
A Dual Formulation for Probabilistic Principal Component Analysis

Henri De Plaen¹ Johan A. K. Suykens¹

Abstract

In this paper, we characterize *Probabilistic Principal Component Analysis* in Hilbert spaces and demonstrate how the optimal solution admits a representation in dual space. This allows us to develop a generative framework for kernel methods. Furthermore, we show how it englobes *Kernel Principal Component Analysis* and illustrate its working on a toy and a real dataset.

1. Introduction

Classical datasets often consist of many features, making dimensionality reduction methods particularly appealing. *Principal Component Analysis* (PCA) is one of the most straightforward frameworks to that goal and it is hard to find a domain in machine learning or statistics where it has not proven to be useful. PCA considers new decorrelated features by computing the eigendecomposition of the covariance matrix.

Probabilistic models on another side participate to the building of a stronger foundation for machine learning models. By considering models as probability distributions, we are able to natively access notions such as variance or sampling, *i.e.* generation. A probabilistic approach to PCA, known as *Probabilistic Principal Component Analysis* (Prob. PCA), has been formulated by (Tipping & Bishop, 1999). Its principles can be visualized in the primal part of Table 1.

Even endowed with a probabilistic interpretation, PCA remains restricted to linear relations between the different features. *Kernel Principal Component Analysis* (KPCA) (Mika et al., 1998; Schölkopf et al., 1998) was an attempt to give a non-linear extension to (non-probabilistic) PCA by decomposing a kernel matrix instead of the covariance matrix. An earlier attempt to give a probabilistic formulation of KPCA has been done by (Zhang et al., 2004). As developed further,

the latter model does not consist in a kernel equivalent of the Prob. PCA, but rather in another model based on similar principles.

More recently, *Restricted Kernel Machines* (Suykens, 2017) opened a new door for a probabilistic version of PCA both in primal and dual. They essentially use the Fenchel-Young inequality on a variational formulation of KPCA (Suykens et al., 2003; Alaíz et al., 2018) to obtain an energy function, closely resembling to *Restricted Boltzmann Machines*. The framework has been further extended to generation (Schreurs & Suykens, 2018; Winant et al., 2020), incorporating robustness (Pandey et al., 2020), multi-view models (Pandey et al., 2021), deep explicit feature maps (Pandey et al., 2022b) or times-series (Pandey et al., 2022a).

1.1. Contributions

1. We characterize the Prob. PCA framework in Hilbert spaces and give a dual interpretation to the model.
2. We develop a new extension of KPCA incorporating a noise assumption on the explicit feature map.
3. We give a probabilistic interpretation of the generation in KPCA.
4. We illustrate how the dual model works on a toy and a real dataset and show its connections to KPCA².

2. Primal and Dual Spaces

The key idea behind the duality in PCA is that outer and inner products share the same eigenvalues. The consequence is that instead of decomposing the covariance matrix of any given feature map, we can decompose the associated Gram matrix, *i.e.* the kernel matrix. The former is considered as the *primal* formulation and the latter as the *dual* formulation and they are both equivalent. Extending Prob. PCA to a dual formulation is however not straightforward: if all feature maps have an associated kernel, the converse is trickier. Some kernels correspond to feature maps in infinite dimensional spaces, where probability distributions cannot be properly defined. We therefore need to choose well defined finite subspaces to work in and consider linear operators instead of matrices. All formal definitions, propositions and proofs are provided in Appendix A.

¹KU Leuven, Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, Kasteelpark Arenberg 10, 3001 Leuven, Belgium. Correspondence to: Henri De Plaen <henri.deplaen@esat.kuleuven.be>.

²Resources: <https://hdeplaen.github.io/kppca>.




Figure 1. Global overview of the Probabilistic Principal Component Analysis in both primal and dual formulations. The primal spaces, or feature \mathcal{H} , $\mathcal{H}_\mathcal{E}$ and $\mathcal{H}_\mathcal{L}$ are in blue. The dual, or kernel and latent spaces \mathcal{E} and \mathcal{L} are in brown. The input space \mathcal{X} is in green. The color or the applications (arrows) is just for the readability and has nothing to do with the color of the spaces.

Distribution	Interpretation	Primal (features)	Dual (kernels)
latent observation	latent projection	$h \phi \sim \mathcal{N}(\Sigma_{h \phi}^{-1} \circ \mathbf{W}_{\text{ML}}^* (\phi - \phi_c), \sigma^2 \Sigma_{h \phi}^{-1})$	$h k_c \sim \mathcal{N}(\Sigma_{h k_c}^{-1} \circ \mathbf{A}_{\text{ML}} k_c, \Sigma_{h k_c}^{-1})$
observation latent	latent-based generation	$\phi h \sim \mathcal{N}(\mathbf{W}_{\text{ML}} h - \phi_c, \sigma^2 \mathbf{I}_{\mathcal{H}_\mathcal{E}})$	$k_c h \sim \mathcal{N}((\Phi_c^* \circ \Phi_c) \circ \mathbf{A}_{\text{ML}} h, \sigma^2 \Phi_c^* \circ \Phi_c)$
latent	latent prior	$h \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_\mathcal{L})$	$h \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_\mathcal{L})$
observation	absolute generation	$\phi \sim \mathcal{N}(\mu, \mathbf{W}_{\text{ML}} \circ \mathbf{W}_{\text{ML}}^* + \sigma^2 \mathbf{I}_{\mathcal{H}_\mathcal{E}})$	$k_c \sim \mathcal{N}(\mathbf{0}, \mathbf{A}_{\text{ML}}^* \circ \mathbf{A}_{\text{ML}} + \sigma^2 (\Phi_c^* \circ \Phi_c)^{-1})$

Table 1. Interpretation of the different distributions of the Prob. PCA framework after training, in both primal and dual formulations. The covariance operators are given by $\Sigma_{h|\phi} = (\mathbf{W}_{\text{ML}}^* \circ \mathbf{W}_{\text{ML}} + \sigma^2 \mathbf{I}_\mathcal{L})^{-1}$ and $\Sigma_{h|k_c} = (\mathbf{A}_{\text{ML}}^* \circ (\Phi_c^* \circ \Phi_c) \circ \mathbf{A}_{\text{ML}} + \sigma^2 \mathbf{I}_\mathcal{L})^{-1}$, with maximum likelihood estimators for the primal and dual interconnection operators \mathbf{W}_{ML} and \mathbf{A}_{ML} .

2.1. Primal Spaces

Feature Space \mathcal{H} . Given an input space \mathcal{X} , we first consider any feature map $\varphi : \mathcal{X} \rightarrow \mathcal{H}$. Following (Alaíz et al., 2018), we will consider a separable, possibly infinite dimensional, Hilbert space $(\mathcal{H}, \langle \cdot, \cdot \rangle_\mathcal{H})$. By φ , we denote an element of \mathcal{H} and its adjoint by $\varphi^* = \langle \varphi, \cdot \rangle \in \mathcal{H}^*$, with $\mathcal{H}^* \sim \mathcal{H}$ its Fréchet-Riesz dual space. Essentially, it corresponds to the transpose φ^\top in real, finite dimensional spaces as $\varphi_1^\top \varphi_2 = \langle \varphi_1, \varphi_2 \rangle_\mathcal{H}$, but generalizes it for the possibly infinite dimensional spaces that will be necessary for the introduction of kernels. Furthermore, we assume our space to be defined over the reals such that $\langle \cdot, \cdot \rangle_\mathcal{H} : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}$ and its inner product is symmetric $\langle \varphi_1, \varphi_2 \rangle_\mathcal{H} = \langle \varphi_2, \varphi_1 \rangle_\mathcal{H}$. If \mathcal{H} is of finite dimension d , we can therefore identify its canonical basis $\mathbf{u}_1, \dots, \mathbf{u}_d$ with the canonical basis of \mathbb{R}^d .

Finite Feature Space $\mathcal{H}_\mathcal{E}$. Considering a set of N observations $\{\mathbf{x}_i \in \mathcal{X}\}_{i=1}^N$, the idea is to work directly in \mathcal{H} by considering instead the feature map of the datapoints $\varphi_i = \varphi(\mathbf{x}_i)$. We can however not define a normal distribution onto the full \mathcal{H} yet as it is possibly infinite dimensional. We therefore have to consider a finite subspace $\mathcal{H}_\mathcal{E} \subset \mathcal{H}$. A natural choice would be $\mathcal{H}_\mathcal{E} = \text{span}\{\varphi_1, \dots, \varphi_N\}$. We now first have to find an orthonormal basis for $\mathcal{H}_\mathcal{E}$.

2.2. Dual Spaces

Kernels. For each feature map, there is an induced positive semi-definite kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R} : k(\mathbf{x}, \mathbf{y}) = \langle \varphi(\mathbf{x}), \varphi(\mathbf{y}) \rangle_\mathcal{H} = \varphi(\mathbf{x})^* \varphi(\mathbf{y})$. Inversely, to each positive semi-definite kernel corresponds a, possibly infinite dimensional, feature map, even if not explicitly defined. This follows from the theory of *Reproducing Kernel Hilbert Spaces*. We refer to (Schölkopf & Smola, 2001) for further info.

Kernel Space \mathcal{E} . We now consider a finite dimensional Hilbert space $(\mathcal{E}, \langle \cdot, \cdot \rangle_\mathcal{E})$ of dimension N , the number of observations. It is defined similarly as above, with orthonormal basis $\mathbf{e}_1, \dots, \mathbf{e}_N$. The basis also defines the identity over \mathcal{E} as $\mathbf{I}_\mathcal{E} = \sum_{i=1}^N \mathbf{e}_i \mathbf{e}_i^*$. The goal for \mathcal{E} is to represent the space of the kernel representations. We therefore define the linear operator $\Phi : \mathcal{E} \rightarrow \mathcal{H} : \sum_{i=1}^N \varphi_i \mathbf{e}_i^*$ and its adjoint $\Phi^* : \mathcal{H} \rightarrow \mathcal{E} : \sum_{i=1}^N \mathbf{e}_i \varphi_i^*$. Essentially, Φ^* returns the kernel value with each datapoint: $\Phi^* \varphi(\mathbf{x}) = \sum_{i=1}^N \mathbf{e}_i (\varphi_i^* \varphi(\mathbf{x})) = \sum_{i=1}^N \mathbf{e}_i k(\mathbf{x}_i, \mathbf{x})$ for any $\mathbf{x} \in \mathcal{X}$. Similarly, Φ projects this value back as a linear combination of the different φ_i 's, thus mapping back to $\mathcal{H}_\mathcal{E} \subset \mathcal{H}$. For this reason, the covariance $\Phi \circ \Phi^* = \sum_{i=1}^N \varphi_i \varphi_i^*$ acts as a projector from $\mathcal{H} \rightarrow \mathcal{H}_\mathcal{E}$. Its eigenvectors therefore form an orthonormal basis of the finite feature space $\mathcal{H}_\mathcal{E}$,

which acts as the primal equivalent of the kernel space \mathcal{E} .

Centered Kernels. In most applications however, we prefer to work with the centered feature map, which we define as $\varphi_c(\cdot) = \varphi(\cdot) - \varphi_c$ with $\varphi_c = \frac{1}{N} \sum_{i=1}^N \varphi_i$. We denote the associated kernel associated centered kernel $k_c : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R} : k_c(\mathbf{x}_1, \mathbf{x}_2) = \varphi_c(\mathbf{x}_1)^* \varphi_c(\mathbf{x}_2)$. This leads to the definition of a new centered operator $\Phi_c = \sum_{i=1}^N (\varphi_i - \varphi_c) e_i^* = \Phi (\mathbf{I}_{\mathcal{E}} - \frac{1}{N} \mathbf{1}_{\mathcal{E} \times \mathcal{E}})$, with $\mathbf{1}_{\mathcal{E} \times \mathcal{E}} = \sum_{i,j=1}^N e_i e_j^*$. As always, we also consider its adjoint Φ_c^* . Considering the dual operator, we have $\Phi_c^* \circ \Phi_c = \sum_{i=1}^N (\varphi_i - \varphi_c)^* (\varphi_i - \varphi_c) e_i e_i^* = \sum_{i=1}^N k_c(\mathbf{x}_i, \mathbf{x}_j) e_i e_i^*$. We notice now that $\mathcal{H}_{\mathcal{E}} = \text{span}\{\varphi_1, \dots, \varphi_N\} = \text{span}\{\varphi_1 - \varphi_c, \dots, \varphi_N - \varphi_c\}$ because φ_c is a linear combination of the elements of the basis. Therefore, the primal operator $\Phi_c \circ \Phi_c^* = \sum_{i=1}^N (\varphi_i - \varphi_c) (\varphi_i - \varphi_c)^*$ also acts as a projector from $\mathcal{H} \rightarrow \mathcal{H}_{\mathcal{E}}$ and we can choose its eigenvectors instead as an orthonormal basis of $\mathcal{H}_{\mathcal{E}}$.

Covariance and Kernels. We now consider the key idea behind the duality in PCA: the operators $\Phi_c \circ \Phi_c^*$ and $\Phi_c^* \circ \Phi_c$ are self-adjoint, positive semi-definite and share the same non-zero eigenvalues. We have $\Phi_c \circ \Phi_c^* = \sum_{i=1}^N \lambda_i v_i v_i^*$ and $\mathcal{H}_{\mathcal{E}} = \text{span}\{v_1, \dots, v_N\}$. Similarly, we have $\Phi_c^* \circ \Phi_c = \sum_{i=1}^N \lambda_i \epsilon_i \epsilon_i^*$ and $\mathcal{E} = \text{span}\{\epsilon_1, \dots, \epsilon_N\}$. The identity over the (primal) finite feature space $\mathcal{H}_{\mathcal{E}}$ can now be defined as $\mathbf{I}_{\mathcal{H}_{\mathcal{E}}} = \sum_{i=1}^N v_i v_i^*$ and the identity over the (dual) kernel space \mathcal{E} as $\mathbf{I}_{\mathcal{E}} = \sum_{i=1}^N \epsilon_i \epsilon_i^*$. This is synthesized in the two first columns of Table 2. The identity over \mathcal{H} reads $\mathbf{I}_{\mathcal{H}} = \mathbf{I}_{\mathcal{H}_{\mathcal{E}}} + \mathbb{P}_{\mathcal{H}_{\mathcal{E}}^\perp}$, with $\mathbb{P}_{\mathcal{H}_{\mathcal{E}}^\perp}$ a projector over the null space of $\Phi_c \circ \Phi_c^*$. It must be noted that it may happen that these basis may contain too much basis vectors if the two operators $\Phi_c^* \circ \Phi_c$ and $\Phi_c \circ \Phi_c^*$ are not of full rank. In particular, this is the case when $\dim(\mathcal{H}) = d$ is finite and $d < N$. In this particular case, we would also have $\dim(\mathcal{H}_{\mathcal{E}}) = \dim(\mathcal{E}) = d$. Without loss of generality, we will assume that this is not the case. Similarly, we will neglect the case $N > d$ as we could just neglect the null space of $\Phi_c^* \circ \Phi_c$.

Notations. We can now define our probabilistic model over $\mathcal{H}_{\mathcal{E}}$. We will therefore use the notation ϕ instead of φ to consider the feature map in our finite dimensional subspace $\mathcal{H}_{\mathcal{E}}$. More formally, we have $\phi : \mathcal{X} \rightarrow \mathcal{H}_{\mathcal{E}} : \mathbf{I}_{\mathcal{H}_{\mathcal{E}}} \circ \varphi$ and following from that $\phi_c : \mathcal{X} \rightarrow \mathcal{H}_{\mathcal{E}} : \mathbf{I}_{\mathcal{H}_{\mathcal{E}}} \circ \varphi_c$. In particular, we have the observations $\phi_i = \phi(\mathbf{x}_i) = \varphi_i$ and $\phi_c = \varphi_c$, as those are linear combinations of the basis. For the sake of readability, we will write $\phi = \phi(\mathbf{x})$, the image of a random variable $\mathbf{x} \in \mathcal{X}$ and refer to it as a *feature* observation or representation. Given any Hilbert space, \mathbf{a} an element of it and a linear operator Σ from and to that space, we consider the *multivariate normal distribution* $\mathbf{a} \sim \mathcal{N}(\mathbf{b}, \Sigma)$ as the distribution with density $\frac{1}{Z} \exp(-\frac{1}{2}(\mathbf{a} - \mathbf{b})^* \Sigma^{-1}(\mathbf{a} - \mathbf{b}))$. It is well defined if Z is non-zero and finite.

		Dimension	d	N	q
Primal	Space		$\mathcal{H}(= \mathbb{R}^d)$	$\mathcal{H}_{\mathcal{E}} \subset \mathcal{H}$	$\mathcal{H}_{\mathcal{L}} \subset \mathcal{H}_{\mathcal{E}}$
	Canon. Basis		$\{\mathbf{u}_i\}_{i=1}^d$	N.A.	N.A.
	Other Basis		N.A.	$\{\mathbf{v}_i\}_{i=1}^N$	$\{\mathbf{e}_p\}_{p=1}^q$
Dual	Space		N.A.	$\mathcal{E} = \mathbb{R}^N$	$\mathcal{L} \subset \mathcal{E}$
	Canon. Basis		N.A.	$\{\mathbf{e}_i\}_{i=1}^N$	N.A.
	Other Basis		N.A.	$\{\mathbf{\epsilon}_i\}_{i=1}^N$	$\{\mathbf{r}_p\}_{p=1}^q$

Table 2. Different spaces with their dimension and the canonical orthonormal basis if it applies, as well as another (non-canonical) basis when used throughout this paper. The equality $\mathcal{H} = \mathbb{R}^d$ only makes sense if $\dim(\mathcal{H}) = d$ is finite.

3. Primal Model

We will now essentially follow the work of (Tipping & Bishop, 1999) and redefine the model distributions. This section corresponds to the primal formulation and we only consider the feature representations. It does not yet introduce the kernel representations, which will appear in the dual formulation (Section 4).

3.1. Model and Latent Space

Factor Analysis. The starting point is to consider a *factor analysis* relationship (Bartholomew et al., 2011; Basilevsky, 2009) between the feature observations ϕ and the latent variables \mathbf{h} . In particular, we consider

$$\phi = \mathbf{W}\mathbf{h} + \boldsymbol{\mu} + \boldsymbol{\zeta}. \quad (1)$$

The observations ϕ live in the primal space $\mathcal{H}_{\mathcal{E}}$ of dimension N . We consider an isotropic normal noise $\boldsymbol{\zeta} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_{\mathcal{H}_{\mathcal{E}}})$ of variance $\sigma^2 \in \mathbb{R}_{>0}$ and a mean $\boldsymbol{\mu} \in \mathcal{H}_{\mathcal{E}}$.

Latent Space \mathcal{L} . The latent variables \mathbf{h} on the other hand live in a latent dual space $\mathcal{L} \subset \mathcal{E}$ of dimension $q \leq N$. They are related by a primal *interconnection linear operator* \mathbf{W} . As it was the case before with Φ , the interconnection operator does not project to the full space $\mathcal{H}_{\mathcal{E}}$ because of its reduced dimensionality. It therefore projects to yet another feature space $\mathcal{H}_{\mathcal{L}} \subset \mathcal{H}_{\mathcal{E}}$, which acts as the primal equivalent of the latent space \mathcal{L} . The equality of these two spaces only holds if $q = N$. We will therefore consider the mappings $\mathbf{W}^* : \mathcal{H}_{\mathcal{E}} \rightarrow \mathcal{L}$ and $\mathbf{W} : \mathcal{L} \rightarrow \mathcal{H}_{\mathcal{L}}$. The identity over \mathcal{L} can be written as $\mathbf{I}_{\mathcal{L}} = \sum_{p=1}^q \mathbf{r}_p \mathbf{r}_p^*$, over $\mathcal{H}_{\mathcal{L}}$ as $\mathbf{I}_{\mathcal{H}_{\mathcal{L}}} = \sum_{p=1}^q \mathbf{e}_p \mathbf{e}_p^*$ and finally the identity over $\mathcal{H}_{\mathcal{E}}$ rewritten as $\mathbf{I}_{\mathcal{H}_{\mathcal{E}}} = \mathbf{I}_{\mathcal{H}_{\mathcal{L}}} + \mathbb{P}_{\mathcal{H}_{\mathcal{L}}^\perp}$, with $\mathbb{P}_{\mathcal{H}_{\mathcal{L}}^\perp}$ as a projector over the null space of $\mathbf{W}^* \circ \mathbf{W}$. This is summarized in the last column of Table 2.

3.2. Feature Distributions

Latent-Based Generation. The relation between the feature observations and the latent variables being set up

(Eq. (1)), we can derive the conditional probability of the feature observations given a latent variable:

$$\phi|\mathbf{h} \sim \mathcal{N}(\mathbf{W}\mathbf{h} - \boldsymbol{\mu}, \sigma^2 \mathbf{I}_{\mathcal{H}_\mathcal{E}}). \quad (2)$$

As discussed earlier, we see that the latent variables do not participate to the full scope of the observations in $\mathcal{H}_\mathcal{E}$, but only to their component in $\mathcal{H}_\mathcal{L}$. The rest is only constituted from the isotropic normal noisy mean. This distribution can be interpreted as a generative one: given a latent variable, we can sample a variable in feature space.

Absolute Generation. Considering the latent prior $\mathbf{h} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_\mathcal{L})$, we can derive the marginal distribution of the observations in feature space:

$$\phi \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{W} \circ \mathbf{W}^* + \sigma^2 \mathbf{I}_{\mathcal{H}_\mathcal{E}}). \quad (3)$$

It can be considered as the data distribution of the model. Sampling from it also means generating feature representations in a more absolute way, *i.e.*, without considering any latent variable, or more precisely considering a random latent variable according to its prior. As a consequence of Eq. (2) and the isotropic aspect of the latent prior, we see that the observations are only non-isotropically distributed in $\mathcal{H}_\mathcal{L}$. Again, the rest is only the isotropically normally noisy mean. In other words, this means that the model parameter \mathbf{W} only influences ϕ for its components in $\mathcal{H}_\mathcal{L}$.

3.3. Training the Model

Maximum Likelihood. As we now have the marginal distribution of the model (Eq. (3)), the goal is to find the optimal hyperparameters \mathbf{W} and $\boldsymbol{\mu}$ to match the set of observations $\{\phi_i\}_{i=1}^N$. One way to determine them is by maximizing the likelihood of our observations. The *Maximum Likelihood* (ML) estimator for the hyperparameters is given by:

$$\boldsymbol{\mu}_{\text{ML}} = \phi_c, \quad (4)$$

$$\mathbf{W}_{\text{ML}} = \sum_{p=1}^q \sqrt{\lambda_p/N - \sigma^2} \mathbf{v}_p \mathbf{r}_p^*, \quad (5)$$

with $\{(\lambda_p, \mathbf{v}_p)\}_{p=1}^q$ the q dominant eigenpairs of $\Phi_c \circ \Phi_c^*$ ($\lambda_1 \geq \dots \geq \lambda_q \geq \dots \geq \lambda_N$), and $\{\mathbf{r}_p\}_{p=1}^q$ and arbitrary orthonormal basis of the latent space \mathcal{L} . The choice for the latter basis is arbitrary and makes the model rotational invariant in latent space. An additional condition is that $\sigma^2 \leq \lambda_q/N$. It is not surprising to see that the optimal mean $\boldsymbol{\mu}_{\text{ML}}$ corresponds to the mean of the observations ϕ_c . We observe that \mathbf{W}_{ML} corresponds to the eigendecomposition of the centered covariance, at the exception that the noise assumption is subtracted from its spectrum. By looking back at Eq. (1), it makes sense to avoid the noise in \mathbf{W}_{ML} as it is still going to be added by the term ζ .

Noise Variance. Maximizing the likelihood as a function of σ^2 leads to

$$\sigma_{\text{ML}}^2 = \frac{1}{N(N-q)} \sum_{p=q+1}^N \lambda_p. \quad (6)$$

The eigenvalue λ_p corresponds to the variance for each component \mathbf{v}_p of the covariance $\Phi_c \circ \Phi_c^*$. The total variance of the data, noise included, is equal to $\frac{1}{N} \sum_{p=1}^N \lambda_p$ and the variance learned by the model through the primal interconnection operator to $\frac{1}{N} \sum_{p=1}^q \lambda_p$. Hence, the maximum likelihood estimator for the noise variance σ_{ML}^2 can be interpreted as the mean of the variance that is discarded by the model. It also verifies the earlier condition that $\sigma^2 \leq \lambda_q/N$, as the eigenvalues are taken in descending order. It can be interpreted as the normalized mean variance of the left over eigendirections, *i.e.* the orthogonal space of the latent space: $\mathcal{L}^\perp = \mathcal{E} \setminus \mathcal{L}$. By consequence, we may decide to choose the latent dimension $q = \dim(\mathcal{L})$ and deduct σ_{ML}^2 . In the opposite, we may also decide to set an arbitrary σ^2 and deduct the latent dimension q instead. We therefore can consider either σ^2 or q as an additional hyperparameter. We must however keep in mind that this is strongly going to be influenced by the distribution of the eigenvalues and that the latent dimension q for the same σ^2 may heavily vary from application to application.

Uncentered Features. We may also consider not to consider the mean as an optimizable hyperparameter and set it arbitrarily to $\boldsymbol{\mu} = \mathbf{0}$. In this case, Eq. (5) would be the same at the difference that the \mathbf{W}_{ML} would be constructed from the dominant eigenpairs of the uncentered covariance $\Phi \circ \Phi^*$ instead of its centered counterpart $\Phi_c \circ \Phi_c^*$.

3.4. Dimensionality Reduction in Feature Space

Latent Projection. Up to now, we only considered the distribution of the feature variables ϕ . We can also calculate the posterior distribution of the latent variable \mathbf{h} given the primal feature variable ϕ :

$$\mathbf{h}|\phi \sim \mathcal{N}(\boldsymbol{\Sigma}_{\mathbf{h}|\phi}^{-1} \circ \mathbf{W}^*(\phi - \boldsymbol{\mu}), \sigma^2 \boldsymbol{\Sigma}_{\mathbf{h}|\phi}^{-1}), \quad (7)$$

with $\boldsymbol{\Sigma}_{\mathbf{h}|\phi} = (\mathbf{W}^* \circ \mathbf{W} + \sigma^2 \mathbf{I}_\mathcal{L})^{-1}$. The mean of the distribution can be considered as a pseudo-inverse of the observation ϕ , but regularized by σ^2 . This regularization ensures to avoid the noise. If the prior of the latent variables was isotropic, this is not the case anymore for the posterior. If we consider the maximum likelihood estimator for the primal interconnection operator \mathbf{W}_{ML} , the variance becomes $\sigma^2 \boldsymbol{\Sigma}_{\mathbf{h}|\phi}^{-1} = N \sigma^2 \sum_{p=1}^q \lambda_p^{-1} \mathbf{r}_p \mathbf{r}_p^*$. It can be interpreted as the uncertainty for each component of the latent variable \mathbf{h} (w.r.t. the eigendirection \mathbf{r}_p), due to the noise assumption. By consequence, the greater the explained variance λ_p for the eigendirection \mathbf{v}_p of the covariance $\Phi_c \circ \Phi_c^*$, the smaller

the corresponding uncertainty on the component \mathbf{r}_p of the latent variable \mathbf{h} . For each observation in feature space ϕ , this returns a distribution for the latent variable ϕ and can therefore be considered as a sort of probabilistic projection in latent space \mathcal{L} .

Maximum A Posteriori. Up to now, we were only considering distributions. The only way to go from a feature representation to a latent variable or the opposite was probabilistic. In order to have a deterministic approach, we need proper mappings. One way is to consider the *Maximum A Posteriori* (MAP) of \mathbf{h} given ϕ . It maps the feature observation $\phi \in \mathcal{H}_{\mathcal{E}}$ to latent variable $\mathbf{h}_{\text{MAP}} \in \mathcal{L}$, hence reducing the dimensionality of any input to that of the latent space. To allow it to work for any input $\varphi \in \mathcal{H}$, we may again consider the projection $\phi = \mathbf{I}_{\mathcal{H}_{\mathcal{E}}} \varphi$. As $\mathbf{W}_{\text{ML}}^* \circ \mathbf{I}_{\mathcal{H}_{\mathcal{E}}} = \mathbf{W}_{\text{ML}}^*$:

$$\mathbf{h}_{\text{MAP}} = (\mathbf{W}_{\text{ML}}^* \circ \mathbf{W}_{\text{ML}} + \sigma^2 \mathbf{I}_{\mathcal{L}})^{-1} \circ \mathbf{W}_{\text{ML}}^* (\varphi - \varphi_c). \quad (8)$$

To map back to the feature space $\mathcal{H}_{\mathcal{L}}$, we may consider the *maximum a posteriori* of ϕ given \mathbf{h} (Eq. (3)). This gives

$$\phi_{\text{MAP}} = \mathbf{W}_{\text{MAP}} \mathbf{h} + \phi_c. \quad (9)$$

The final projection reads

$$\phi_{\text{MAP}} = \mathbf{W}_{\text{ML}} \circ (\mathbf{W}_{\text{ML}}^* \circ \mathbf{W}_{\text{ML}} + \sigma^2 \mathbf{I}_{\mathcal{L}})^{-1} \circ \mathbf{W}_{\text{ML}}^* (\varphi - \varphi_c) + \phi_c. \quad (10)$$

No Noise. We may also decide not to consider σ^2 as a parameter to optimize and set it to an arbitrary value. The latent dimensions q could also be set an arbitrary value, without it to be related to the latent dimension q according to Eq. (6). We notice that in the limit of $\sigma^2 \rightarrow 0$, we recover the classical Principal Component Analysis reconstruction scheme. Indeed the conditional probability distributions become exact relations. We also notice that the condition $\sigma^2 \leq \lambda_q/N$ (Prop. 3) is then always satisfied. Furthermore, when $q = \dim(\mathcal{H}_{\mathcal{E}})$, the reconstruction is perfect in $\mathcal{H}_{\mathcal{E}}$ and in particular for our original observations $\{\varphi_i\}_{i=1}^N$ and φ_c (as we have $\phi_i = \varphi_i$). Indeed, we would have

$$\mathbf{h}_{\text{MAP}} = \mathbf{W}_{\text{ML}}^+ (\varphi - \varphi_c), \quad (11)$$

with \mathbf{W}_{ML}^+ the Moore-Penrose pseudo-inverse of \mathbf{W}_{ML} . We note here the symmetry with Eq. (9). If the maximum likelihood estimator for σ^2 is to be respected (Eq. (6)), this would mean that all components are kept ($\mathcal{L} = \mathcal{E}$) and the model reconstructs the full feature variance. In this case, the primal interconnection operator would become $\mathbf{W}_{\text{ML}} = \sum_{p=1}^N \sqrt{\lambda_p/N} \mathbf{v}_p \mathbf{r}_p^*$ and be invertible. Its Moore-Penrose pseudo-inverse would become an exact inverse. Eqs. (9) and (11) would become exact opposites and there would be no loss due to the dimensionality reduction as there would be no noise to discard. By consequence, the reduction would become an identity over $\mathcal{H}_{\mathcal{E}}$: $\phi_{\text{MAP}} - \phi_c = \mathbf{I}_{\mathcal{H}_{\mathcal{L}}} (\varphi - \varphi_c)$.

4. Dual Model

Kernels without Dual. In (Zhang et al., 2004), the authors made the kernel matrix appear by considering the new observations $\{\sum_{i=1}^d \mathbf{u}_i \mathbf{u}_j^* \phi(\mathbf{x}_i)\}_{j=1}^N$. In other words, each new datapoint consists in one particular feature of the feature map, for each original datapoint. If the original datapoints were organized as a matrix in $\mathbb{R}^{N \times d}$, this would correspond to taking its transpose as new datapoints. The outer product of the covariance matrix is transformed to the inner product of the kernel matrix. If indeed this formulation makes the kernel appear, it is not a dual formulation of the original problem, but another problem. In this section, we show how the spaces defined hereabove help us build an equivalent dual formulation of the problem.

Dual Formulation. While keeping an equivalence with the primal model, we will now see that we can directly work in dual spaces \mathcal{E} and \mathcal{L} without considering the feature spaces at all, *i.e.* resorting to the primal space \mathcal{H} and its subsets. As we did for the primal feature variable ϕ , we will consider $\mathbf{k}_c = \Phi_c^* (\phi - \phi_c) = \sum_{i=1}^N k_c(\mathbf{x}, \mathbf{x}_i) \mathbf{e}_i$ to represent the image in \mathcal{E} , of a random variable $\mathbf{x} \in \mathcal{X}$. We will refer to it as a *dual feature variable*.

4.1. Representation

Considering the dual spaces, we can always express the interconnection operator \mathbf{W} in the (non-orthonormal) basis $\{\phi_1 - \phi_c, \dots, \phi_N - \phi_c\}$. As a consequence, we can always write

$$\mathbf{W} = \Phi_c \circ \mathbf{A}, \quad (12)$$

with $\mathbf{A} : \mathcal{L} \rightarrow \mathcal{L}$, the dual interconnection operator. Given the maximum likelihood estimator for the primal interconnection operator \mathbf{W}_{ML} , we can directly deduce the dual one:

$$\mathbf{A}_{\text{ML}} = \sum_{p=1}^q \sqrt{1/N - \sigma^2 \lambda_p^{-1}} \epsilon_p \mathbf{r}_p^*, \quad (13)$$

with $\{(\lambda_p, \epsilon_p)\}_{p=1}^q$ the q dominant eigenpairs of $\Phi_c^* \circ \Phi_c$ and $\{\mathbf{r}_p\}_{p=1}^q$ an arbitrary orthonormal basis of the latent space \mathcal{L} . The rotational invariance of the dual interconnection operator \mathbf{A}_{ML} is inherited from its primal counterpart \mathbf{W}_{ML} . Again, if we consider an optimized mean $\boldsymbol{\mu} = \mathbf{0}$, we would have the relation $\mathbf{W}_{\text{ML}} = \Phi \circ \mathbf{A}_{\text{ML}}$ with \mathbf{A}_{ML} then based on the eigenpairs of the non-centered $\Phi^* \circ \Phi$ instead. Using the same structure for \mathbf{A}_{ML} , the optimal (primal) interconnection operator \mathbf{W}_{ML} could be expressed in the (non-orthonormal) basis $\{\phi_1, \dots, \phi_N\}$.

4.2. Kernel Distributions

Projection and Generation. We can also consider the dual counterparts of the distributions of the primal model (Eqs. (2) and (7)). For the sake of simplicity and to avoid

heavier equations with non-centered kernels, we will only consider here the equations of the trained model, in particular with $\mu_{\text{ML}} = \phi_c$ leading to centered kernels:

$$\mathbf{k}_c | \mathbf{h} \sim \mathcal{N}((\Phi_c^* \circ \Phi_c) \circ \mathbf{A}_{\text{ML}} \mathbf{h}, \sigma^2 \Phi_c^* \circ \Phi_c), \quad (14)$$

$$\mathbf{h} | \mathbf{k}_c \sim \mathcal{N}(\Sigma_{\mathbf{h}|\mathbf{k}_c}^{-1} \circ \mathbf{A}_{\text{ML}} \mathbf{k}_c, \Sigma_{\mathbf{h}|\mathbf{k}_c}^{-1}), \quad (15)$$

with $\Sigma_{\mathbf{h}|\mathbf{k}_c} = (\mathbf{A}_{\text{ML}}^* \circ (\Phi_c^* \circ \Phi_c) \circ \mathbf{A}_{\text{ML}} + \sigma^2 \mathbf{I}_{\mathcal{L}})^{-1}$.

4.3. Dimensionality Reduction in Kernel Space

Maximum A Posteriori. This now allows us to consider the dimensionality reduction in kernel space in a similar fashion as in Section 3.4. Again we consider the MAP of the latent variable \mathbf{h} given the kernel representation \mathbf{k}_c :

$$\mathbf{h}_{\text{MAP}} = (\mathbf{A}_{\text{ML}}^* \circ (\Phi_c^* \circ \Phi_c) \circ \mathbf{A}_{\text{ML}} + \sigma^2 \mathbf{I}_{\mathcal{L}})^{-1} \circ \mathbf{A}_{\text{ML}} \mathbf{k}_c, \quad (16)$$

and similarly with the MAP of the kernel representation \mathbf{k}_c given the latent variable \mathbf{h} :

$$(\mathbf{k}_c)_{\text{MAP}} = (\Phi_c^* \circ \Phi_c) \circ \mathbf{A}_{\text{ML}} \mathbf{h}. \quad (17)$$

As for the primal model, the dimensionality reduction in dual is computed as $(\mathbf{k}_c)_{\text{MAP}} = (\Phi_c^* \circ \Phi_c) \circ \mathbf{A}_{\text{ML}} \mathbf{h}_{\text{MAP}}$.

No Noise. Again, considering $\sigma^2 \rightarrow 0$ makes both dual conditional distributions become exact relations. In a ML context for σ^2 (Eq. (6)), this would imply that $q = \dim(\mathcal{E})$ and we would recover an identity $(\mathbf{k}_c)_{\text{MAP}} = \mathbf{k}_c$, *i.e.* no reduction. Without considering a ML context for $\sigma^2 \rightarrow 0$ and choosing an arbitrary $q \leq \dim(\mathcal{E})$, the reduction become exactly the reconstruction done in KPCA.

4.4. Kernel Sampling

Probabilistic Sampling. The dual counterpart of Eq. (3) after training is given by

$$\mathbf{k}_c \sim \mathcal{N}(\mathbf{0}, \mathbf{A}_{\text{ML}}^* \circ \mathbf{A}_{\text{ML}} + \sigma^2 (\Phi_c^* \circ \Phi_c)^{-1}). \quad (18)$$

The covariance $\mathbf{A}_{\text{ML}}^* \circ \mathbf{A}_{\text{ML}} + \sigma^2 (\Phi_c^* \circ \Phi_c)^{-1}$ can be decomposed as $\mathbf{B} \circ \mathbf{B}^*$, with $\mathbf{B} : \mathcal{E} \rightarrow \mathcal{E} : N^{-1/2} \sum_{p=1}^q \lambda_p \epsilon_p \epsilon_p^* + \sum_{p=q+1}^N \sigma \lambda_p^{1/2} \epsilon_p \epsilon_p^*$ and $\{\epsilon_i\}_{i=1}^N$ any arbitrary orthonormal basis of the latent space \mathcal{E} . This decomposition allows us to sample \mathbf{k}_c on the trained model with $\mathbf{k}_c = \mathbf{B} \boldsymbol{\xi}$ with $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{\mathcal{E}})$. We see that \mathbf{B} is rotational invariant, which is not surprising as this is also the case for the distribution from which $\boldsymbol{\xi}$ is sampled. In practice and for simplicity, we may decide too choose the canonical basis for $\{\epsilon_i\}_{i=1}^N$ as any choice would be identified to the same covariance and to the same sampling of \mathbf{k}_c . We will therefore assume that $\epsilon_i = e_i$ for all $i = 1, \dots, N$. In that particular case, \mathbf{B} is self-adjoint

and by consequence corresponds to the matrix square root of $\mathbf{A}_{\text{ML}}^* \circ \mathbf{A}_{\text{ML}} + \sigma^2 (\Phi_c^* \circ \Phi_c)^{-1}$.

KPCA Sampling The classical sampling done by KPCA (Schreurs & Suykens, 2018) corresponds to the limit of $\sigma^2 \rightarrow 0$ for an arbitrary latent dimension q . Unless the latent dimension is chosen as $q = \dim(\mathcal{E})$, the sampling in that case can never cover \mathcal{E} fully, but rather \mathcal{L} , as \mathbf{B} is not a bijection. The second term of \mathbf{B} ($\sum_{p=q+1}^N \sigma \lambda_p^{1/2} \epsilon_p \epsilon_p^*$) allows \mathbf{B} to be a bijection no matter what is the choice of the latent dimension q , as long as $\sigma^2 > 0$. We thus always sample in the full \mathcal{E} . This can be observed at Fig. 2.




Figure 2. Schematic overview of the dual sampling in Prob. PCA compared to the generation in KPCA.

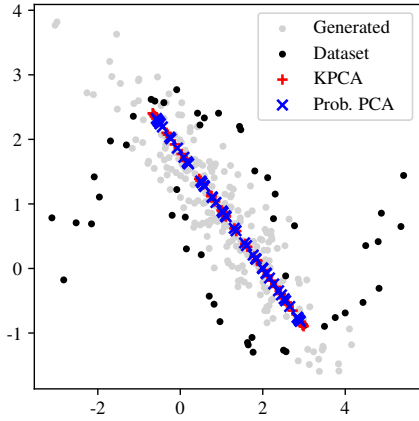
5. Experiments

Hilbert Spaces to Matrices. Working in Hilbert spaces is helpful to treat possibly infinite dimensional feature maps, but not very useful for practical applications. Matrix representations are possible in primal if d is finite and in dual if N is finite. It suffices to consider the different canonical basis. For the latent space \mathcal{L} , this enforces a unique representation for \mathbf{W}_{ML} and \mathbf{A}_{ML} , but we must keep in mind that they are rotational invariant. All the operators and elements described before are then represented in matrix or vector format (Table 3). We will use the tilde to denote these matrices and use software-like notation by denoting with $(\cdot)_{i_1:i_2, j_1:j_2}$ the matrix truncated to its i_1 to i_2 rows and j_1 to j_2 columns.

Preimage. Given a dual representation, we will also consider the *kernel smoother* preimage method, as suggested by (Schreurs & Suykens, 2018):

$$\hat{\mathbf{x}} = \frac{\sum_{i=1}^N (\tilde{\mathbf{k}})_i \mathbf{x}_i}{\sum_{i=1}^N (\tilde{\mathbf{k}})_i}. \quad (19)$$


In practice, as we work with centered feature maps and kernels, it may be that the kernel smoother may be unstable due to its normalization term. We therefore may consider to add a stabilization term.



(a) With $q = 1$ component, the explained variance is 31.23% and $\sigma^2 = 1.40\%$.




(a) Sample of original datapoints.




(b) With $q = 3$ components, the explained variance is 54.03% and $\sigma^2 = 0.98\%$.



(b) Datapoints of Fig. 4a after dimensionality reduction.



(c) With $q = 10$ components, the explained variance is 91.93% and $\sigma^2 = 0.20\%$.



(c) Generated datapoints. The sample $\tilde{\mathbf{u}}$ is uniform on $[-1, 1]$ for the two first components and zero for the others. The horizontal axis varies in the first component and the vertical one in the second component.

Figure 3. Visualisation of the Probabilistic PCA reconstruction (in blue) the classical KPCA (in red). Samples generated by are also given (in grey). The dataset contains $N = 20$ points (in black).

Figure 4. The Probabilistic PCA dual formulation on the MNIST dataset restricted to 0's and 1's, with $N = 500$ datapoints, with $q = 2$ components. The explained variance is 27.97% and $\sigma^2 = 0.14\%$.

	Name	Space	Values
Dual	$\tilde{\mathbf{K}}_c$	$\mathbb{R}^{N \times N}$	$(\tilde{\mathbf{k}}_c)_{i,j} = k_c(\mathbf{x}_i, \mathbf{x}_j)$
	$\tilde{\mathbf{E}}$	$\mathbb{R}^{N \times N}$	$(\tilde{\mathbf{E}})_{i,j} = \mathbf{e}_i^* \mathbf{e}_j$
	$\tilde{\mathbf{R}}$	$\mathbb{R}^{q \times q}$	$\tilde{\mathbf{R}} = \mathbf{I}_q$
	$\tilde{\mathbf{h}}$	\mathbb{R}^q	$(\tilde{\mathbf{h}})_p = \mathbf{e}_p^* \mathbf{h}$
	$\tilde{\mathbf{k}}_c$	\mathbb{R}^N	$(\tilde{\mathbf{k}}_c)_i = \mathbf{e}_i^* \mathbf{k}_c$
Both	$\tilde{\mathbf{\Lambda}}$	$\mathbb{R}_{\geq 0}^{N \times N}$	$\tilde{\mathbf{\Lambda}} = \text{diag}(\lambda_1, \dots, \lambda_N)$
	$\tilde{\mathbf{S}}$	$\mathbb{R}_{\geq 0}^{q \times q}$	$\tilde{\mathbf{S}} = \text{diag}(s_1, \dots, s_q)$
Primal	$\tilde{\mathbf{C}}_c$	$\mathbb{R}^{d \times d}$	$(\tilde{\mathbf{C}}_c)_{i,j} = (\mathbf{u}_i^* \Phi_c) \circ (\mathbf{u}_j^* \Phi_c)^*$
	$\tilde{\Phi}_c$	$\mathbb{R}^{d \times N}$	$(\tilde{\Phi}_c)_{i,j} = \mathbf{u}_i^* \Phi_c \mathbf{e}_j$
	$\tilde{\mathbf{V}}$	$\mathbb{R}^{d \times N}$	$(\tilde{\mathbf{V}})_{i,j} = \mathbf{u}_i^* \mathbf{v}_j$
	$\tilde{\mathbf{P}}$	$\mathbb{R}^{d \times q}$	$(\tilde{\mathbf{P}})_{i,p} = \mathbf{v}_i^* \mathbf{p}_p$
	$\tilde{\phi}$	\mathbb{R}^d	$(\tilde{\phi})_i = \mathbf{v}_i^* \phi$
	$\tilde{\phi}_c$	\mathbb{R}^d	$(\tilde{\phi}_c)_i = \mathbf{v}_i^* \phi_c$

Table 3. Representation of the various operators and elements in their respective canonical basis, as matrices and vectors. The primal representations exist if and only if $\dim(\mathcal{H}) = d$ is finite.

Name	Space	Trained
$\tilde{\mathbf{W}}$	$\mathbb{R}^{d \times q}$	$\tilde{\mathbf{V}}_{1:N,1:q} (\tilde{\mathbf{\Lambda}}_{1:q,1:q} / N - \sigma^2 \mathbf{I}_q)^{1/2}$
$\tilde{\mathbf{A}}$	$\mathbb{R}^{N \times q}$	$\tilde{\mathbf{E}}_{1:N,1:q} (\mathbf{I}_q / N - \sigma^2 (\tilde{\mathbf{\Lambda}}_{1:q,1:q})^{-1})^{1/2}$
$\tilde{\mathbf{B}}$	$\mathbb{R}^{N \times q}$	$\tilde{\mathbf{E}} \tilde{\mathbf{\Lambda}}^{1/2} \begin{bmatrix} (N)^{-1/2} \tilde{\mathbf{\Lambda}}_{1:q,1:q}^{1/2} & \mathbf{0} \\ \mathbf{0} & \sigma \mathbf{I}_{N-q} \end{bmatrix}$

Table 4. Value of the different operators in the canonical basis, after training.

5.1. Model

The direct application of the theoretical discussions of the previous sections leads to the decompositions $\tilde{\mathbf{K}}_c = \tilde{\mathbf{E}} \tilde{\mathbf{\Lambda}} \tilde{\mathbf{E}}^\top$, $\tilde{\mathbf{C}}_c = \tilde{\mathbf{V}} \tilde{\mathbf{\Lambda}} \tilde{\mathbf{V}}^\top$, $\tilde{\Phi}_c = \tilde{\mathbf{V}} \tilde{\mathbf{\Lambda}}^{1/2} \tilde{\mathbf{E}}^\top$. The value of the operators after training are given in Table 4. Once the model is trained, we can verify that $\tilde{\mathbf{W}} = \tilde{\Phi}_c \tilde{\mathbf{A}}$. We can also have a look at the hidden variables. A way to do it is to consider the MAP of \mathbf{h} given ϕ or \mathbf{k} . We have

$$\mathbf{h}_{\text{MAP}} = N \tilde{\mathbf{\Lambda}}_{1:q,1:q}^{-1} \tilde{\mathbf{A}}^\top \tilde{\mathbf{k}}_c \quad (\text{if } \text{rank}(\tilde{\mathbf{K}}_c) \geq q), \quad (20)$$

$$= N \tilde{\mathbf{\Lambda}}_{1:q,1:q}^{-1} \tilde{\mathbf{W}}^\top (\tilde{\phi} - \tilde{\phi}_c) \quad (\text{if } \mathcal{H} \text{ is finite}), \quad (21)$$

and

$$(\mathbf{k}_c)_{\text{MAP}} = \tilde{\mathbf{K}}_c \tilde{\mathbf{A}} \tilde{\mathbf{h}} \quad (\text{if } \text{rank}(\tilde{\mathbf{K}}_c) \geq q), \quad (22)$$

$$\phi_{\text{MAP}} = \tilde{\mathbf{W}} \tilde{\mathbf{h}} + \tilde{\phi}_c \quad (\text{if } \mathcal{H} \text{ is finite}). \quad (23)$$

As developed in Section 4, we can easily generate samples in both feature and kernel representations. For the latter and in canonical basis, it becomes

$$\tilde{\mathbf{k}}_c = \tilde{\mathbf{B}} \tilde{\mathbf{u}}, \quad \text{with } \tilde{\mathbf{u}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_N). \quad (24)$$

5.2. Examples

As the primal case is already treated by (Tipping & Bishop, 1999), we consider here the model in its dual formulation. A toy example can be found in Fig. 3. We use an RBF kernel $k(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|_2^2 / (2\gamma^2))$ with bandwidth $\gamma = 2$. As the number of components increases, the mean variance of the $N - q$ unused components σ^2 becomes smaller and the model tends to the classical KPCA model. Another way to reduce σ^2 is to increase the number of components q , with $\sigma^2 \rightarrow 0$ when $q \rightarrow N$. This can be observed in Fig. 3c: the Probabilistic PCA model resembles closely the KPCA model, whereas more variance is left over, *i.e.* not projected back, in Figs. 3a and 3b. The results of the generation is Gaussian, which is a consequence of the linearity of the preimage method chosen (Eq. (19)). Here again, as the number of components increases and σ^2 decreases, the model is allowed to project back more variance and the distribution becomes wider. Another example on the MNIST dataset (LeCun & Cortes, 2010) with the RBF kernel with $\gamma = 4$ is given at Fig. 4.

6. Conclusion

Probabilistic Interpretation. By reformulating the Prob. PCA model in Hilbert space, we were able to define a formulation of it. Likewise Prob. PCA in primal was englobing classical PCA (with $\sigma^2 \rightarrow 0$), Prob. PCA in dual is also englobing KPCA in the same limit. Furthermore, we are now able to sample in dual space, enhancing the understanding of the generation done with KPCA.

Limitations. As most kernel methods, the model is still limited by the need of a preimage method to go back to the input space once a sample is projected or generated. Furthermore, training the model in dual required to find the q first eigenvalues of the kernel matrix, which may become expensive as the number of datapoints N increases. Generating renders the problem even worse as it requires the computation of all eigenvalues. The model also requires to determine a σ^2 or alternatively a latent dimension q .

Acknowledgements

EU: The research leading to these results has received funding from the European Research Council under the European Union’s Horizon 2020 research and innovation program / ERC Advanced Grant E-DUALITY (787960). This paper reflects only the authors’ views and the Union is not liable for any use that may be made of the contained information. Research Council KUL: Tensor Tools for Taming the Curse iBOF/23/064, Optimization frameworks for deep kernel machines C14/18/068. Flemish Government: FWO projects: GOA4917N (Deep Restricted Kernel Machines: Methods and Foundations), PhD/Postdoc grant. This research received funding from the Flemish Government (AI Research Program). Henri De Plaen and Johan A. K. Suykens are also affiliated to Leuven.AI – KU Leuven institute for AI, B-3000, Leuven, Belgium.

References

- Alaíz, C. M., Fanuel, M., and Suykens, J. A. K. Convex formulation for kernel PCA and its use in semisupervised learning. *IEEE Transactions on Neural Networks and Learning Systems*, 29(8):3863–3869, 2018. doi: 10.1109/TNNLS.2017.2709838.
- Bartholomew, D. J., Knott, M., and Moustaki, I. *Latent variable models and factor analysis: A unified approach*, volume 904. John Wiley & Sons, 2011.
- Basilevsky, A. T. *Statistical factor analysis and related methods: theory and applications*. John Wiley & Sons, 2009.
- LeCun, Y. and Cortes, C. MNIST handwritten digit database. 2010. URL <http://yann.lecun.com/exdb/mnist/>.
- Mika, S., Schölkopf, B., Smola, A., Müller, K.-R., Scholz, M., and Rätsch, G. Kernel PCA and de-noising in feature spaces. *Advances in neural information processing systems*, 11, 1998.
- Pandey, A., Schreurs, J., and Suykens, J. A. K. Robust generative restricted kernel machines using weighted conjugate feature duality. In *Machine Learning, Optimization, and Data Science: 6th International Conference, LOD 2020, Siena, Italy, July 19–23, 2020, Revised Selected Papers, Part I 6*, pp. 613–624. Springer, 2020.
- Pandey, A., Schreurs, J., and Suykens, J. A. K. Generative restricted kernel machines: A framework for multi-view generation and disentangled feature learning. *Neural Networks*, 135:177–191, 2021.
- Pandey, A., De Meulemeester, H., De Plaen, H., De Moor, B., and Suykens, J. A. K. Recurrent restricted kernel machines for time-series forecasting. *Proceedings of ESANN 2022*, 2022a.
- Pandey, A., Fanuel, M., Schreurs, J., and Suykens, J. A. K. Disentangled representation learning and generation with manifold optimization. *Neural Computation*, 34(10): 2009–2036, 2022b.
- Schölkopf, B. and Smola, A. J. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001. ISBN 0262194759.
- Schölkopf, B., Smola, A., and Müller, K.-R. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation*, 10(5):1299–1319, 1998.
- Schreurs, J. and Suykens, J. A. K. Generative kernel PCA. *26th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, pp. 129–134, 2018.
- Suykens, J. A. K. Deep Restricted Kernel Machines Using Conjugate Feature Duality. *Neural Computation*, 29(8): 2123–2163, 08 2017. ISSN 0899-7667. doi: 10.1162/neco_a_00984.
- Suykens, J. A. K., Van Gestel, T., Vandewalle, J., and De Moor, B. A support vector machine formulation to PCA analysis and its kernel version. *IEEE Transactions on Neural Networks*, 14(2):447–450, 2003. doi: 10.1109/TNN.2003.809414.
- Tipping, M. E. and Bishop, C. M. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- Winant, D., Schreurs, J., and Suykens, J. A. K. Latent space exploration using generative kernel PCA. In *Artificial Intelligence and Machine Learning: 31st Benelux AI Conference, BNAIC 2019, and 28th Belgian-Dutch Machine Learning Conference, BENELEARN 2019, Brussels, Belgium, November 6-8, 2019, Revised Selected Papers 28*, pp. 70–82. Springer, 2020.
- Zhang, Z., Wang, G., Yeung, D.-Y., and Kwok, J. T. Probabilistic kernel principal component analysis. *Department of Computer Science, The Hong Kong University of Science and Technology, Tech. Rep*, 2004.

A. Theoretical Development

For brevity of notations, we will define the *expectation value* as $\|\mathbf{a}\|_{\Sigma}^2 = \mathbf{a}^* \Sigma \mathbf{a} \in \mathbb{R}_{\geq 0}$, with \mathbf{a} an element of a Hilbert space and Σ a linear operator from and to that space. The norm is a particular case with the identity as operator $\Sigma = \mathbf{I}$. The density function of the multivariate normal distribution $\mathbf{a} \sim \mathcal{N}(\mathbf{b}, \Sigma)$ can be rewritten as $\frac{1}{Z} \exp(-\frac{1}{2} \|\mathbf{a} - \mathbf{b}\|_{\Sigma^{-1}}^2)$.

A.1. Primal and Dual Spaces

Lemma 1. *The operators $\Phi_c \circ \Phi_c^*$ and $\Phi_c^* \circ \Phi_c$ are self-adjoint, positive semi-definite and share the same non-zero eigenvalues. In particular, we have the eigenvector relations $\mathbf{v}_i = (\lambda_i)^{-1/2} \Phi_c \boldsymbol{\epsilon}_i$ and $\boldsymbol{\epsilon}_i = (\lambda_i)^{-1/2} \Phi_c^* \mathbf{v}_i$.*

Proof. The result stated here is inspired from (Alaíz et al., 2018). Self-adjointness is a consequence of the definition of the inner product, which also guarantees the positive semi-definiteness. (\implies) Let us suppose that $(\Phi_c^* \circ \Phi_c) \boldsymbol{\epsilon}_i = \lambda_i \boldsymbol{\epsilon}_i$ with $\lambda_i \neq 0$. We then have $\Phi_c \circ (\Phi_c^* \circ \Phi_c) \boldsymbol{\epsilon}_i = (\Phi_c \circ \Phi_c^*) \circ \Phi_c \boldsymbol{\epsilon}_i = \lambda_i \Phi_c \boldsymbol{\epsilon}_i$. Hence, we have that $\Phi_c \boldsymbol{\epsilon}_i$ is eigenvector, but not necessarily normalized. In fact, its norm is given by $\boldsymbol{\epsilon}_i^* \Phi_c^* \circ \Phi_c \boldsymbol{\epsilon}_i = \lambda_i$. We thus have the relation $\mathbf{u}_i = (\lambda_i)^{-1/2} \Phi_c \boldsymbol{\epsilon}_i$. (\impliedby) We suppose now that $(\Phi_c \circ \Phi_c^*) \mathbf{u}_i$, which leads to $\Phi_c^* \circ (\Phi_c \circ \Phi_c^*) \mathbf{u}_i = (\Phi_c^* \circ \Phi_c) \circ \Phi_c^* \mathbf{u}_i = \lambda_i \mathbf{u}_i$. Again, we have the relation $\boldsymbol{\epsilon}_i = (\lambda_i)^{-1/2} \Phi_c^* \mathbf{u}_i$. \square

A.2. Primal Model

A.2.1. FEATURE DISTRIBUTION

Definition 1 (Conditional Feature Distribution). *Considering $\phi, \boldsymbol{\mu} \in \mathcal{H}_{\mathcal{E}}$, $\mathbf{h} \in \mathcal{L}$, a linear operator $\mathbf{W} : \mathcal{L} \rightarrow \mathcal{H}_{\mathcal{L}}$ and its adjoint $\mathbf{W}^* : \mathcal{H}_{\mathcal{L}} \rightarrow \mathcal{L}$. We define the conditional probability distribution of the primal feature variable ϕ with variance $\sigma^2 \in \mathbb{R}_{>0}$ as*

$$p(\phi | \mathbf{h}) = \frac{1}{(\sigma \sqrt{2\pi})^N} \exp\left(\frac{-1}{2\sigma^2} \|\phi - \mathbf{W}\mathbf{h} - \boldsymbol{\mu}\|_{\mathcal{H}_{\mathcal{E}}}^2\right). \quad (25)$$

The following proposition verifies that the conditional distribution of ϕ given \mathbf{h} (Def. 1) is well defined. To ease the readability, we first consider a Lemma.

Lemma 2. *Given values $a \in \mathbb{R}_{>0}$ and $b, c \in \mathbb{R}$, we have the following integral*

$$\int_{\mathbb{R}} \exp\left(-\frac{1}{2}ax^2 + bx + c\right) dx = \sqrt{\frac{2\pi}{a}} \exp\left(\frac{b^2}{2a} + c\right). \quad (26)$$

Proof. We first find the primitive $\int \exp(-\frac{1}{2}ax^2 + bx + c) dx = \sqrt{\frac{\pi}{2a}} \exp(\frac{b^2}{2a} + c) \operatorname{erf}\left(\frac{ax-b}{\sqrt{2a}}\right)$, with the error function defined as $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$. Indeed, we have $\frac{d}{dx} \operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} e^{-x^2}$. It suffices then to derivate the primitive to verify them. To conclude the proof, it suffices to notice that the error function is symmetric and that $\lim_{x \rightarrow +\infty} \operatorname{erf}(x) = 1$. \square

From now on, we will consider the singular value decomposition of the primal interconnection linear operator $\mathbf{W} = \sum_{p=1}^q s_p \boldsymbol{\rho}_p \mathbf{r}_p^*$, with $\boldsymbol{\rho}_p \in \mathcal{H}_{\mathcal{E}}$ and $\mathbf{r}_p^* \in \mathcal{L}^*$ two sets of orthonormal variables and $s_p \in \mathbb{R}_{>0}$ the singular values.

Proposition 1. *Def. 1 is a well-defined distribution. More specifically, its measure is normalized:*

$$\int_{\mathcal{L}} p(\phi | \mathbf{h}) d\phi = 1. \quad (27)$$

Proof. By considering the singular value decomposition of \mathbf{W} , the term inside the exponential becomes

$$-\frac{1}{2\sigma^2} \sum_{i=1}^N \left\{ (\mathbf{v}_i^* \phi)^2 - (\mathbf{v}_i^* \phi) \left(2 \sum_{p=1}^q ((\mathbf{r}_p^* \mathbf{h}) (\boldsymbol{\rho}_p^* \mathbf{v}_i) s_p) + (\mathbf{v}_i^* \boldsymbol{\mu}) \right) \right\} + C, \quad (28)$$

with $C = -\frac{1}{2\sigma^2} (\mathbf{W}\mathbf{h} + \boldsymbol{\mu})^* (\mathbf{W}\mathbf{h} + \boldsymbol{\mu}) = -\frac{1}{2\sigma^2} \sum_{i=1}^N \sum_{p=1}^q ((\mathbf{r}_p^* \mathbf{h})(\boldsymbol{\rho}_p^* \mathbf{v}_i)_{s_p} + (\mathbf{v}_i^* \boldsymbol{\mu})^2)$. Integrating over \mathcal{H}_ε means integrating over span $\{\mathbf{v}_1, \dots, \mathbf{v}_N\}$, thus for all $(\mathbf{v}_i^* \boldsymbol{\phi}) \in \mathbb{R}$. Using Lemma 2 and by Fubini's theorem, we have

$$\int_{\mathcal{F}} p(\boldsymbol{\phi} | \mathbf{h}, \boldsymbol{\rho}) d\boldsymbol{\phi} = (\sigma\sqrt{2\pi})^{-N} \exp(-C) \int_{\mathbb{R}^N} \exp\left(\sum_{i=1}^N \left(-\frac{1}{2\sigma^2} x_i^2 + b_i x_i\right)\right) d\mathbf{x}, \quad (29)$$

$$= (\sigma\sqrt{2\pi})^{-N} \exp(-C) \prod_{i=1}^N \int_{\mathbb{R}} \exp\left(-\frac{1}{2\sigma^2} x_i^2 + b_i x_i\right) dx_i, \quad (30)$$

$$= (\sigma\sqrt{2\pi})^{-N} \exp(-C) \prod_{i=1}^N \sigma\sqrt{2\pi} \exp\left(\frac{1}{2}\sigma^2 b_i^2\right), \quad (31)$$

$$= \exp\left(\frac{1}{2}\sigma^2 \sum_{i=1}^N b_i^2 - C\right), \quad (32)$$

with $b_i = \frac{1}{\sigma^2} ((\mathbf{r}_p^* \mathbf{h})(\boldsymbol{\rho}_p^* \mathbf{v}_i)_{s_p} + (\mathbf{v}_i^* \boldsymbol{\mu}))$. The proof is concluded by observing that $C = \frac{1}{2}\sigma^2 \sum_{i=1}^N b_i^2$. \square

We can now consider the marginal distribution of the feature representation $\boldsymbol{\phi}$.

Lemma 3. *Both linear operators $\mathbf{W} \circ \mathbf{W}^* + \sigma^2 \mathbf{I}_{\mathcal{H}_\varepsilon}$ and $\mathbf{W}^* \circ \mathbf{W} + \sigma^2 \mathbf{I}_{\mathcal{L}}$ are positive definite, of full rank and invertible. The linear operator $\mathbf{W} \circ \mathbf{W}^* + \sigma^2 \mathbf{I}_{\mathcal{H}_\varepsilon}$ shares the q non-zero eigenvalues of $\mathbf{W}^* \circ \mathbf{W} + \sigma^2 \mathbf{I}_{\mathcal{L}}$, with the remaining $N - q$ eigenvalues being equal to σ^2 . In particular, we have*

$$\text{eig}(\mathbf{W}^* \circ \mathbf{W} + \sigma^2 \mathbf{I}_{\mathcal{L}}) = \{s_p^2 + \sigma^2\}_{p=1}^q, \quad (33)$$

$$\text{eig}(\mathbf{W} \circ \mathbf{W}^* + \sigma^2 \mathbf{I}_{\mathcal{H}_\varepsilon}) = \{s_p^2 + \sigma^2\}_{p=1}^q \cup \{\sigma^2\}^{N-q}. \quad (34)$$

Proof. We consider the eigenvalues of $\mathbf{W}^* \circ \mathbf{W} + \mathbf{I}_{\mathcal{L}} = \sum_{p=1}^q (s_p^2 + \sigma^2) \boldsymbol{\rho}_p \boldsymbol{\rho}_p^*$. Its eigenvalues are thus given by $s_p^2 + \sigma^2$ for $p = 1, \dots, q$ and we directly conclude that they are all strictly positive. In a similar fashion, the eigenvalues of $\mathbf{W} \circ \mathbf{W}^* + \sigma^2 \mathbf{I}_{\mathcal{H}_\varepsilon} = \sum_{p=1}^q (s_p^2 + \sigma^2) \mathbf{r}_p \mathbf{r}_p^* + \sigma^2 \mathbb{P}_{\mathcal{H}_\varepsilon}$ are given by $s_p^2 + \sigma^2$ for $p = 1, \dots, q$ and σ^2 for $k = q + 1, \dots, N$. All the eigenvalues are strictly positive. Hence are the operators positive definite, full rank and invertible. \square

Proposition 2 (Marginal Feature Distribution). *Assuming the conditional distribution of $\boldsymbol{\phi}$ given \mathbf{h} (Def. 1), assumed normally distributed, i.e., $p(\boldsymbol{\phi}) = (2\pi)^{-q/2} \exp(-\frac{1}{2} \mathbf{h}^* \mathbf{h})$, the posterior distribution of the primal feature vector $\boldsymbol{\phi}$ is given by*

$$p(\boldsymbol{\phi}) = \frac{1}{Z_\phi} \exp\left(-\frac{1}{2} \|\boldsymbol{\phi} - \boldsymbol{\mu}\|_{\boldsymbol{\Sigma}_{\phi|\mathbf{W}}}^2\right), \quad (35)$$

with $Z_\phi = (2\pi)^{N/2} ((\sigma^2)^{N-q} \prod_{p=1}^q (s_p^2 + \sigma^2))^{1/2}$ and $\boldsymbol{\Sigma}_{\phi|\mathbf{W}} = \mathbf{W} \circ \mathbf{W}^* + \sigma^2 \mathbf{I}_{\mathcal{H}_\varepsilon}$

Proof. The joint probability distribution is given by Bayes' theorem as $p(\boldsymbol{\phi}, \mathbf{h}) = p(\boldsymbol{\phi} | \mathbf{h}) p(\mathbf{h})$. The integration proceeds in a very similar fashion as the proof of Prop. 1. Similarly, the term inside the exponential becomes

$$-\frac{1}{2\sigma^2} \sum_{p=1}^q \left\{ (\mathbf{r}_p^* \mathbf{h})^2 (s_p^2 + \sigma^2) - (\mathbf{r}_p^* \mathbf{h}) \left(2s_p (\boldsymbol{\rho}_p^* (\boldsymbol{\phi} - \boldsymbol{\mu}))\right) \right\} + D, \quad (36)$$

with

$$D = -\frac{1}{2\sigma^2} (\boldsymbol{\phi} - \boldsymbol{\mu})^* (\boldsymbol{\phi} - \boldsymbol{\mu}), \quad (37)$$

$$= -\frac{1}{2\sigma^2} \sum_{p=1}^q (\boldsymbol{\rho}_p^* (\boldsymbol{\phi} - \boldsymbol{\mu}))^2 - \frac{1}{2\sigma^2} (\boldsymbol{\phi} - \boldsymbol{\mu})^* \mathbb{P}_{\mathcal{H}_\varepsilon} (\boldsymbol{\phi} - \boldsymbol{\mu}). \quad (38)$$

Here again, integrating on the whole latent space \mathcal{L} means integrating for all $(\mathbf{r}_p^* \mathbf{h}) \in \mathbb{R}$. Again, we use Lemma 2 and Fubini's theorem. After some simplification inside the exponential, we can use the development made in Lemma 3 to find the following:

$$-\frac{1}{2} \left\{ \sum_{p=1}^q \frac{(\boldsymbol{\varrho}_p^*(\boldsymbol{\phi} - \boldsymbol{\mu}))^2}{s_p^2 + \sigma^2} + \frac{1}{\sigma^2} (\boldsymbol{\phi} - \boldsymbol{\mu})^* \mathbb{P}_{\mathcal{H}_\varepsilon} (\boldsymbol{\phi} - \boldsymbol{\mu}) \right\} = -\frac{1}{2} (\boldsymbol{\phi} - \boldsymbol{\mu})^* (\mathbf{W} \circ \mathbf{W}^* + \sigma^2 \mathbf{I}_{\mathcal{H}_\varepsilon})^{-1} (\boldsymbol{\phi} - \boldsymbol{\mu}). \quad (39)$$

The normalization term follows similarly to Proposition 1 and we can verify the consistency of the obtained distribution by looking at the spectrum using Lemma 3. \square

A.2.2. MAXIMUM LIKELIHOOD

Training the model in primal corresponds to maximizing the likelihood of the observations in the finite feature space \mathcal{H}_ε .

Proposition 3 (Primal Maximum Likelihood). *Provided $\sigma^2 \leq \lambda_q/N$, where λ_q is the q^{th} greatest eigenvalue of $\boldsymbol{\Phi}_c \circ \boldsymbol{\Phi}_c^*$, the Maximum Likelihood (ML) of \mathbf{W} and $\boldsymbol{\mu}$ given the observations $\{\phi_i\}_{i=1}^N$ is given by*

$$\boldsymbol{\mu}_{\text{ML}} = \boldsymbol{\phi}_c, \quad (40)$$

$$\mathbf{W}_{\text{ML}} = \sum_{p=1}^q \sqrt{\lambda_p/N - \sigma^2} \mathbf{v}_p \mathbf{r}_p^*, \quad (41)$$

with $\{(\lambda_p, \mathbf{v}_p)\}_{p=1}^q$ the greatest eigenpairs of $\boldsymbol{\Phi}_c \circ \boldsymbol{\Phi}_c^*$ (w.r.t. the eigenvalues).

Proof. The maximum likelihood of the observations $\{\phi_i\}_{i=1}^N$ is computed as $\text{argmax}_{\mathbf{W}, \boldsymbol{\mu}} \log(\prod_{i=1}^N p(\phi_i | \mathbf{W}, \boldsymbol{\mu})) = \text{argmax}_{\mathbf{W}, \boldsymbol{\mu}} L_\phi$, with L_ϕ the likelihood function, which can be written as

$$L_\phi = -\frac{N}{2} \left\{ N \log(2\pi) + (N - q) \log(\sigma^2) + \sum_{p=1}^q \log(s_p^2 + \sigma^2) + \frac{1}{N\sigma^2} \sum_{i=1}^N (\phi_i - \boldsymbol{\mu})^* \mathbb{P}_{\mathcal{H}_\varepsilon} (\phi_i - \boldsymbol{\mu}) + \frac{1}{N} \sum_{p=1}^q \frac{1}{s_p^2 + \sigma^2} \left(\sum_{i=1}^N (\boldsymbol{\varrho}_p^*(\phi_i - \boldsymbol{\mu})) \right)^2 \right\}. \quad (42)$$

The optimization of the mean $\boldsymbol{\mu}$ is trivial and we have $\boldsymbol{\mu}_{\text{ML}} = \boldsymbol{\phi}_c$. The optimization of $\mathbf{W} = \sum_{p=1}^q s_p \boldsymbol{\varrho}_p \mathbf{r}_p^*$ is less trivial. We first note that optimizing for $\{\boldsymbol{\varrho}_p\}_{p=1}^q$, $\{s_p^2\}_{p=1}^q$ and $\{\mathbf{r}_p\}_{p=1}^q$ is not identifiable: indeed, two singular values may be equal. Furthermore, the likelihood function L_ϕ is independent of the basis $\{\mathbf{r}_p\}_{p=1}^q$ and for any solution of $\boldsymbol{\varrho}_p$, we must also admit its opposite $-\boldsymbol{\varrho}_p$ too. In addition to that, computing the saddle points of the likelihood is not straightforward as we cannot consider the vectors $\{\boldsymbol{\varrho}_p\}_{p=1}^q$ to be independent variables as they must remain orthonormal: the variation of one basis vector $\boldsymbol{\varrho}_p$ must keep it normalized and has an influence on the other basis vectors. For s_p however, the variations may happen with the only constraint of strict positivity. We may thus consider

$$\frac{\partial L_\phi}{\partial s_p} = 0 \iff s_p^2 + \sigma^2 = \frac{1}{N} \left(\boldsymbol{\varrho}_p^* \left(\sum_{n=1}^N \phi_n - \boldsymbol{\mu} \right) \right)^2. \quad (43)$$

We now consider the fact that $\mathbb{P}_{\mathcal{H}_\varepsilon} = \mathbf{I}_{\mathcal{H}_\varepsilon} - \sum_{p=1}^q \boldsymbol{\varrho}_p \boldsymbol{\varrho}_p^*$. The likelihood function restricted to the terms dependent on $\boldsymbol{\varrho}_p$ reduces to

$$\frac{1}{2} \left\{ \sum_{p=1}^q \left(\frac{1}{s_p^2 + \sigma^2} - \frac{1}{\sigma^2} \right) \boldsymbol{\Phi}_c^* \circ (\boldsymbol{\varrho}_p \boldsymbol{\varrho}_p^*) \circ \boldsymbol{\Phi}_c \right\}. \quad (44)$$

From there, we can deduce that the likelihood is maximized when $\{\boldsymbol{\varrho}_p\}_{p=1}^q$ forms a basis of $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_p)$, with $\{\mathbf{v}_p\}_{p=1}^q$ the q greatest eigenvectors of $\boldsymbol{\Phi}_c \circ \boldsymbol{\Phi}_c^*$ (w.r.t. the eigenvalues λ_p 's). We may thus identify both basis: $\boldsymbol{\varrho}_p = \mathbf{v}_p$ and by consequence $s_p^2 + \sigma^2 = \lambda_p/N$ (from Eq. (43)), for all $p = 1, \dots, q$.

What about the other choices of basis? At the end, it would not change anything as it would be identified to the same solution. Indeed, as the $\boldsymbol{\varrho}_p$'s could always be written as a linear combination of $\{\mathbf{v}_p\}_{p=1}^q$, we could always write the interconnection operator as $\mathbf{W}_{\text{ML}} = \sum_{p=1}^q (\lambda_p/N - \sigma^2) \mathbf{v}_p \mathbf{r}_p^*$ as the choice of the orthonormal basis $\{\mathbf{r}_p\}_{p=1}^q$ is arbitrary. \square

A.2.3. DIMENSIONALITY REDUCTION IN FEATURE SPACE

Proposition 4 (Primal Conditional Latent Distribution). *The conditional distribution of \mathbf{h} given ϕ is given by*

$$p(\mathbf{h}|\phi) = \frac{1}{Z_{\mathbf{h}|\phi}} \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{h} - \mathbf{m}\|_{\Sigma_{\mathbf{h}|\phi}^{-1}}^2\right), \quad (45)$$

with $\Sigma_{\mathbf{h}|\phi} = (\mathbf{W}^* \circ \mathbf{W} + \sigma^2 \mathbf{I}_{\mathcal{L}})^{-1}$, $\mathbf{m} = (\Sigma_{\mathbf{h}|\phi})^{-1} \mathbf{W}^*(\phi - \boldsymbol{\mu})$ and $Z_{\mathbf{h}|\phi} = (\sigma\sqrt{2\pi})^q (\prod_{p=1}^q (s_p^2 + \sigma^2))^{-1/2}$.

Proof. The methodology is analogous to Props. 1 and 2. \square

A.3. Dual Model

A.3.1. REPRESENTATION

Proposition 5 (Dual Representation). *Given any interconnection operator $\mathbf{W} : \mathcal{L} \rightarrow \mathcal{H}_{\mathcal{L}}$, we have the following representation:*

$$\mathbf{W} = \Phi_c \circ \mathbf{A}, \quad (46)$$

with $\mathbf{A} : \mathcal{L} \rightarrow \mathcal{L}$, the dual interconnection operator.

Proof. Following our definitions, we know that $\Phi_c \circ \Phi_c^* \geq q$. By consequence, the linear operator $\Phi_c : \mathcal{H} \rightarrow \mathcal{E}$ has at least q non-zero singular values. As we have $\dim(\mathcal{L}) = \dim(\mathcal{H}_{\mathcal{L}})$, the proof is concluded by recalling that the primal interconnection operator $\mathbf{W} : \mathcal{L} \rightarrow \mathcal{H}_{\mathcal{L}}$ is also linear. \square

Proposition 6 (Dual Maximum Likelihood). *Given the operator \mathbf{W}_{ML} (Prop. 3) and provided $\sigma^2 \leq \lambda_q/N$, where λ_q is the q^{th} greatest eigenvalue of $\Phi_c^* \circ \Phi_c$, the dual interconnection operator $\mathbf{A}_{\text{ML}} : \mathcal{L} \rightarrow \mathcal{L}$ is given by:*

$$\mathbf{A}_{\text{ML}} = \sum_{p=1}^q \sqrt{1/N - \sigma^2 \lambda_p^{-1}} \boldsymbol{\epsilon}_p \mathbf{r}_p^*, \quad (47)$$

with $\{(\lambda_p, \boldsymbol{\epsilon}_p)\}_{p=1}^q$ the greatest eigenpairs of $\Phi_c^* \circ \Phi_c$ (w.r.t. the eigenvalues) and $\{\mathbf{r}_p\}_{p=1}^q$ an arbitrary orthonormal basis of the latent space \mathcal{L} .

Proof. Using the relation $\mathbf{v}_p = (\lambda_p)^{-1/2} \Phi_c \boldsymbol{\epsilon}_p$, we have

$$\mathbf{W}_{\text{ML}} = \sum_{p=1}^q \sqrt{\lambda_p/N - \sigma^2} \mathbf{v}_p \mathbf{r}_p^*, \quad (48)$$

$$= \sum_{p=1}^q (\lambda_p)^{-1/2} \sqrt{\lambda_p/N - \sigma^2} \Phi_c \boldsymbol{\epsilon}_p \mathbf{r}_p^*, \quad (49)$$

$$= \Phi_c \circ \left(\sum_{p=1}^q \sqrt{1/N - \sigma^2 \lambda_p^{-1}} \boldsymbol{\epsilon}_p \mathbf{r}_p^* \right). \quad (50)$$

The proof is concluded by insuring that $1/N - \sigma^2 \lambda_p^{-1}$ is never negative. \square

A.3.2. DIMENSIONALITY REDUCTION IN KERNEL SPACE

Proposition 7 (Conditional Kernel Distribution). *Under the same assumption as Prop. 6 and provided that $\Phi_c^* \circ \Phi_c$ is of full rank, the posterior distribution of the dual feature variable \mathbf{k}_c given the latent variable \mathbf{h} is given by*

$$p(\mathbf{k}_c|\mathbf{h}) = \frac{1}{Z_{\mathbf{k}_c|\mathbf{h}}} \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{k}_c - (\Phi_c^* \circ \Phi_c) \circ \mathbf{A}_{\text{ML}} \mathbf{h}\|_{\Sigma_{\mathbf{k}_c|\mathbf{h}}^{-1}}^2\right), \quad (51)$$

with $Z_{\mathbf{k}_c|\mathbf{h}} = (\sigma\sqrt{2\pi})^N \prod_{i=1}^N \lambda_i$ and $\Sigma_{\mathbf{k}_c|\mathbf{h}} = \Phi_c^* \circ \Phi_c$.

Proof. We consider the optimal primal interconnection operator $\mathbf{W}_{\text{ML}} = \Phi_c \circ \mathbf{A}_{\text{ML}}$ together with the optimal mean $\mu_{\text{ML}} = \phi_c$ and fill it in Def. 1. The covariance is deduced by considering $\phi - \phi_c = (\Phi_c^*)^{-1} \circ \Phi_c^*(\phi - \phi_c) = (\Phi_c^*)^{-1} \mathbf{k}_c$. \square

Proposition 8 (Dual Conditional Latent Distribution). *Under the same assumption as Prop. 6, the posterior distribution of the latent variable \mathbf{h} given the dual feature variable \mathbf{k}_c is given by*

$$p(\mathbf{h}|\mathbf{k}_c) = \frac{1}{Z_{\mathbf{h}|\mathbf{k}_c}} \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{h} - \mathbf{m}'\|_{\Sigma_{\mathbf{h}|\mathbf{k}_c}^{-1}}^2\right), \quad (52)$$

with $\Sigma_{\mathbf{h}|\mathbf{k}_c} = (\mathbf{A}_{\text{ML}}^* \circ (\Phi_c^* \circ \Phi_c) \circ \mathbf{A}_{\text{ML}} + \sigma^2 \mathbf{I}_{\mathcal{L}})^{-1}$, $Z_{\mathbf{h}|\mathbf{k}_c} = Z_{\mathbf{h}|\phi}$ (Prop. 4) and $\mathbf{m}' = (\Sigma_{\mathbf{h}|\mathbf{k}_c}^{-1} \circ \mathbf{A}_{\text{ML}}^*) \mathbf{k}_c$.

Proof. It suffices to consider the optimal primal interconnection operator $\mathbf{W}_{\text{ML}} = \Phi_c \circ \mathbf{A}_{\text{ML}}$, as well as the optimal mean $\mu_{\text{ML}} = \phi_c$ and fill both in Prop. 4. \square

A.3.3. KERNEL SAMPLING

Proposition 9 (Marginal Kernel Distribution). *Provided that $\Phi_c^* \circ \Phi_c$ is of full rank and given the assumptions of Prop. 6, the posterior distribution of the the dual feature \mathbf{k}_c after optimization is given by*

$$p(\mathbf{k}_c) = \frac{1}{Z_{\mathbf{k}_c}} \exp\left(-\frac{1}{2} \|\mathbf{k}_c\|_{\Sigma_{\mathbf{k}_c}^{-1}}^2\right), \quad (53)$$

with $\Sigma_{\mathbf{k}_c} = \mathbf{A}_{\text{ML}}^* \circ \mathbf{A}_{\text{ML}} + \sigma^2 (\Phi_c^* \circ \Phi_c)^{-1}$ and $Z_{\mathbf{k}_c} = (2\pi)^{N/2} N^{-q/2} \prod_{p=1}^q \lambda_p \prod_{p=q+1}^N (\sigma^2 \lambda_p)^{1/2}$.

Proof. We start from Prop. 2 and develop the covariance:

$$\mathbf{W}_{\text{ML}} \mathbf{W}_{\text{ML}}^* + \sigma^2 \mathbf{I}_{\mathcal{H}_{\mathcal{E}}} = \Phi_c \circ \left(\mathbf{A}^* \circ \mathbf{A} + \sigma^2 (\Phi_c^* \circ \Phi_c)^{-1} \right) \circ \Phi_c^*, \quad (54)$$

$$= \Phi_c \circ \left(\frac{1}{N} \sum_{p=1}^q \epsilon_p \epsilon_p^* + \sigma^2 \sum_{p=q+1}^N \lambda_p^{-1} \epsilon_p \epsilon_p^* \right) \circ \Phi_c^*. \quad (55)$$

We set $\mu_{\text{ML}} = \phi_c$ and observe that

$$(\phi - \phi_c)^* (\mathbf{W}_{\text{ML}} \mathbf{W}_{\text{ML}}^* + \sigma^2 \mathbf{I}_{\mathcal{H}_{\mathcal{E}}})^{-1} (\phi - \phi_c), = \mathbf{k}_c^* (\Phi_c^* \circ (\mathbf{W}_{\text{ML}} \mathbf{W}_{\text{ML}}^* + \sigma^2 \mathbf{I}_{\mathcal{H}_{\mathcal{E}}}) \circ \Phi_c)^{-1} \mathbf{k}_c, \quad (56)$$

$$= \mathbf{k}_c^* \left(\frac{1}{N} \sum_{p=1}^q \lambda_p^2 \epsilon_p \epsilon_p^* + \sigma^2 \sum_{p=q+1}^N \lambda_p \epsilon_p \epsilon_p^* \right) \mathbf{k}_c. \quad (57)$$

\square