and backward couplings between levels, eigenvalue problems in (6) cannot be independently solved in series, which motivates the DKPCA training algorithm by residual minimization of the set of nonlinear equations (6) described in Section 3.2.

4.2. Deep approximation analysis

For theoretical analysis, we consider the two-level DKPCA with $k_2(\mathbf{z}, \mathbf{y}) = \mathbf{z}^T \mathbf{y}$, as the optimization can be simplified. In this case, \mathbf{M}_j does not depend on \mathbf{H}_j such that $\mathbf{M}_1(\mathbf{H}_2) \mathbf{H}_1 = \mathbf{H}_1 \Lambda_1$ and $\mathbf{M}_2(\mathbf{H}_1) \mathbf{H}_2 = \mathbf{H}_2 \Lambda_2$, where \mathbf{H}_1 and \mathbf{H}_2 are implemented as the eigenvectors in Level 1 and Level 2, respectively:

$$\begin{cases} \text{Level 1: } \left(\frac{1}{\eta_1} \mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^T \right) \mathbf{H}_1 = \mathbf{H}_1 \Lambda_1, \\ \text{Level 2: } \left(\frac{1}{\eta_2} \mathbf{H}_1 \mathbf{H}_1^T \right) \mathbf{H}_2 = \mathbf{H}_2 \Lambda_2, \end{cases} \quad (13)$$

where the first level performs KPCA of $\frac{1}{\eta_1} \mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^T$ and the second level performs KPCA of $\frac{1}{\eta_2} \mathbf{H}_1 \mathbf{H}_1^T$. Here, the second level can

be regarded as playing a regularization role: the second level leads to a regularized \mathbf{K}_1 with the regularization constant $\frac{\eta_1}{\eta_2}$. Note that \mathbf{H}_2 is unknown a priori, so one has to solve the sets of nonlinear equations (13) for both levels rather than first solving the eigenvalue problem for level 1 and then for level 2, reflecting the forward and backward dependency.

We analyze approximation error bounds for the conceived two-level architectures through the Eckart-Young theorem (Eckart & Young, 1936), as both of the matrices to be factorized are symmetric, providing additional insights into the DKPCA.

Lemma 4.1 (Error Bounds). *Applying the Eckart-Young theorem to both levels in (13) with orthonormality constraints, the following bound for the deep approximation of \mathbf{K}_1 is obtained*

$$\begin{aligned} \sqrt{\sum_{i=s_1+1}^{r_1} \lambda_i^{(1)2} - \frac{\sqrt{s_2}}{|\eta_2|}} &\leq \|\mathbf{K}_1 - \mathbf{H}_1 \Lambda_1 \mathbf{H}_1^T\|_F \\ &\leq \begin{cases} \sqrt{\sum_{i=s_1+1}^{r_1} \lambda_i^{(1)2} - \left(\frac{s_2}{\eta_2} + 2 \sum_{i=1}^{s_2} \tilde{\lambda}_i \right) \frac{1}{\eta_2}} & \eta_2 < 0, \\ \sqrt{\sum_{i=s_1+1}^{r_1} \lambda_i^{(1)2} - \left(\frac{1}{\eta_2} - 2s_1 \sum_{i=1}^{s_1} \lambda_i^{(1)} \right) \frac{s_2}{\eta_2}} & \eta_2 > 0, \end{cases} \end{aligned} \quad (14)$$

with $s_1 \leq r_1$, where $r_1 = \text{rank}(\mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^T)$ and $\tilde{\lambda}_i$ is the i th largest eigenvalue of \mathbf{K}_1 .

Lemma 4.1 gives the error of approximating the data kernel matrix \mathbf{K}_1 with the low-rank matrix of hidden features \mathbf{H}_1 of the first level as a lower bound depending on the remaining eigenvalues of \mathbf{K}_1 regularized with the matrix of hidden features \mathbf{H}_2 of the second level. The smaller η_2 , the greater the effect of the second level. On the other hand, a very large η_2 indicates high regularization on the second level, reducing its effect, in which the deep architecture behaves resembling a shallow low-rank approximation. If the number of columns s_1 of the approximating matrix is greater than the rank r_1 of the matrix to be approximated, one can choose $s_1 = r_1$ achieving an error-free approximation. See Fig. 4 for numerical evaluation and Appendix A.3.1 for the proof.

In the next Lemma, we study the cumulative explained variance given by the principal components of the considered two-level DKPCA

with comparisons to shallow KPCA, analytically showing the higher explained variance of DKPCA.

Lemma 4.2 (Explained Variance of Deep KPCA). *In the full decomposition case ($s_1 = s_2 = N$), when $\eta_2 < -\frac{1}{\tilde{\lambda}_N}$, the explained variance of the top n principal components of DKPCA in (13) is strictly greater than the variance explained by the top n principal components of shallow kernel PCA, i.e.,*

$$\frac{\sum_{j=1}^n \tilde{\lambda}_j}{\sum_{i=1}^N \lambda_i} > \frac{\sum_{j=1}^n \tilde{\lambda}_j}{\sum_{i=1}^N \tilde{\lambda}_i}, \quad (15)$$

where $\tilde{\lambda}_i > 0$ is the i th largest eigenvalue of the kernel matrix \mathbf{K}_1 , which is taken positive-definite, and λ_i is the i th largest eigenvalue of $\mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^\top$, for all $1 \leq n < N$.

The above Lemma gives conditions on η_2 under which the considered two-level DKPCA is advantageous compared to shallow KPCA in terms of the explained variance of the first n principal components. When choosing $\eta_2 < -\frac{1}{\tilde{\lambda}_N}$, where $\tilde{\lambda}_N$ is the smallest eigenvalue of the data kernel matrix, the cumulative variance explained by the first n components of the first DKPCA level is strictly greater than the variance explained by the first n components of shallow KPCA. In other words, DKPCA can capture more information in fewer components. See the next Section for associated numerical experiments and Appendix A.3.2 for the proof.

5. Numerical experiments

We present a series of experiments to assess and explore DKPCA, showing the efficacy and advantages of the proposed deep method from different aspects in the following subsections. DKPCA is implemented in Python using the PyTorch library. The code is available at <https://github.com/taraloc/deepkpc>, where all datasets used in this study and the setup details are publicly available and described in the repository.

Datasets. Both synthetic and real-world data are used to assess the proposed method with empirical evidence. Three synthetic datasets are presented: a 2D square dataset (Synth 1), a complex 2D dataset consisting of one square, two spirals and one ring (Synth 2), and a 140-dimensional multivariate Normal dataset (Synth 3), where samples are drawn randomly from mixed Gaussian distributions. For real-world data, we consider MNIST (LeCun et al., 2010), 3DShapes (Kim & Mnih, 2018), Cars3D (Reed et al., 2015), and SmallNORB (LeCun et al., 2004). In particular, we evaluate disentanglement on the 3DShapes, Cars3D, and SmallNORB, which are popular benchmarks for evaluating variation factors.

Evaluation metrics and compared methods. Different related unsupervised learning methods are adopted to comprehensively evaluate DKPCA. A comparison to the shallow KPCA is presented with the learned principal components on multiple aspects. We also consider the state-of-the-art methods β -VAE (Higgins et al., 2017), FactorVAE (Kim & Mnih, 2018), and β -TCVAE (Chen et al., 2018) for general disentangled feature learning. For quantitative evaluations, we employ the IRS metric (Suter et al., 2019), where a higher value indicates better robustness to changes in variation factors. The shared hyperparameters among all methods are fixed to be the same. For the model-specific hyperparameters, we used the suggested values in their papers. It is worth mentioning that the compared methods are sensitive to hyperparameter selections, as shown in Locatello et al. (2019). Our method does not suffer from such issue as Λ_j is automatically determined by the solution of the deep KPCA problem.

Hyperparameters and tuning setups. In unsupervised learning setups, the experiments with linear kernels are parameter-free. In the experiments comparing the explained variance of KPCA, DKPCA and Deng et al. (2019), the shared hyperparameters, i.e., the RBF bandwidth, are set the same for all levels in the methods for consistent comparisons and selected based on reconstruction error on a validation set. We fix $\eta_j = 1$, $j = 1, \dots, n_{\text{levels}}$ and $\gamma = 1$ in (11) to equally balance the AE and deep KPCA error. For the more challenging SmallNORB, we set $\gamma = 100$. The RBF kernel in the downstream supervised experiments is tuned in the range $\exp(-2)$ and $\exp(7)$ using gridsearch on the separate validation set. In the generative experiments, the maximum number of epochs for the Riemannian Adam optimizer was set to 80000 with learning rate $2 \cdot 10^{-4}$, and we keep the same encoder φ_1 and decoder ψ_1 architecture for all compared methods. Thorough setup details are given in Appendix B.1.

5.1. DKPCA provides interpretable deep principal components

This part examines the roles of each individual deep principal component and of the components in each level. Contrary to shallow KPCA owning one set of eigenvectors/eigenvalues, DKPCA have multiple sets of eigenvectors/eigenvalues for each level. Thus, the features can be represented in a more hierarchical way that benefits the interpretation explorations. In fact, via the proposed deep generative procedure (10), sampled hidden features and their pre-image mappings to the input space can be computed. By traversing the latent space in some specific dimensions, i.e., varying a single deep principal component while keeping the others fixed and generating the corresponding sample in input space, what each component learns can be observed. In DKPCA, with the extracted deep eigenvectors \mathbf{H}_j , the model can well disentangle the factors of variation in the data. This is verified quantitatively and qualitatively, comparing the traversals on the learned principal components with the state-of-the-art FactorVAE.

Notably, we show that DKPCA effectively facilitates hierarchical data exploration, as the role of each principal component in each level can be investigated through the generation of new data. Specifically, we consider images of 3D objects with different generative factors, i.e., colors, sizes, etc. For individual components, our method can find new principal components such that, when sampling along one of them, only one generative factor changes, e.g., only the object scale changes, while its color and other factors remain fixed. For the components in each level, our deep method creates a learning hierarchy: prevailing features are typically learned in the shallower levels, e.g., colors, while the deeper levels capture more subtle features, e.g., the specific object shape.

Fig. 5 summarizes the main results for 3DShapes. Detailed analysis is given in the following for each principal component in all levels and for each level separately.

Individual principal components. In Figs. 6 and 5(a), we show the traversals in the latent spaces of a DKPCA with explicit feature maps. Aside of high visual reconstruction quality in the second row, other rows show the generated images while traversing along the individual principal component of the first level ($h^{(1)}$) or of the second level ($h^{(2)}$) of the proposed DKPCA that explains the corresponding generative factor. In FactorVAE, a single latent space is obtained, and the images are generated by traversing along each dimension in the latent space of FactorVAE. For instance, in 3DShapes (Fig. 5(a)), the component in Row 3 captures the factor of wall hue, as both the floor and object hue remain almost constant. In 3DShapes, DKPCA better disentangles the scale of the object, which only slightly varies in FactorVAE. In Cars3D (Fig. 6(a)), the three factors of elevation, car type, and azimuth by DKPCA are well captured and disentangled, while FactorVAE gives entanglement in differentiating the learning of azimuth with the two other components of elevation and car type. A similar analysis is conducted for the other rows, showing that the deep components well capture the

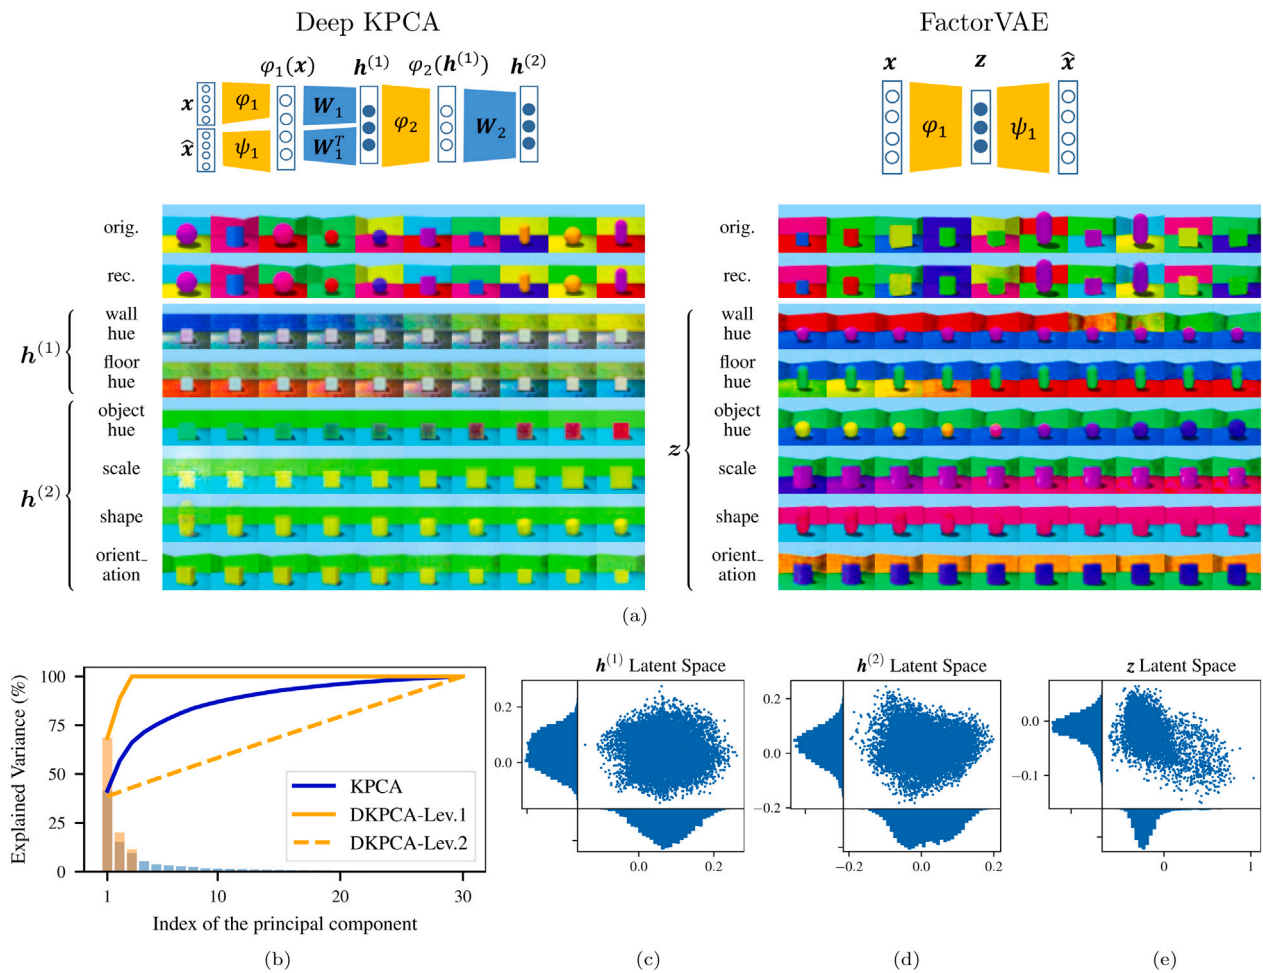


Fig. 5. Results on the 3DShapes dataset. (a) **Role of the deep principal components.** The ground-truth on the 1st row, reconstructions on the 2nd row, and traversals on other rows in the latent spaces induced by DKPCA and FactorVAE. The factors extracted by DKPCA are better disentangled than FactorVAE. Unlike FactorVAE, DKPCA shows a hierarchy of details, where the second level learns more complex factors of variation than the first level. (b) **Explained variance (%)** of both DKPCA and shallow KPCA using the same kernel. DKPCA captures considerably greater explained variance (informative features) in the first principal components than KPCA, where the lines denote the cumulative explained variance and the bars denote the variance explained by each component. (c)(d)(e) **Scatter plots of the latent variable distribution**, where DKPCA learns one latent space for each level. The FactorVAE distribution shows partial irregularity, while the distributions learned by DKPCA follow a more compact Gaussian profile, centered around the origin in the second level.

factors of variation of the data. Besides, thanks to the eigenvalues Λ_j obtained in the optimization, DKPCA can identify an ordering of the components, providing a way to reflect their relative importance. This cannot be done with the considered VAE-based methods (Chen et al., 2018; Higgins et al., 2017; Kim & Mnih, 2018).

Principal components in each level. Besides individual components, we further explore the level-wise interpretation of the learned deep principal components in DKPCA. In Fig. 5(a), the two components of the first level capture the background, which corresponds to the factors of the highest variation, i.e., the wall and floor hue, as they involve the most pixels in the images. The two components of the second level capture subtle characteristics of the object, e.g., scale and orientation, as the deeper components capture generative factors for more detailed information with less variation among samples. In other words, DKPCA learns a hierarchy of abstraction in its deep components, from less abstract, i.e., background, to more abstract, i.e., object. Similar conclusions hold for Cars3D (Fig. 6(a)): the first level learns the car type, which is the factor of highest variation, while the second level learns more sophisticated factors capturing the elevation and azimuth of the car.

Disentanglement learning. A quantitative evaluation of disentangled feature learning is performed by comparing with the state-of-the-art methods β -VAE (Higgins et al., 2017), FactorVAE (F-VAE) (Kim & Mnih,

2018), and β -TCVAE (Chen et al., 2018) on the commonly used IRS metric (Suter et al., 2019). We also incorporate the KPCA approach for comparison between the classical shallow kernel machine and our proposed deep architecture of DKPCA. The studied DKPCA architecture has $n_{levels} = 2$, $s_1 = s_2$ set to the number of generative factors, and the latent representation of a data point x_i is given by the concatenation of $h_1^{(1)}$ and $h_2^{(2)}$. The dimension of the latent space of the compared methods is set to $s_1 + s_2$. Fig. 7 gives the performance evaluation with models trained on a subset of $N = 200$ samples. The proposed DKPCA shows overall favorable performance for disentanglement learning on the tested datasets, notably outperforming the state-of-the-art VAE-based methods in Cars3D. Those advantageous results of DKPCA achieved under this setting reflects better sample efficiency in this set of experiments: from only hundreds of data points, the DKPCA can learn more disentangled representations than the compared data-hungry deep learning methods. In real-life scenarios, this property can be of particular interest, as the training examples might be available in limited quantity or expensive to collect, so models better capturing the true generative factors from a limited number of data are desirable.

DKPCA can be implemented with out-of-sample extensions for large-scale cases by selecting a subset $M \ll N$ for training and then obtaining the latent representations of the remaining data. To evaluate the performance of DKPCA on the full dataset, in Fig. 8 we evaluate the

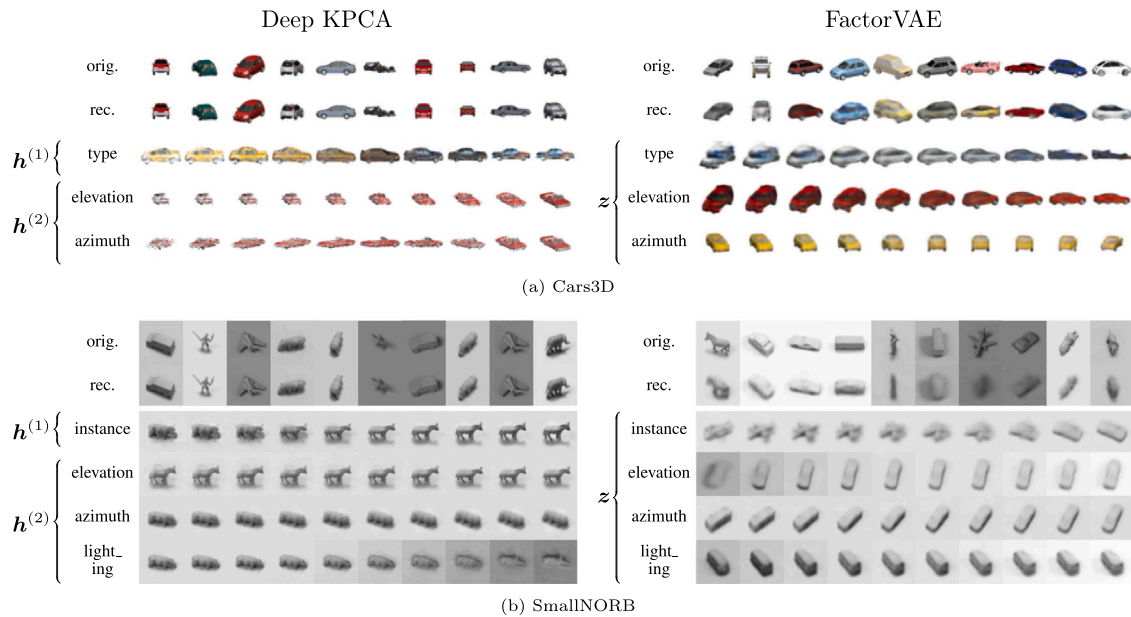


Fig. 6. Role of the deep principal components. First row: ground-truth. Second row: reconstructions. Other rows: traversals in the latent spaces induced by DKPCA. The DKPCA shows a hierarchy of details, where the second level learns more complex factors of variation than the first level.




Fig. 7. Disentanglement with small data (†). Distribution (mean and standard deviation) of disentanglement scores (IRS) for different methods with $N = 200$ samples. Higher is better (†).




Fig. 8. Disentanglement with large-scale extensions (†). Disentanglement score (IRS) on the full datasets. A subset of $M = 200$ is selected for training (D)KPCA and the remaining data is inferred with out-of-sample extensions, which is particularly beneficial under limited computational resources.

entire corresponding datasets through out-of-sample extensions using $M = 200$ samples for the training. The results shows a higher mean IRS is attained over all compared VAE-based methods which are trained on the full dataset. The results also demonstrate that DKPCA consistently outperforms the shallow KPCA. This comparison further verifies the disentanglement of the hidden features learned by our method, as well as its sample efficiency: only hundreds of samples are needed by DKPCA to effectively learn disentangled representations and outperform the deep learning methods trained on thousands of data points.

5.2. DKPCA learns more informative features

In this section, we further investigate the features learned by DKPCA. DKPCA gives higher explained variance than shallow KPCA, indicating that more information is captured in fewer components. We therefore show the superiority of DKPCA as a feature extractor for downstream supervised tasks for multiple data types. We also investigate the problem of selecting the number of principal components in each level and the number of levels, providing a selection strategy in an




Fig. 9. Interpretation of the deep eigenvalues. Explained variance (%) of DKPCA. (a) The compared method is kernel PCA with RBF kernel with the same bandwidth. Our method is able to capture considerably greater explained variance in the first principal component than shallow KPCA, showing that the proposed deeper architecture outputs more informative principal components even with the same kernel function as the shallow KPCA. (b) Illustration of Lemma 4.2: the first DKPCA level maintains higher cumulative explained variance than KPCA for all n , capturing more information in fewer components. (c) Four-level DKPCA with RBF kernels on MNIST. In all plots, bars: explained variance, lines: cumulative explained variance.

unsupervised setting, in contrast with typical trial and error tuning in deep learning. Besides, we also compare with the extracted principal components which are learned from the deep KPCA in Deng et al. (2019) employing a level-wisely sequential learning scheme.

Deep eigenvalues. As presented in Section 3.2, deep eigenvalues Λ_j are learned by DKPCA in different levels $j = 1, \dots, n_{\text{levels}}$, compared to Λ of the single level in shallow KPCA. We now investigate the learned deep eigenvalues in terms of the percentage of variance explained and compare with shallow KPCA, where the nonlinear case is considered by using the RBF kernel in all levels. Figs. 5(b), 9(a), and 9(b) plot the variance explained by each component by DKPCA (orange bars) and by shallow KPCA (blue bars), as well as the cumulative variance explained by DKPCA (orange line) and by shallow KPCA (blue line).

In Fig. 9(a) for Synth 2 with 30 components in each level, it shows that the cumulative explained variance reaches almost 100% after around 10 deep principal components, while a much slower explained variance growth in the shallow case. Even if both methods use the same kernels, the first principal component of DKPCA explains around 20% of the variance compared to only around 8% of KPCA. This experiment shows that our method can lead to more informative principal components, ultimately resulting in a more powerful representation in fewer components with the deep architecture. Comparing the deep eigenvalues Λ_1 (solid orange line) of the first level with the ones Λ_2 (dotted orange line) of the second level, the former shows faster initial growth, while the latter gives a flatter cumulative explained variance. A similar analysis is conducted for 3DShapes in Fig. 5(b), while Fig. 9(b) presents the numerical evaluation of Lemma 4.2 in the two-level DKPCA with RBF first level and linear second level on Synth 3, with η_2 chosen to be the largest value satisfying the conditions of 4.2.

Additionally, a 4-level DKPCA with 10 principal components in each level is trained on the handwritten digit images dataset MNIST (LeCun et al., 1998) in Fig. 9(c): the first and the second levels follow a similar pattern, and each subsequent level shows a flatter curve with increasingly higher explained variance in the top components. The fourth level explains almost the entire variance in the first few components, indicating that the current four levels are sufficient. In this way, the minimum number of levels to fully explain a given dataset can be determined. This observation can also be a useful suggestion for tuning the kernel settings in the different levels: the kernel settings might need to be better tuned when introducing additional levels does not lead to a sufficient increase in explained variance.

We additionally compare with the sequential deep KPCA of Deng et al. (2019) through the ratio of explained variance. The results, shown in Fig. B.2 in Appendix B.2, demonstrate that, with our end-to-end optimization with both forward and backward learning in the levels, DKPCA can attain higher explained variance in the first principal components than straightforwardly applying KPCA sequentially (Deng et al., 2019).

Selection of principal components in each level. Contrary to shallow KPCA, different numbers of principal components can be selected for each level in deep architectures of DKPCA. In this experiment, we train a two-level DKPCA, introduced in Eq. (13), with linear kernels on the synthetic datasets to investigate the influence of the numbers of selected principal components s_1 and s_2 of the first and second levels, respectively. In practice, one would like to select the smallest number of principal components to suffice the required small enough reconstruction error that depends on the specific applications, so a general method for selection of s_j is needed for practitioners. This selection can be performed by analyzing the relative importance of each deep component through its explained variance.

A two-dimensional synthetic dataset located as a noisy square is exemplified (Synth 1). As shown in Fig. 10(a), the eigenvalues of the first level drop distinctively after the second principal component, and the percentages of explained variance by the first and second component are similar. This is consistent with the ground-truth properties of this two-dimensional dataset. In Fig. 10(b), the reconstruction error decreases with s_2 increasing and shows its largest drop after the first two components in the second level, where the MSE reaches 0 with $s_1 = s_2 = 6$. In fact, our method can always achieve 0 reconstruction error in the case of the full decomposition with $s_1 = s_2 = N$, also as verified on the real-world 3DShapes in Fig. 11. For 3DShapes, the ground-truth number of variation factors is 6, so the cumulative explained variance climbs quickly as most variance has been captured by only a few components. The reconstruction error shows the opposite behavior, dropping sharply after around 10 principal components and reaching 0 for the full decomposition. Such evaluations are conducted in an unsupervised setting, and thus practitioners can accordingly use these evaluations to determine s_j of the DKPCA architecture for faithful reconstructions.

Extracted principal components for downstream tasks. KPCA is often used as a feature extraction step for downstream supervised tasks. Similarly, DKPCA can extract multiple levels of disentangled features that can facilitate different tasks. Specifically, it has been suggested that disentangled features could be useful for the supervised downstream problems due to the compact structure of the representation of the input distribution (Locatello et al., 2019). Besides the shallow PCA and KPCA, we also compare with the deep KPCA method proposed in Deng et al. (2019) in the downstream task. We fed the concatenation of the deep representation learned by the unsupervised two-level deep KPCA machines of our work and of Deng et al. (2019) to a linear classifier/regressor and compared with shallow (K)PCA with s principal components using the same overall number of components, i.e. $s_1 + s_2 = s$. For all datasets $s_1 = 3, s_2 = 2$. The nonlinear kernel machines, i.e., KPCA, DKPCA and Deng et al. (2019), employ RBF kernels; hyperparameters are tuned on a validation set using a 60/20/20 split for training/validation/test sets.

Results are shown in Table 1. Our proposed DKPCA outperforms all compared methods, including the shallow KPCA and the deep KPCA

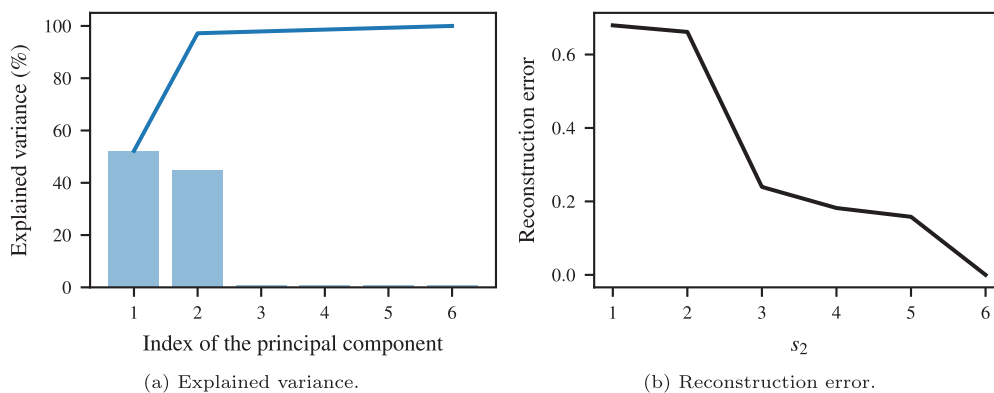


Fig. 10. Explained variance by the first level and **reconstruction error** (training MSE) for the Synth 1 dataset with $s_1 = 6$. This experiment shows the minimum number of components such that the approximation error is small enough so that practitioners have a guarantee on the faithfulness of the representation learned by the proposed model. For this dataset, the reconstruction error is 0 with $s_1 = s_2 = 6$.

Table 1
Comparison of test performance for classification/regression and disentangled feature learning by DKPCA on real-world datasets of various data types. Higher scores (↑) are better for ACC (%) and WINDIN, lower scores (↓) are better for RMSE. The best performance is in bold. All datasets are UCI datasets from [Dua and Graff \(2017\)](#).

Dataset	Metric	PCA	KPCA	Deng et al.	DKPCA
Diabetes	ACC(↑)	70.83	67.89	71.42	72.02
	WINDIN(↑)	0.001	0.082	0.001	0.287
Ionosphere	ACC(↑)	86.09	92.17	88.70	93.04
	WINDIN(↑)	0.001	0.22	0.001	1.05
Liver	ACC(↑)	70.94	72.65	73.50	74.36
	WINDIN(↑)	0.001	0.069	0.001	0.72
Cholesterol	RMSE(↓)	61.67	61.65	61.42	60.59
	WINDIN(↑)	0.0001	0.0001	0.0001	0.003
Yacht Hydrodynamics	RMSE(↓)	8.41	8.42	8.38	8.02
	WINDIN(↑)	0.0001	0.0001	0.0001	0.19




Fig. 11. Full decomposition for a subset of 3DShapes ($N = 480$). Cumulative explained variance (%) from the deep eigenvalues of the first level of KPCA and reconstruction error. A sharp increase in the explained variance corresponds to a distinctive drop in reconstruction error, which reaches 0 for the full decomposition.

proposed in [Deng et al. \(2019\)](#) in all evaluated metrics. In general, deep kernel machines, i.e., DKPCA and [Deng et al. \(2019\)](#), achieve better classification and regression results than the shallow KPCA, showing the potential of deep KPCA methods. The WINDIN metric ([Do & Tran, 2022](#)) evaluates the disentanglement of a representation z when the ground truth factors of variations are not known: it measures both the informativeness and the separability of the representation through the conditional mutual information between the input x and its latent representation z . DKPCA produces significantly more disentangled representations; for instance, in the Liver dataset DKPCA improves the WINDIN by approximately 10 times over KPCA. The sequentially learned deep KPCA method in [Deng et al. \(2019\)](#) displays inferior performance in disentanglement. Overall, DKPCA leads to better supervised performances and disentanglement, showing the improved informativeness of our DKPCA, which can more effectively capture the trends of the data that are most relevant for supervised prediction.

6. Discussion and conclusion

6.1. Discussion

Our proposed DKPCA establishes a novel framework for deep non-linear principal component analysis by leveraging the RKM formulation. DKPCA exploits the Fenchel–Young inequality introducing conjugate feature duality, and extends the classical shallow KPCA to multiple levels, where both neural network feature mappings and kernel functions can be adopted in different levels for flexible modeling. In contrast to shallow KPCA involving a single eigendecomposition to the kernel matrix, DKPCA gives different eigenvalue problems across levels and yields the so-called deep eigenvectors and deep eigenvalues, as characterized by the stationary conditions. DKPCA can be applied to general feature learning tasks in place of classical KPCA or VAE-based methods in various applications. Conventional KPCA may need many components to attain a high explained variance, while DKPCA can capture information more efficiently in fewer components. Compared to the black-box optimization in deep learning-based methods, the optimization problem of DKPCA explicitly formulates a set of nonlinear equations for each level resembling an eigenvalue problem.

DKPCA formalizes the couplings between levels in terms of the conjugated hidden features, playing the roles of principal components in the latent spaces with dual formulations. The proposed deep kernel method is not a simple forward level-wise algorithm, but the optimization of features flows backwards in the deep architecture, so that components in levels of lower abstraction can benefit from the representation learned in levels of higher abstraction. This property has been theoretically verified as essential for effective hierarchical learning, and yet has not been explored in deep kernel methods. We then devise a multi-level optimization algorithm for DKPCA, where

the deep eigenvectors and deep eigenvalues regarding the level-wise principal components are taken as optimization variables. For a specific case with two-level architectures, the optimization is simplified with solutions being the singular vectors in each level, which facilitates theoretical analysis for greater insights: the Eckart-Young theorem is applied to establish approximation bounds, interpreting the role of the second level as a regularizer, and the explained variance by DKPCA is analytically compared with shallow KPCA.

We also develop the generative DKPCA, so that hidden features in multiple levels can be sampled from the latent spaces and their correspondingly newly generated data can be attained. The role of each component or each level can be explored by traversing it in the latent space and keeping the others fixed, providing diversified aspects to explore the meaning of principal components and the variation factors of data. The pre-image problem is a well-known challenging problem in KPCA, and its solution to general cases of multi-level KPCA was not investigated before. In DKPCA, we incorporate the reconstruction errors, minimized to approximate the pre-image feature mappings, so that the reconstruction procedures can be conducted. Compared to the generation and reconstruction in VAE-based methods, DKPCA creates multiple latent spaces, which not only enhances modeling flexibility with deep architectures but also provides multi-level feature learning. Out-of-sample extensions are also allowed in DKPCA to predict unseen data. The scalability issue commonly exists in kernel-based methods, but this can be well resolved by the out-of-sample extensions owned by DKPCA. When a small subset with M samples is used in training and the rest $N - M$ samples are predicted via out-of-sample extensions, the maximal storage complexity of level j drops from $\mathcal{O}(N^2)$ to $\mathcal{O}(M^2)$.

6.2. Conclusion

In this paper, the proposed DKPCA introduces a novel deep architecture for unsupervised multi-level feature learning, where deep kernel machines and neural networks can both be exploited. DKPCA realizes forward and backward learning and provides more informative features enabling the exploration and interpretations on hierarchical feature abstractions. Both theoretical derivations and numerical evaluations verify the effectiveness of DKPCA. The data representations learned by DKPCA can be utilized in various tasks and on different types of data with promising practical values in the era of versatile data. Deep kernel machines combine the advantages of both flexibility and interpretability. In future work, several interesting directions could be worth investigating: variants of KPCA, e.g., sparse or robust KPCA, can be studied in deep architectures when sparsity or reliability are required; a more adaptive scheme of choosing the kernel functions in different levels can be favored for real-world applications, e.g., through parametric techniques; similar to deep learning, more efficient end-to-end training algorithms would be of great significance.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data/code is available on the GitHub link in the manuscript.

Acknowledgments

This work is jointly supported by ERC Advanced Grant E-DUALITY (787960), KU Leuven Grant CoE PFV/10/002, and Grant FWO G0A49 17N, EU H2020 ICT-48 Network TAILOR (Foundations of Trustworthy AI - Integrating Reasoning, Learning and Optimization), and Leuven.AI Institute. This work was also supported by the Research Foundation Flanders (FWO) research projects G086518N, G086318N, and G0A0920N; Fonds de la Recherche Scientifique — FNRS and the Fonds Wetenschappelijk Onderzoek — Vlaanderen under EOS Project No. 30468160 (SeLMA).

Appendix A. Proofs and derivations

In this section, mathematical derivations to the modeling, optimization, and analytical properties of the proposed DKPCA are elaborated. DKPCA establishes a novel deep architecture of KPCA, which has long been an important unsupervised feature learning methodology. Appendix A.1 provides the formulations leading to the optimization interpreted by a set of eigendecompositions. It demonstrates how DKPCA leverages the RKM formulations bridging neural networks and kernels and enjoys the merits of flexible deep architectures and more interpretable kernel methods. In what follows, technical details of the generative modeling are presented in Appendix A.2, showing promising potentials for versatile scenarios in real-world applications. Proofs for the lemmas in Section 4.2 are given in Appendix A.3, providing more details and insights towards the proposed DKPCA under the considered settings with analytical properties.

A.1. Derivation of DKPCA

The objective (5) of DKPCA in the primal formulations is given by the compositions of latent spaces of multiple levels, and its dual formulations can be attained by characterizing the stationary points to (5):

$$\left\{ \begin{array}{l} \frac{\partial J}{\partial \mathbf{h}_i^{(1)}} = 0 \Rightarrow \mathbf{W}_1^\top \varphi_1(\mathbf{x}_i) = \Lambda_1 \mathbf{h}_i^{(1)} \\ \qquad \qquad \qquad - \frac{\partial}{\partial \mathbf{h}_i^{(1)}} \left[\varphi_2(\mathbf{h}_i^{(1)})^\top \mathbf{W}_2 \mathbf{h}_i^{(2)} \right], \\ \frac{\partial J}{\partial \mathbf{W}_1} = 0 \Rightarrow \mathbf{W}_1 = \frac{1}{n_1} \sum_{i=1}^N \varphi_1(\mathbf{x}_i) \mathbf{h}_i^{(1)\top}, \\ \frac{\partial J}{\partial \mathbf{h}_i^{(j)}} = 0 \Rightarrow \mathbf{W}_j^\top \varphi_j(\mathbf{h}_i^{(j)}) = \Lambda_j \mathbf{h}_i^{(j)} \\ \qquad \qquad \qquad - \frac{\partial}{\partial \mathbf{h}_i^{(j)}} \left[\varphi_{j+1}(\mathbf{h}_i^{(j)})^\top \mathbf{W}_{j+1} \mathbf{h}_i^{(j+1)} \right], \\ \qquad \qquad \qquad \forall j = 2, \dots, n_{\text{levels}} - 1, \\ \frac{\partial J}{\partial \mathbf{W}_j} = 0 \Rightarrow \mathbf{W}_j = \frac{1}{n_j} \sum_{i=1}^N \varphi_j(\mathbf{h}_i^{(j-1)}) \mathbf{h}_i^{(j)\top}, \\ \qquad \qquad \qquad \forall j = 2, \dots, n_{\text{levels}} - 1, \\ \frac{\partial J}{\partial \mathbf{h}_i^{(n_{\text{levels}})}} = 0 \Rightarrow \mathbf{W}_{n_{\text{levels}}}^\top \varphi_{n_{\text{levels}}}(\mathbf{h}_i^{(n_{\text{levels}}-1)}) = \Lambda_{n_{\text{levels}}} \mathbf{h}_i^{(n_{\text{levels}})}, \\ \frac{\partial J}{\partial \mathbf{W}_{n_{\text{levels}}}} = 0 \Rightarrow \mathbf{W}_{n_{\text{levels}}} = \frac{1}{n_{n_{\text{levels}}}} \sum_{i=1}^N \varphi_{n_{\text{levels}}}(\mathbf{h}_i^{(n_{\text{levels}}-1)}) \\ \qquad \qquad \qquad \times \mathbf{h}_i^{(n_{\text{levels}})\top}. \end{array} \right. \tag{A.1}$$

By eliminating the weight matrices \mathbf{W}_j , one obtains the following non-linear equations in the hidden features $\mathbf{h}_i^{(j)}$ (see Box II): with $j = 2, \dots, n_{\text{levels}} - 1$.

By organizing (A.2) into matrices, the dual formulation of DKPCA in (6) is obtained equivalently.

A.2. Derivation of generative DKPCA

For the challenging pre-image problem for multi-level nonlinear PCA, we propose a procedure for generative DKPCA from the sampled hidden features $\mathbf{h}^{(j)}$ in latent spaces with explicit feature maps: the feature map φ_j of each level is known and can also be parametric with learnable parameters.

Assume that φ_j is invertible, with the inverse map denoted as φ_j^{-1} , and that $\mathbf{h}^{(n_{\text{levels}})}$ is given, which can be the hidden feature vector of a training or test point, or newly sampled from the latent space. First, given the learned $\mathbf{h}_i^{(j)}$ from the training, we introduce an additional term per level to the objective (5) for a point \mathbf{x} : $\frac{1}{2} \varphi_1(\mathbf{x})^\top \varphi_1(\mathbf{x})$ for the first level and $\frac{1}{2} \varphi_j(\mathbf{h}^{(j-1)})^\top \varphi_j(\mathbf{h}^{(j-1)})$ for level $j = 2, \dots, n_{\text{levels}}$.

$$\left\{ \begin{array}{l} \text{Level 1:} \\ \text{Level } j: \\ \text{Level } n_{\text{levels}}: \end{array} \right. \begin{array}{l} \frac{1}{\eta_1} \sum_{n=1}^N \mathbf{h}_n^{(1)} k_1(\mathbf{x}_n, \mathbf{x}_i) + \frac{1}{\eta_2} \sum_{n=1}^N \frac{\partial k_2(\mathbf{h}_i^{(1)}, \mathbf{h}_n^{(1)})}{\partial \mathbf{h}_i^{(1)}} \mathbf{h}_n^{(2)\top} \mathbf{h}_i^{(2)} \\ \frac{1}{\eta_j} \sum_{n=1}^N \mathbf{h}_n^{(j)} k_j(\mathbf{h}_n^{(j)}, \mathbf{h}_i^{(j)}) + \frac{1}{\eta_{j+1}} \sum_{n=1}^N \frac{\partial k_{j+1}(\mathbf{h}_i^{(j)}, \mathbf{h}_n^{(j)})}{\partial \mathbf{h}_i^{(j)}} \mathbf{h}_n^{(j+1)\top} \mathbf{h}_i^{(j+1)} \\ \frac{1}{\eta_{n_{\text{levels}}}} \sum_{n=1}^N \mathbf{h}_n^{(n_{\text{levels}})} k_{n_{\text{levels}}}(\mathbf{h}_n^{(n_{\text{levels}}-1)}, \mathbf{h}_i^{(n_{\text{levels}}-1)}) \end{array} = \begin{array}{l} \Lambda_1 \mathbf{h}_i^{(1)}, \\ \Lambda_j \mathbf{h}_i^{(j)}, \\ \Lambda_{n_{\text{levels}}} \mathbf{h}_i^{(n_{\text{levels}})}, \end{array} \quad (\text{A.2})$$

Box II.

Characterizing the stationary points w.r.t. $\varphi_1(\mathbf{x})$ and $\varphi_j(\mathbf{h}^{(j-1)})$, we obtain

$$\left\{ \begin{array}{l} \frac{\partial J}{\partial \varphi_1(\mathbf{x})} = 0 \\ \frac{\partial J}{\partial \varphi_j(\mathbf{h}^{(j-1)})} = 0 \end{array} \right. \Rightarrow \begin{array}{l} \varphi_1(\mathbf{x}) = \mathbf{W}_1 \mathbf{h}^{(1)}, \\ \varphi_j(\mathbf{h}^{(j-1)}) = \mathbf{W}_j \mathbf{h}^{(j)}, \quad \forall j = 2, \dots, n_{\text{levels}}, \end{array} \quad (\text{A.3})$$

so that the feature map $\varphi_j(\cdot)$ of each level can be calculated from the given hidden features $\mathbf{h}^{(j)}$. DKPCA then generates new samples through the inverse maps of the multiple levels, and accordingly a generated sample $\hat{\mathbf{x}}$ is attained through φ_1^{-1} in the first level that maps $\mathbf{W}_1 \mathbf{h}^{(1)}$ back to the input space, as shown in (10).

In case the inverse map φ_1^{-1} is unknown explicitly, one can learn a pre-image map by minimization of the AutoEncoder reconstruction as described in Section 3.3.

We also developed an extension to attain the hidden features in each level corresponding to an out-of-sample point \mathbf{x}^* from the first, third, and fifth equations in (A.1). For the two-level case with linear k_2 , where it is more straightforward to obtain the out-of-sample extension, we obtain by eliminating the interconnection matrices: $\mathbf{h}^{(2)*} = \frac{1}{\eta_1 \eta_2} (\Lambda_2 - \frac{1}{\eta_2} \mathbf{H}_2^\top \mathbf{H}_1 \Lambda_1^{-1} \mathbf{H}_1^\top \mathbf{H}_2)^{-1} \mathbf{H}_2^\top \mathbf{H}_1 \Lambda_1^{-1} \mathbf{H}_1^\top \Phi_1 \varphi_1(\mathbf{x}^*)$ and $\mathbf{h}^{(1)*} = \Lambda_1^{-1} (\frac{1}{\eta_1} \mathbf{H}_1^\top \Phi_1 \varphi_1(\mathbf{x}^*) + \frac{1}{\eta_2} \mathbf{H}_2^\top \mathbf{H}_1 \mathbf{h}^{(2)*})$.

A.3. Proof of deep approximation analysis

In this section, we give the proofs of Lemmas 4.1 and 4.2 in the two-level case of (13).

A.3.1. Proof of approximation bounds

With the level-wise SVD interpretation to the discussed two-level cases in (13), the Eckart-Young theorem can be applied to both levels, deriving the approximation errors:

$$\left\{ \begin{array}{l} \left\| \mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^\top - \mathbf{H}_1 \Lambda_1 \mathbf{H}_1^\top \right\|_F = \sqrt{\sum_{i=s_1+1}^{r_1} \lambda_i^{(1)2}} \\ \left\| \frac{1}{\eta_2} \mathbf{H}_1 \mathbf{H}_1^\top - \mathbf{H}_2 \Lambda_2 \mathbf{H}_2^\top \right\|_F = \sqrt{\sum_{i=s_2+1}^{r_2} \lambda_i^{(2)2}} \end{array} \right. \quad (\text{A.4})$$

with $r_1 = \text{rank}(\mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^\top)$ and $r_2 = \text{rank}(\frac{1}{\eta_2} \mathbf{H}_1 \mathbf{H}_1^\top)$. We fix $\eta_1 = 1$ and vary the regularization factor η_2 . With orthonormality constraints in the second level, $\left\| \mathbf{H}_2 \mathbf{H}_2^\top \right\|_F = \sqrt{s_2}$, the lower bound in Lemma 4.1 is obtained.

Using the orthonormality constraints of the second level, we square (A.4) and rewrite it as

$$\begin{aligned} \left\| \mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^\top - \mathbf{H}_1 \Lambda_1 \mathbf{H}_1^\top \right\|_F^2 &= \text{Tr} \left(\mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^\top - \mathbf{H}_1 \Lambda_1 \mathbf{H}_1^\top \right)^2 \\ &= \left\| \mathbf{K}_1 - \mathbf{H}_1 \Lambda_1 \mathbf{H}_1^\top \right\|_F^2 + \text{Tr} \left[\left(\frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^\top \right)^2 + \frac{2}{\eta_2} \mathbf{H}_2^\top \mathbf{K}_1 \mathbf{H}_2 - \frac{2}{\eta_2} \mathbf{H}_2^\top \mathbf{H}_1 \Lambda_1 \mathbf{H}_1^\top \mathbf{H}_2 \right] \\ &= \left\| \mathbf{K}_1 - \mathbf{H}_1 \Lambda_1 \mathbf{H}_1^\top \right\|_F^2 + \frac{s_2}{\eta_2} + \frac{2}{\eta_2} \text{Tr}(\mathbf{H}_2^\top \mathbf{K}_1 \mathbf{H}_2) - \frac{2}{\eta_2} \text{Tr}(\mathbf{H}_2^\top \mathbf{H}_1 \Lambda_1 \mathbf{H}_1^\top \mathbf{H}_2), \end{aligned} \quad (\text{A.5})$$

with $\text{Tr}(\mathbf{H}_2^\top \mathbf{H}_1 \Lambda_1 \mathbf{H}_1^\top \mathbf{H}_2) = \left\| \sqrt{\Lambda_1} \mathbf{H}_1^\top \mathbf{H}_2 \right\|_F^2$. By the Cauchy–Schwarz inequality, we further obtain

$$\begin{aligned} \left\| \sqrt{\Lambda_1} \mathbf{H}_1^\top \mathbf{H}_2 \right\|_F^2 &\leq \left\| \sqrt{\Lambda_1} \right\|_F^2 \left\| \mathbf{H}_1^\top \right\|_F^2 \left\| \mathbf{H}_2 \right\|_F^2 \\ &= \left\| \sqrt{\Lambda_1} \right\|_F^2 \left\| \mathbf{H}_1 \right\|_F^2 \left\| \mathbf{H}_2 \right\|_F^2 \\ &\leq \left\| \sqrt{\Lambda_1} \right\|_F^2 s_1 s_2 = \left(\sum_{i=1}^{s_1} \lambda_i^{(1)} \right) s_1 s_2. \end{aligned} \quad (\text{A.6})$$

Recalling that \mathbf{K}_1 is positive semi-definite, the inequality for the upper bound when $\eta_2 > 0$ in Lemma 4.1 is obtained by using (A.6) in (A.5). When $\eta_2 < 0$ in Lemma 4.1, with symmetric \mathbf{K}_1 , note that $\max_{\mathbf{H}_2^\top \mathbf{H}_2 = \mathbf{I}} \text{Tr}(\mathbf{H}_2^\top \mathbf{K}_1 \mathbf{H}_2) = \sum_{i=1}^{s_2} \tilde{\lambda}_i$, which gives the upper bound by combining with (A.5). Therefore, the proof of deriving the bounds for the approximation analysis in Section 4.2 is completed.

A.3.2. Proof of explained variance lemma

In the two-level architecture of (13), let $\tilde{\lambda}_i$ be the i th largest eigenvalue of \mathbf{K}_1 and λ_i be the i th largest eigenvalue of $\mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^\top$.

In the full decomposition case ($s_1 = s_2 = N$), \mathbf{H}_2 is an orthogonal matrix due to the orthogonality constraints. We denote $\mathbf{h}_{:i}^{(1)}$ the i th column of \mathbf{H}_1 , i.e., the eigenvector corresponding to λ_i . Then for each eigenvalue λ_i , we have

$$\begin{aligned} \left(\mathbf{K}_1 + \frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^\top \right) \mathbf{h}_{:i}^{(1)} &= \lambda_i \mathbf{h}_{:i}^{(1)} \\ \mathbf{K}_1 \mathbf{h}_{:i}^{(1)} &= (\lambda_i \mathbf{I}_N - \frac{1}{\eta_2} \mathbf{H}_2 \mathbf{H}_2^\top) \mathbf{h}_{:i}^{(1)} \\ \mathbf{H}_2^\top \mathbf{K}_1 \mathbf{h}_{:i}^{(1)} &= (\lambda_i \mathbf{H}_2^\top - \frac{1}{\eta_2} \mathbf{H}_2^\top) \mathbf{h}_{:i}^{(1)} \\ \mathbf{K}_1 \mathbf{h}_{:i}^{(1)} &= (\lambda_i - \frac{1}{\eta_2}) \mathbf{h}_{:i}^{(1)}, \end{aligned} \quad (\text{A.7})$$

yielding $\tilde{\lambda}_i = \lambda_i - \frac{1}{\eta_2}$. Note that we consider $1 \leq n < N$ for the cumulative explained variance in this lemma, as otherwise one would explain 100% of the variance, resulting in equality in (15). Further, λ_i is constrained to be non-negative, i.e., $\lambda_i \geq 0$, which keeps the ratio $\frac{\sum_{j=1}^n \tilde{\lambda}_j}{\sum_{j=1}^N \tilde{\lambda}_j}$ between 0 and 1 for a solid analysis of the explained variance. This leads to the condition

$$\eta_2 \leq -\frac{1}{\tilde{\lambda}_N}, \quad (\text{A.8})$$

where $\tilde{\lambda}_N$ is the smallest eigenvalue of \mathbf{K}_1 .

With (A.7), the explained variance by the first n deep principal components can be rewritten as

$$\frac{\sum_{j=1}^n \lambda_j}{\sum_{i=1}^N \lambda_i} = \frac{\sum_{j=1}^n \tilde{\lambda}_j + \frac{n}{\eta_2}}{\sum_{i=1}^N \tilde{\lambda}_i + \frac{N}{\eta_2}}, \quad (\text{A.9})$$

which is greater than the variance explained by the shallow principal components $\frac{\sum_{j=1}^n \tilde{\lambda}_j}{\sum_{i=1}^N \tilde{\lambda}_i}$ when satisfying

$$\frac{\sum_{j=1}^n \tilde{\lambda}_j + \frac{n}{\eta_2}}{\sum_{i=1}^N \tilde{\lambda}_i + \frac{N}{\eta_2}} - \frac{\sum_{j=1}^n \tilde{\lambda}_j}{\sum_{i=1}^N \tilde{\lambda}_i} = \frac{n \sum_{i=1}^N \tilde{\lambda}_i - N \sum_{j=1}^n \tilde{\lambda}_j}{\eta_2 \left(\sum_{i=1}^N \tilde{\lambda}_i \right)^2 + N \sum_{i=1}^N \tilde{\lambda}_i} > 0. \quad (\text{A.10})$$

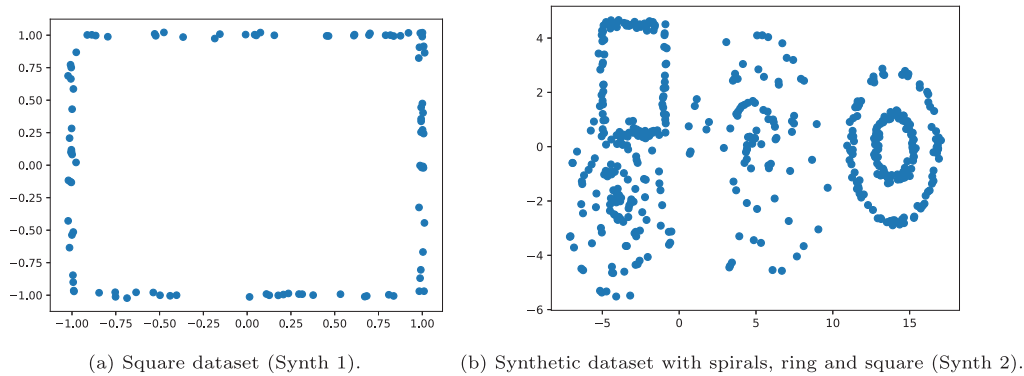


Fig. B.1. Plots of the 2D synthetic data used in the experiments.

In (A.10), there exists two cases either with denominator and numerator both positive or negative. For the former case, a positive numerator gives $\frac{\sum_{j=1}^n \tilde{\lambda}_j}{\text{Tr } \mathbf{K}_1} < \frac{n}{N}$. For $1 \leq n < N$, the ratio between the first largest n eigenvalues and the summation of all eigenvalues, i.e., $\frac{\sum_{j=1}^n \tilde{\lambda}_j}{\text{Tr } \mathbf{K}_1}$, is always greater than $\frac{n}{N}$, which is contradictory with the condition for a positive numerator. We thereby consider the case with denominator and numerator both negative, from which one obtains the conditions $\eta_2 < -\frac{N}{\text{Tr } \mathbf{K}_1}$ and $1 \leq n < \frac{N \sum_{j=1}^n \tilde{\lambda}_j}{\text{Tr } \mathbf{K}_1}$, the latter of which always holds. Combining with (A.8), as $-\frac{N}{\text{Tr } \mathbf{K}_1} > -\frac{1}{\tilde{\lambda}_N}$, the required condition of Lemma 4.2 on the explained variance remains as (A.8), i.e., $\eta_2 < -\frac{1}{\tilde{\lambda}_N}$.

Note that the increase of DKPCA in explained variance can now be written as $\frac{n-N}{\eta_2 \text{Tr } \mathbf{K}_1 + N}$. Given fixed η_2, n , and $\text{Tr } \mathbf{K}_1$, the explained variance boost of DKPCA is more pronounced when the decay or the cumulative ratio of the first n eigenvalues of \mathbf{K}_1 is not sharp, which is often the case in complex real-world data.

Appendix B. Supplementary empirical evaluations

B.1. Detailed experimental setups

Datasets. Three synthetic datasets are tested (see Fig. B.1): a square synthetic dataset (Synth 1, see Fig. B.1(a)), a complex 2D synthetic dataset consisting of one square, two spirals and one ring (Synth 2, see Fig. B.1(b)), and a multivariate Normal synthetic dataset (Synth 3), where samples are drawn randomly from a multivariate normal distribution with zero mean and fixed covariance matrix. The test set of Synth 3 consists of $N_{\text{test}} = 10000$ samples drawn from the same distribution. For real-world data, detailed descriptions on the used datasets can be found in Table B.1. The downstream supervised tasks are performed on publicly available UCI benchmark datasets (Dua & Graff, 2017).

Evaluation metrics and compared methods. Different related unsupervised learning methods are adopted to comprehensively evaluate our proposed deep KPCA. A comparison to the shallow kernel PCA is presented in terms of the explained variance, demonstrating the higher informativeness of the principal components learned by our method. For the general disentangled feature learning, we consider the state-of-the-art methods β -VAE (Higgins et al., 2017), FactorVAE (Kim & Mnih, 2018), and β -TCVAE (Chen et al., 2018). In the qualitative disentanglement experiments, convolutional-based network architectures are used for the data feature maps, with details shown in Table B.2. We keep the same encoder φ_1 and decoder ψ_1 architecture for all compared methods and use $k_2(z, y) = z^T y$ for the second DKPCA level. To quantitatively evaluate the disentanglement learning, we employ the

IRS metric (Suter et al., 2019), where a higher value indicates better robustness to changes in generative factors. In other words, if a latent variable is associated with some generative factor, the inferred value of this latent variable shows little change when that factor remains the same, regardless of interventions to the other generative factors. Other metrics for disentanglement evaluation have been proposed, but it has been shown that they are closely correlated with each other (Locatello et al., 2019).

Hyperparameter selection. In unsupervised learning experiments, for consistent evaluations, the shared hyperparameters among all methods are fixed to be the same, e.g., the RBF bandwidth in KPCA. For the model-specific hyperparameters, we used the suggested values in the papers of the compared methods. Specifically, we used $\beta = 4$ for β -VAE, $\beta = 6, \alpha = 1, \gamma = 1$ for β -TCVAE, and $\gamma = 10$ for FactorVAE. We fix $\eta_j = 1, j = 1, \dots, n_{\text{levels}}$ and $\gamma = 1$ in (11) to equally balance the AE and deep KPCA error. For the more challenging SmallNORB, we set $\gamma = 100$ for better reconstruction. In the qualitative disentanglement experiments, we use Adam optimizer (Kingma & Ba, 2015) with learning rate 10^{-4} and 50,000 maximum epochs for the compared methods. For DKPCA, in the qualitative disentanglement experiments we use Riemannian Adam (Becigneul & Ganea, 2019) with learning rate $2 \cdot 10^{-4}$ and 80,000 maximum epochs. Concerning the principal components, $s_1 = s_2$ is set to the true number of generative factors and the factors of variations involving the fewest pixels are trained on a subset with fixed factors of highest variation as these factors dominate the principal components as they have the largest number of pixels. In the quantitative disentanglement experiments, we employ the two-level architecture of (13) with linear kernels, which are parameter-free. In the experiments comparing the explained variance of KPCA, DKPCA and Deng et al. (2019), the shared hyperparameters, i.e., the RBF bandwidth, are set the same for all levels in the methods for consistent comparisons and selected based on reconstruction error on a validation set. In the supervised experiments, the RBF kernel is used for all datasets. Tuning is carried out through grid search based on validation performance. The σ^2 of RBF kernels is tuned between $\exp(-2)$ and $\exp(7)$. For DKPCA, we tune η_2 between -10 and 10 . The hidden features of the test points are obtained through a kernel smoother approach (Hastie et al., 2009) for the supervised and the large-scale disentanglement experiments. In the explained variance experiments, subsampling of 3DShapes and MNIST is performed with $N = 50$ and $N = 100$, respectively. The shared hyperparameters in the compared methods are tuned under the same settings, e.g., the kernel parameters are tuned in the same range.

B.2. Additional results

In Table B.3, we empirically evaluate the DKPCA optimization algorithms discussed in Section 3.2 in terms of speed and quality

Table B.1
 Details of the datasets used in the experimental evaluation of unsupervised learning with DKPCA, where “#” stands for “number of”.

Dataset	Input dimensions	# Factors of variation	Meaning of the factors of variation and # possible values	Total #data points
Synth 1	2	–	–	100
Synth 2	2	–	–	100
Synth 3	140	–	–	10 100
MNIST	1 x 28 x 28	–	–	60 000
Cars3D	3 x 64 x 64	3	- elevation (4 possible values) - azimuth (24 possible values) - object type (183 possible values)	17 568
3DShapes	3 x 64 x 64	6	- floor color (10 possible values) - wall color (10 possible values) - object color (10 possible values) - scale (8 possible values) - shape (4 possible values) - orientation (15 possible values)	480 000
SmallNORB	1 x 64 x 64	4	- category (5 possible values) - elevation (9 possible values) - azimuth (18 possible values) - lighting condition (6 possible values)	4860

Table B.2

Model architectures for the disentangled feature learning experiments with computer vision datasets. For all, $c = 40$ and $\hat{k} = 3$. All convolutions (*Conv*) and transposed convolutions (*ConvTr*) are with stride 2 and padding 1, except the last convolutional layer of φ_1 and the first transposed convolutional layer of ψ_1 , which have stride 1 and no padding. Layers have Parametric-RELU ($\alpha = 0.2$) activation functions, except the output layer of the pre-image map ψ_1 that has Sigmoid activation function (since input data is normalized in $[0, 1]$).

Encoder $\varphi_1(\cdot)$	Decoder $\psi_1(\cdot)$
$\left\{ \begin{array}{l} Conv [c] \times 4 \times 4; \\ Conv [c \times 2] \times 4 \times 4; \\ Conv [c \times 4] \times \hat{k} \times \hat{k}; \\ FC 256; \\ FC 50 (Linear) \end{array} \right.$	$\left\{ \begin{array}{l} FC 256; \\ FC [c \times 4] \times \hat{k} \times \hat{k}; \\ ConvTr [c \times 4] \times 4 \times 4; \\ ConvTr [c \times 2] \times 4 \times 4; \\ ConvTr [c] (Sigmoid) \end{array} \right.$

of the solutions on the Ionosphere dataset under the settings of Table 1 in main text. The termination condition for both algorithms is $\|H^{k+1} - H^k\|_{\max} / \alpha_k \leq 10^{-4}$ where α_k is the stepsize. Both algorithms achieve a similar final cost and all computed solutions show good feasibility. The lowest average cost is found by Riemannian Adam. Riemannian Adam is significantly faster than PGD due to the computationally expensive SVD used in the latter. Riemannian Adam can therefore be a more efficient alternative training algorithm for large-scale high-dimensional problems and was employed in the generative disentanglement experiments. We also show the training time of shallow KPCA and the deep KPCA of Deng et al. (2019) in Table B.4. It can be seen that deep kernel machines require more training time, which is reasonable as our DKPCA algorithm performs end-to-end training, employing a significantly different training scheme than the compared methods that simply conduct eigendecompositions. While

our DKPCA takes relatively longer time in training, our new insights in the unsupervised deep kernel methods and boosted performance are of significance.

Table B.5 gives the test reconstruction errors on a 140-dimensional synthetic dataset with different numbers of principal components in the two-level DKPCA of Eq. (13) with linear kernels. In Table B.5, for a fixed s_1 , the best test error is obtained with $s_2 = s_1$: the test error does not further decrease for $s_2 > s_1$. In fact, the rank of $H_1 H_1^T$ is at most s_1 , so an H_2 with rank higher than s_1 cannot lead to lower approximation error. Therefore, for a fixed s_1 in this conceived two-level architecture with linear kernels, s_1 should be set as $s_1 \leq s_2$ in terms of reconstruction error, in which increasing s_1 leads to lower reconstruction error as more principal components are incorporated.

We additionally compare the principal components from our DKPCA and Deng et al. (2019) through the ratio of explained variance. The results are shown in Fig. B.2. These experiments show that the learned features of the proposed DKPCA can explain more variance in the first principal components than the sequential deep KPCA of Deng et al. (2019).

In Table B.6, we provide the ablation study on the number of levels in our DKPCA, based on the setups of the downstream task in Table 1. The number of levels is currently determined by the validation performance: a two-level architecture gives good performance for all evaluated datasets and tasks in this work; performance does not distinctively vary when increasing the depth to three or four levels and yet it may result in overfitting and more computation. This is similar to determining the number of layers in neural networks. The ablation study demonstrates that 3-level and 4-level models perform similarly to the two-level model in terms of the supervised and unsupervised metrics. For simplified setups and more consistent evaluations, we default the number of levels as 2 in Table 1 in main text. A deeper architecture may be useful on problems where greater model complexity could boost classification accuracy. One can always tune each dataset separately on

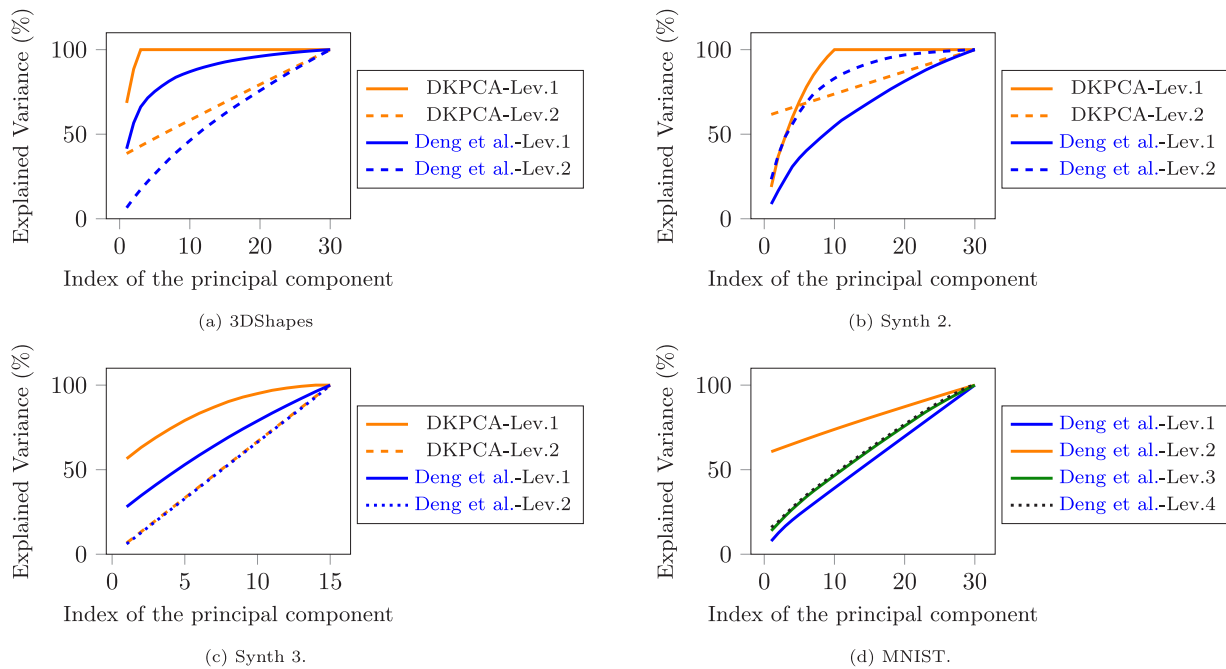


Fig. B.2. Explained variance of principal components from the deep KPCA of Deng et al. (2019) and our DKPCA. (a, b, c) Compared to Deng et al. (2019) using the same kernel, our method is able to capture greater explained variance in the first principal component, showing that the proposed forward–backward training algorithm outputs more informative principal components. (d) Four-level architecture from Deng et al. (2019) trained on MNIST, which does not consistently show increased explained variance with deeper levels compared to our DKPCA in Fig. 9(c).

Table B.3

Comparison of PGD and Riemannian Adam training algorithms for Problem (9) on the Ionosphere dataset. Mean (standard deviation) over five initializations. Lower is better on all metrics.

Algorithm	Final cost	Final feasibility	Runtime (s)
PGD	0.44045 (0.00003)	2.29×10^{-13} (10^{-12})	1.03 (0.018)
Riemannian Adam	0.43938 (0.0001)	1.93×10^{-12} (10^{-12})	0.172 (0.002)

Table B.4

Average training time of 5 runs (in seconds).

Dataset	KPCA	Deng et al. (2019)	DKPCA
Diabetes	0.016	0.038	0.161
Ionosphere	0.015	0.025	0.172
Liver	0.017	0.095	0.117
Cholesterol	0.014	0.037	0.126
Yacht Hydrodynamic	0.009	0.053	0.143

Table B.5

Test reconstruction error (MSE) on the 140D Synth 3 dataset for different numbers of principal components of the two levels in the proposed deep KPCA. All numbers are $\times 10^2$.

$s_2 \backslash s_1$	2	4	16	32	64	100
2	1.59	1.60	1.59	1.63	1.65	1.66
4	1.59	1.53	1.55	1.62	1.62	1.64
16	1.59	1.53	1.26	1.38	1.55	1.59
32	1.59	1.53	1.26	1.08	1.35	1.44
64	1.59	1.53	1.26	1.08	0.98	1.19
100	1.59	1.53	1.26	1.08	0.98	0.97

Table B.6

Ablation study on the number of levels in DKPCA in terms of the supervised downstream tasks and disentangled feature learning. Higher scores (\uparrow) are better for ACC (%) and WINDIN, lower scores (\downarrow) are better for RMSE.

Dataset	Metric	2-level	3-level	4-level
Diabetes	ACC(\uparrow)	72.02	72.07	72.10
	WINDIN(\uparrow)	0.287	0.275	0.276
Ionosphere	ACC(\uparrow)	93.04	93.01	92.99
	WINDIN(\uparrow)	1.05	1.08	1.13
Liver	ACC(\uparrow)	74.36	74.37	74.37
	WINDIN(\uparrow)	0.72	0.74	0.73
Cholesterol	RMSE(\downarrow)	60.59	60.25	60.19
	WINDIN(\uparrow)	0.003	0.004	0.004
Yacht Hydrodynamics	RMSE(\downarrow)	8.02	8.21	8.27
	WINDIN(\uparrow)	0.19	0.18	0.16

References

Allen-Zhu, Z., & Li, Y. (2023). Backward feature correction: How deep learning performs deep (hierarchical) learning. In G. Neu, & L. Rosasco (Eds.), *Proceedings of machine learning research: 195, Proceedings of thirty sixth conference on learning theory* (p. 4598). PMLR.

Becigneul, G., & Ganea, O.-E. (2019). Riemannian adaptive optimization methods. In *International conference on learning representations*.

Bengio, Y. (2009). Learning deep architectures for AI. *Foundations and trends in Machine Learning*, 2(1), 1–127.

its own validation set, and in this way the performance of our DKPCA can be further boosted as shown in the supplemented ablation study.

- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798–1828.
- Bohn, B., Rieger, C., & Griebel, M. (2019). A representer theorem for deep kernel learning. *Journal of Machine Learning Research*, 20(1), 2302–2333.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., & Abbeel, P. (2016). InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*. Curran Associates Inc.
- Chen, T. Q., Li, X., Grosse, R. B., & Duvenaud, D. K. (2018). Isolating sources of disentanglement in variational autoencoders. In *Advances in neural information processing systems*. Curran Associates Inc.
- Deng, X., Tian, X., Chen, S., & Harris, C. J. (2019). Deep principal component analysis based on layerwise feature extraction and its application to nonlinear process monitoring. *IEEE Transactions on Control Systems Technology*, 27(6), 2526–2540. <http://dx.doi.org/10.1109/TCST.2018.2865413>.
- Do, K., & Tran, T. (2022). Theory and evaluation metrics for learning disentangled representations. In *International conference on learning representations*. URL: <https://openreview.net/forum?id=HJgK0h4Ywr>.
- Dua, D., & Graff, C. (2017). *UCI machine learning repository*. University of California, Irvine, School of Information and Computer Sciences, URL: <http://archive.ics.uci.edu/ml>.
- Eckart, C., & Young, G. (1936). The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3), 211–218.
- Espinoza, M., Suykens, J. A., & De Moor, B. (2003). Least squares support vector machines and primal space estimation. In *IEEE international conference on decision and control*, vol. 4 (pp. 3451–3456). IEEE.
- Fanuel, M., Schreurs, J., & Suykens, J. (2021). Diversity sampling is an implicit regularization for kernel methods. *SIAM Journal on Mathematics of Data Science*, 3(1), 280–297.
- Fischer, A., & Igel, C. (2014). Training restricted Boltzmann machines: an introduction. *Pattern Recognition*, 47(1), 25–39.
- Girolami, M. (2002). Orthogonal series density estimation and the kernel eigenvalue problem. *Neural Computation*, 14(3), 669–688. <http://dx.doi.org/10.1162/089976602317250942>.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction*. Springer Science & Business Media.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., & Lerchner, A. (2017). Beta-VAE: learning basic visual concepts with a constrained variational framework. In *International conference on learning representations*.
- Hinton, G. E., Osindero, S., & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7), 1527–1554. <http://dx.doi.org/10.1162/neco.2006.18.7.1527>.
- Holzinger, A., Langs, G., Denk, H., Zatloukal, K., & Müller, H. (2019). Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(4), Article e1312.
- Jolliffe, I. T. (1986). *Principal component analysis*. Springer.
- Kim, H., & Mnih, A. (2018). Disentangling by factorising. In *International conference on machine learning*. PMLR.
- Kingma, D., & Ba, J. L. (2015). Adam: A method for stochastic optimization. In *International conference on learning representations*.
- Kingma, D., & Welling, M. (2014). Auto-encoding variational Bayes. In *International conference on learning representations*.
- Laforge, P., Cléménçon, S., & d'Alché-Buc, F. (2019). Autoencoding any data through kernel autoencoders. In *International conference on artificial intelligence and statistics*. PMLR.
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324.
- LeCun, Y., Cortes, C., & Burges, C. (2010). MNIST handwritten digit database. URL: <http://yann.lecun.com/exdb/mnist>.
- LeCun, Y., Huang, F. J., & Bottou, L. (2004). Learning methods for generic object recognition with invariance to pose and lighting. In *IEEE computer society conference on computer vision and pattern recognition*, vol. 2.
- Lever, J., Krzywinski, M., & Altman, N. (2017). Principal component analysis. *Nature Methods*, 14(7), 641–643.
- Locatello, F., Bauer, S., Lucic, M., Raetsch, G., Gelly, S., Schölkopf, B., & Bachem, O. (2019). Challenging common assumptions in the unsupervised learning of disentangled representations. In *International conference on machine learning*, vol. 97. PMLR.
- Mercer, J. (1909). Functions of positive and negative type, and their connection with the theory of integral equations. *Philosophical Transactions of the Royal Society, Series A*, 209(441–458), 415–446.
- Mika, S., Schölkopf, B., Smola, A., Müller, K.-R., Scholz, M., & Rätsch, G. (1999). Kernel PCA and de-noising in feature spaces. In *Advances in neural information processing systems*. MIT Press.
- Pandey, A., Schreurs, J., & Suykens, J. A. K. (2021). Generative restricted kernel machines: A framework for multi-view generation and disentangled feature learning. *Neural Networks*, 135, 177–191. <http://dx.doi.org/10.1016/j.neunet.2020.12.010>.
- Reed, S., Sohn, K., Zhang, Y., & Lee, H. (2014). Learning to disentangle factors of variation with manifold interaction. In *International conference on machine learning*. PMLR.
- Reed, S. E., Zhang, Y., Zhang, Y., & Lee, H. (2015). Deep visual analogy-making. In *Advances in neural information processing systems*, vol. 1. MIT Press.
- Ringnér, M. (2008). What is principal component analysis? *Nature biotechnology*, 26(3), 303–304.
- Rudi, A., Calandriello, D., Carratino, L., & Rosasco, L. (2018). On fast leverage score sampling and optimal learning. *Advances in Neural Information Processing Systems*.
- Salakhutdinov, R. (2015). Learning deep generative models. *Annual Review of Statistics and Its Application*, 2, 361–385.
- Salakhutdinov, R., & Hinton, G. (2009). Deep Boltzmann machines. In *International conference on artificial intelligence and statistics*.
- Sarhan, M. H., Eslami, A., Navab, N., & Albarqouni, S. (2019). Learning interpretable disentangled representations using adversarial VAEs. In *Domain adaptation and representation transfer and medical image learning with less labels and imperfect data* (pp. 37–44). Springer.
- Schölkopf, B., Smola, A., & Müller, K. R. (1998). Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10(5), 1299–1319. <http://dx.doi.org/10.1162/089976698300017467>.
- Suter, R., Miladinovic, D., Schölkopf, B., & Bauer, S. (2019). Robustly disentangled causal mechanisms: validating deep representations for interventional robustness. In *International conference on machine learning*, vol. 97. PMLR.
- Suykens, J. A. K. (2017). Deep restricted kernel machines using conjugate feature duality. *Neural Computation*, 29(8), 2123–2163. http://dx.doi.org/10.1162/neco_a_00984.
- Suykens, J. A. K., Van Gestel, T., De Brabanter, J., De Moor, B., & Vandewalle, J. (2002). *Least squares support vector machines*. World Scientific, <http://dx.doi.org/10.1142/5089>.
- Suykens, J. A. K., Van Gestel, T., Vandewalle, J., & De Moor, B. (2003). A support vector machine formulation to PCA analysis and its kernel version. *IEEE Transactions on Neural Networks*, 14(2), 447–450.
- Tonin, F., Patrinos, P., & Suykens, J. A. K. (2021). Unsupervised learning of disentangled representations in deep restricted kernel machines with orthogonality constraints. *Neural Networks*, 142, 661–679. <http://dx.doi.org/10.1016/j.neunet.2021.07.023>.
- Wright, J., Ganesh, A., Rao, S., Peng, Y., & Ma, Y. (2009). Robust principal component analysis: exact recovery of corrupted low-rank matrices via convex optimization. In *Advances in neural information processing systems*.