

Dimensionality Reduction and Wasserstein Stability for Kernel Regression

Stephan Eckstein

*Department of Mathematics
ETH Zürich
Zurich, CH*

SECKSTEIN@ETHZ.CH

Armin Iske

*Department of Mathematics
Universität Hamburg
Hamburg, DE*

ARMIN.ISKE@UNI-HAMBURG.DE

Mathias Trabs

*Department of Mathematics
Karlsruhe Institute of Technology
Karlsruhe, DE*

TRABS@KIT.EDU

Editor: Miguel Carreira-Perpinan

Abstract

In a high-dimensional regression framework, we study consequences of the naive two-step procedure where first the dimension of the input variables is reduced and second, the reduced input variables are used to predict the output variable with kernel regression. In order to analyze the resulting regression errors, a novel stability result for kernel regression with respect to the Wasserstein distance is derived. This allows us to bound errors that occur when perturbed input data is used to fit the regression function. We apply the general stability result to principal component analysis (PCA). Exploiting known estimates from the literature on both principal component analysis and kernel regression, we deduce convergence rates for the two-step procedure. The latter turns out to be particularly useful in a semi-supervised setting.

Keywords: Kernel regression, Dimensionality reduction, Principal component analysis, Stability, Wasserstein distance

1. Introduction

In data analysis, it is natural to first understand and reduce the explanatory data before fitting a regression (see, e.g., Abdul-Wahab et al., 2005; Adraghi and Cook, 2009; Andras, 2017; Fradkin and Madigan, 2003; Hua et al., 2007). Aside from principal component analysis (PCA; see, e.g., Wold et al., 1987), a variety of different approaches, like UMAP (McInnes et al., 2018), MDS (Cox and Cox, 2008), autoencoders (Ng, 2011; Tolstikhin et al., 2018; Wang et al., 2016), SIR (Li, 1991), KPCA (Schölkopf et al., 1997), random projections (Bingham and Mannila, 2001), etc., have become increasingly popular lately. There is a similarly large variety of regression and function estimation methods (like neural networks, kernel methods, decision forests, etc.). This paper aims to shed some light on the

estimation errors that occur when first applying a dimensionality reduction method before fitting a regression. We focus on two frequently used and studied methods: PCA and kernel regression.

If X is the explanatory data and Y is an output variable in a standard L^2 -regression framework, the goal is to find a regression function g such that

$$\mathbb{E}[|g(X) - Y|^2]$$

is minimized. The optimal regression function g given perfect knowledge of the underlying distributions is given by $g(X) = \mathbb{E}[Y|X]$. When first applying a dimensionality reduction map P on X , the goal is instead to find a regression function f minimizing

$$\mathbb{E}[|f(P(X)) - Y|^2].$$

On the one hand reducing the dimension before performing a regression has several advantages, especially in view of computational aspects and interpretability of the occurring features and the obtained regression function (see, e.g., Chipman and Gu, 2005; Guillemard and Iske, 2017; Kim et al., 2005). An aspect that is of crucial importance in any regression framework is *sample efficiency*, i.e., how much sample data is required to obtain good approximations of the true regression function. Usually, sample efficiency tends to deteriorate in higher dimensions, which is an instance of the curse of dimensionality. Studying sample efficiency in combination with dimensionality reduction is a focus of the analysis in this paper.

On the other hand, dimensionality reduction decreases the predictive power of the explanatory data. Indeed, in terms of information content, it is clear that for any non-invertible transformation P , $P(X)$ contains at most as much information as X does about Y (Adragni and Cook, 2009). However, in many cases where dimensionality reduction is applied, the implicit assumption is that the reduced data $P(X)$ retains almost all relevant information encoded in X , and thus the fit $f(P(X)) \approx Y$ should be almost as good as $g(X) \approx Y$. For PCA, the structural error $f \circ P - g$ can be controlled via the eigenvalues of the covariance matrix of X , see Proposition 16.

In sample based regression, only finitely many i.i.d. copies of (X, Y)

$$(X_1, Y_1), \dots, (X_n, Y_n)$$

are given without knowing the true distribution of (X, Y) . For kernel regression, the convergence of the sample based regression functions for $n \rightarrow \infty$ is a well studied problem (see, e.g., Caponnetto and De Vito, 2007; Mendelson and Neeman, 2010; Steinwart et al., 2009; Wu et al., 2006). When including a prior dimensionality reduction step, however, new errors occur that need to be controlled.

Let P_n denote the estimator of the dimensionality reduction map using X_1, \dots, X_n and let f_n denote the estimated regression function such that $f_n(P_n(X_i)) \approx Y_i$ for $i = 1, \dots, n$. Comparing $f_n \circ P_n$ to $f \circ P$, we observe two different error sources: First, obviously both f_n and P_n include the statistical error introduced by sampling. Second however, a crucial error that is introduced for f_n is that already the input to its estimation procedure, $P_n(X_1), \dots, P_n(X_n)$, differs *in distribution* from the data $P(X)$ which is used to reconstruct

f . This means, the sample points $P_n(X_1), \dots, P_n(X_n)$ which are used to estimate f_n are not merely i.i.d. copies of $P(X)$ which is used to characterize f . Thus, the question of how close f_n is to f requires a suitable stability of kernel regression. In Theorem 7, we obtain a stability result which can control the difference between two kernel estimators when the difference of the underlying distributions is measured in the first Wasserstein distance. Subsequently, in Theorem 10, the stability result is used to obtain bounds on the joint procedure for general dimensionality reduction methods, which assumes that we can control both the operator norm $\|P_n - P\|_{\text{op}}$ and also the regression error for the regression on the low-dimensional space. These bounds are detailed out for the case of PCA in Theorem 12.

We emphasize that the subspaces spanned by estimated PCA and true PCA map may have entirely different support, and hence it is crucial that the stability result in Theorem 7 works with a metric distance like the Wasserstein distance, while stability results from the literature that work with the total variation norm or similar distances (see, e.g., Bousquet and Elisseeff, 2002; Christmann and Steinwart, 2007; Christmann et al., 2018; Xu et al., 2009) cannot be applied here. While we focus on PCA as a standard and well studied dimensionality reduction approach, Theorem 7 and the results in Section 3.1 allow also to study the impact of more advanced methods. These are especially necessary if the explanatory variables cannot be described by a linear subspace, but, for instance, by a submanifold.

Intuitively, one expects that first reducing the dimension of the data and then performing a regression on the low-dimensional data should be as sample efficient as starting with the low-dimensional data directly. As just discussed, this intuition ignores the added difficulty of the estimation error for the dimensionality reduction map, and in particular that this error is propagated by the stability of the low-dimensional regression. We obtain error estimates in Theorem 12 and Corollary 14 for the overall procedure of PCA and kernel regression where the errors coming from this source are indeed restrictive, and thus the overall error estimates might be worse than starting directly with low-dimensional data. Thus, the analysis in this paper suggests that two considerations for dimensionality reduction in terms of sample efficiency have to be taken into account: On the one hand, the actual regression on the low-dimensional space is usually faster (see Section 2.2 for a quantification thereof), but on the other hand, a new error through estimation of the dimension reduction map is introduced. This comparison is discussed in more detail in Section 3.3. We conclude that the two-step approach is particularly useful if the dimension of the input data is large and the kernel is not too regular. These theoretical findings are confirmed in a numerical illustration.

Another considerable benefit is in semi-supervised learning (see, e.g., Zhu and Goldberg, 2009; Zhu, 2005): If a larger sample for X than for (X, Y) pairs is given, one can use the whole X sample for the estimation of the dimensionality reduction map. If the sample for X (compared to (X, Y) sample pairs) is sufficiently large, Proposition 15 shows that the error arising from the estimation of the PCA map is no longer restrictive for the overall error of the PCA and kernel regression procedure, and thus the sample efficiency of the low-dimensional kernel regression is recovered.

We emphasize that the studied setting in this paper differs from the frequently studied question of *simultaneously* estimating a dimensionality reduction map P and an estimator f such that $f(P(X)) \approx Y$, which is for instance the case in multi-index models and related

fields (see, e.g., Alquier and Biau, 2013; DeVore et al., 2011; Gaïffas and Lecué, 2007; Gu, 2013; Lin and Zhang, 2006; Raskutti et al., 2012). The method of simply applying an off-the-shelf dimensionality reduction method like PCA certainly leads to worse performance on regression tasks, as the dimensionality reduction map is estimated using just the X -data, ignoring the output variable Y . Then again, such an approach has other benefits, like a clear interpretation of the reduced data and being able to use the same reduced data for different regression tasks.

Finally, we mention that related objectives to this paper have recently been studied from different perspectives, like dimensionality reduction for kernel regression based on partitioning the space (Hamm and Steinwart, 2021, 2022), or related stability questions by studying out-of-distribution performance of kernel methods (Canatar et al., 2021; Cui et al., 2021).

The remainder of the paper is structured as follows: We introduce the kernel regression setting in Section 2.1 and analyze the basic influence that dimension and dimensionality reduction methods have on sample efficiency in Section 2.2. The Wasserstein stability result for kernel regression is given in Section 2.3. Section 3 studies the combined procedure of dimensionality reduction and kernel regression. Section 3.1 focuses on general dimensionality reduction methods. Section 3.2 states basic results from the literature on PCA and Section 3.3 contains the results on the estimation error when PCA is combined with kernel regression. Section 4 gives numerical examples illustrating the procedure. All proofs are postponed to Section 5.

Throughout this paper, we denote by $|\cdot|$ the Euclidean norm. For a bounded linear operator P on a Hilbert space, we denote by $\|P\|_{\text{op}}$ the operator norm and $\|P\|_2$ the Hilbert-Schmidt norm, and similarly we denote by $\|P\|_{\text{op}}$ the operator norm for a map P mapping one normed space into another.

2. Kernel Regression: Dimension dependence, Convergence Rates, and Stability

This section gives results related to kernel regression. First, in Subsection 2.1 we introduce the kernel regression framework and state results from the literature on convergence rates. In Subsection 2.2, we give basic intuition and results on how kernel regression behaves under simple transformations of input measures, namely linear dimensionality reduction or inclusion of independent noise. And finally, Subsection 2.3 studies stability of the optimal kernel regression function when the input data is perturbed.

2.1 Kernel Regression Setting

We are given an input space $\mathcal{X} \subseteq \mathbb{R}^D$ and an output space $\mathcal{Y} \subseteq [-M, M]$ for some $M > 0$. The learning problem we are interested in is governed by a probability distribution ρ on $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, where the goal is to find the optimal predictor of the output variables, given the input variables. This means, the goal is to solve

$$f_\rho := \arg \min_f \int (f(x) - y)^2 \rho(dx, dy),$$

where the arg min is taken over all measurable functions mapping \mathcal{X} to \mathcal{Y} . Writing $\rho(dx, dy) = \rho_X(dx)\rho(dy|x)$, the solution is given by the conditional mean

$$f_\rho(x) = \int y \rho(dy|x) = \mathbb{E}[Y|X = x].$$

However, ρ is unknown, and only finitely many sample pairs

$$(X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n) \in \mathcal{Z}$$

are observed, which are independent and identically distributed according to ρ . To approximate f_ρ using these finitely many observations, we introduce the regularized kernel regression problem

$$f_{\lambda,n} := \arg \min_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n (f(X_i) - Y_i)^2 + \lambda \|f\|_{\mathcal{H}}^2,$$

where $1 \geq \lambda > 0$ is a regularization parameter and \mathcal{H} is a *reproducing kernel Hilbert space* (see, e.g., Cucker and Smale, 2002; Iske, 2018) of functions mapping \mathcal{X} to \mathbb{R} . We use the notation $K: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ for the kernel associated with \mathcal{H} , and always assume that the kernel is normalized to $\sup_{x \in \mathcal{X}} K(x, x) = 1$. We further introduce

$$f_{\lambda,\rho} := \arg \min_{f \in \mathcal{H}} \int (f(x) - y)^2 \rho(dx, dy) + \lambda \|f\|_{\mathcal{H}}^2, \tag{1}$$

$$= \arg \min_{f \in \mathcal{H}} \int (f(x) - f_\rho(x))^2 \rho_X(dx) + \lambda \|f\|_{\mathcal{H}}^2 \tag{2}$$

which is called the population counterpart to $f_{\lambda,n}$. If ρ_n denotes the empirical distribution of the given sample, then we have $f_{\lambda,\rho_n} = f_{\lambda,n}$. Defining $T_\rho: \mathcal{H} \rightarrow \mathcal{H}$ and $g_\rho \in \mathcal{H}$ by

$$\begin{aligned} T_\rho f(x) &:= \int K(x, u) f(u) \rho_X(du), \\ g_\rho(x) &:= \int K(x, u) y \rho(du, dy) \end{aligned}$$

we recall that the solution of (1) is given by (Caponnetto and De Vito, 2007)

$$f_{\lambda,\rho} = (T_\rho + \lambda I)^{-1} g_\rho.$$

Finally, we introduce the clipped estimator

$$\mathbf{f}_{\lambda,\rho} := -M \vee (M \wedge f_{\lambda,\rho}),$$

and similarly $\mathbf{f}_{\lambda,n}$. The purpose of clipping is that, since f_ρ only takes values in $[-M, M]$, it holds $|\mathbf{f}_{\lambda,\rho} - f_\rho| \leq |f_{\lambda,\rho} - f_\rho|$ and in certain aspects $\mathbf{f}_{\lambda,\rho}$ is more tractable than $f_{\lambda,\rho}$ (see also the discussion in Steinwart et al., 2009, and references therein).

We will require an analogue of T_ρ acting on L^2 . To this end, for a probability measure μ on \mathcal{X} , we define the kernel integral operator $L_{K,\mu}: L^2(\mu) \rightarrow L^2(\mu)$ by

$$L_{K,\mu} f(x) := \int K(x, x') f(x') \mu(dx'),$$

which is compact and self-adjoint, cf. Cucker and Smale (2002, Chapter III).

To study the approximation of $\|f_{\lambda,n} - f_\rho\|_{L^2(\rho_X)}$, two key parameters were identified in the literature to describe the convergence rate in n for a suitable choice of λ :

Assumption 1

- (i) Let $\alpha \in (0, 1)$ such that there exists a constant $C_\alpha > 0$ so that the sequence of non-increasing eigenvalues $\sigma_1 \geq \sigma_2 \geq \dots$ of L_{K,ρ_X} satisfies $\sigma_n \leq C_\alpha n^{-\frac{1}{\alpha}}$.
- (ii) Let $\beta \in (0, 1]$ such that there exists a constant $C > 0$ so that

$$\|f_{\lambda,\rho} - f_\rho\|_{L^2(\rho_X)}^2 + \lambda \|f_{\lambda,\rho}\|_{\mathcal{H}}^2 \leq C\lambda^\beta. \tag{3}$$

Roughly, the parameter α describes the complexity of the measure ρ_X in terms of the kernel K while β measures the approximation error of the function f_ρ with functions in \mathcal{H} . We briefly discuss α and β as follows.

Remark 1 Since $\int K(x, x)^2 \rho_X(dx) \leq 1$, L_{K,ρ_X} is compact and we have $\alpha \leq 1$ (Steinwart and Christmann, 2008, Theorem 4.27). If the kernel is m -times continuously differentiable and ρ_X is the uniform distribution on the Euclidean unit ball in \mathbb{R}^D , where $D < 2m$, then Assumption 1(i) is satisfied for $\alpha = \frac{D}{2m}$, see Steinwart et al. (2009, page 3). In particular, regular kernels lead to smaller α and a faster decay of (σ_n) . On the other hand, the decay suffers from large dimensions.

In view of (2), Assumption 1(ii) quantifies the approximation quality of f_ρ in a \mathcal{H} -ball. The function $A_2(\lambda) = \|f_{\lambda,\rho} - f_\rho\|_{L^2(\rho_X)}^2 + \lambda \|f_{\lambda,\rho}\|_{\mathcal{H}}^2$ occurring in Assumption 1(ii) is also called the approximation error function. It is directly related to the interpolation properties of the RKHS, see Steinwart and Christmann (2008, Chapter 5.6). For instance, if \mathcal{H} is a Sobolev space of order m , then the kernel is m -times differentiable and we can choose $\beta = \frac{2k}{2m}$ for k -Sobolev regular functions f_ρ where $k \leq m$ (cf. Wendland, 2004, Chapter 11). In particular, for a fixed regularity of the regression function f_ρ very regular kernels lead to small β .

We will require the following assumption for the kernel K .

Assumption 2 Assume the form $K(x_1, x_2) = \phi(|x_1 - x_2|)$ for kernel K , where $\phi: \mathbb{R}_+ \rightarrow \mathbb{R}$ satisfies the growth condition $\phi(0) - \phi(r) \leq \frac{L^2}{2} r^2$.

This assumption immediately implies the bound (cf. Iske, 2018, Section 8.4.2)

$$\|K(x_1, \cdot) - K(x_2, \cdot)\|_{\mathcal{H}} = (2(\phi(0) - \phi(|x_1 - x_2|)))^{1/2} \leq L|x_1 - x_2|. \tag{4}$$

In particular, all $f \in \mathcal{H}$ are bounded and Lipschitz continuous with $\|f\|_\infty \leq \|f\|_{\mathcal{H}}$ (recalling that $\phi(0) = K(x, x) = 1$) and $|f(x) - f(x')| = |\langle K(x, \cdot) - K(x', \cdot), f \rangle_{\mathcal{H}}| \leq L\|f\|_{\mathcal{H}}|x - x'|$. Throughout, we denote the Lipschitz norm of a Lipschitz continuous function $g: \mathcal{X} \rightarrow \mathcal{Y}$ by

$$\|g\|_{\text{Lip}} := \sup_{\substack{x_1, x_2 \in \mathcal{X}: \\ x_1 \neq x_2}} \frac{|g(x_1) - g(x_2)|}{|x_1 - x_2|}.$$

Assumption 2 is satisfied for frequently used choices like the Gaussian kernel $K(x, y) = \exp(-|x - y|^2)$ or various compactly supported radial kernels like $K(x, y) = (1 - |x - y|^2)_+^l$ for $l \in \mathbb{N}$ (Wu, 1995, Example 2), see also Zhu (2012, Table 4.1) or Wendland (2004) for more examples. In general, any function ϕ that is twice continuously differentiable, satisfies $\phi'(0) = 0$ and has bounded second derivative fulfills Assumption 2, as can be seen by using a second order Taylor expansion.

While we will use an error bound by Steinwart et al. (2009), in related settings very similar convergence results are obtained for instance by Caponnetto and De Vito (2007); Fischer and Steinwart (2020); Lin et al. (2020); Mendelson and Neeman (2010); Wendland (2004); Wu et al. (2006). The following result is a special case of Theorem 1 by Steinwart et al. (2009) with $s = p = \frac{1}{b}$ and $q = 2$.

Lemma 2 (Steinwart et al. (2009)) *There exists a constant $C_l > 0$, such that for $\tau > 0$, it holds*

$$\|\mathbf{f}_{\lambda,n} - f_\rho\|_{L^2(\rho_X)}^2 \leq \frac{C_l \tau}{2} \left(\lambda^\beta + \frac{1}{\lambda^{\alpha n}} + \frac{\lambda^\beta}{n} + \frac{1}{n} \right) \leq C_l \tau \left(\lambda^\beta + \frac{1}{\lambda^{\alpha n}} \right)$$

with probability at least $1 - 3 \exp(-\tau)$ with respect to the n -fold product measure $\rho^{\otimes n}$. The dominating terms on the right hand side of the inequality are λ^β and $(\lambda^{\alpha n})^{-1}$ for $n \rightarrow \infty$, $\lambda \rightarrow 0$. Optimizing for λ yields $\lambda_n := n^{-\frac{1}{\beta+\alpha}}$ and a resulting learning rate of

$$\|\mathbf{f}_{\lambda_n,n} - f_\rho\|_{L^2(\rho_X)}^2 \leq C_l \tau n^{-\frac{\beta}{\beta+\alpha}}$$

with probability at least $1 - 3 \exp(-\tau)$.

In the situation of Remark 1 with a Sobolev space \mathcal{H} of order m and a k -Sobolev functions f_ρ , we recover the classical minimax rate $n^{-2k/(2k+D)}$ if $2m > (2k) \vee D$. In particular, the rate deteriorates for large dimensions D .

2.2 Dimensionality Reduction and Transformations of Measures

In the previous section, the two parameters α and β were introduced which characterize the properties of \mathcal{H} in relation to the data distribution ρ necessary for the rate of convergence, see Lemma 2.

When transforming data, it is thus important to study how these parameters change when the data distribution changes. In particular, we focus on two simple yet important cases: First, a linear d -dimensional subspace of \mathbb{R}^D is transformed into a parametrization in \mathbb{R}^d , see Lemma 3. And second, independent noise is added to the input data, see Lemma 4.

For the following Lemma 3, the input data is changed with a linear map $A: \mathbb{R}^D \rightarrow \mathbb{R}^d$. We identify A by the corresponding matrix $A \in \mathbb{R}^{d \times D}$. We assume that A has orthonormal rows, where one may think of the PCA matrix containing the largest d eigenvectors of the covariance matrix of the data. Formally, the transformation of the input data is the following: If (X, Y) is distributed according to ρ , we set $\tilde{\rho}$ as the distribution of (AX, Y) , which is a probability measure on $\tilde{\mathcal{Z}} = \tilde{\mathcal{X}} \times \mathcal{Y}$, where $\tilde{\mathcal{X}} \subset \mathbb{R}^d$. We also assume that A is invertible on $\text{supp}(\rho_X)$, meaning that there is an inverse A_{inv} having orthonormal columns and $A_{\text{inv}}AX = X$. The interpretation of this assumption is that the data distribution ρ_X completely lies on a d -dimensional plane, see the left and middle images in Figure 2.2.




Figure 1: The figure on the left shows a two-dimensional uniform distribution. The figure in the middle illustrates a case where high-dimensional data lies on a lower-dimensional linear subspace and can be transformed with an orthonormal mapping as assumed in Lemma 3. The figure on the right shows a distribution which arises when independent noise is added to data lying on a linear subspace as assumed in Lemma 4.

Lemma 3 *Assume that the kernel K is of the form $K(x_1, x_2) = \phi(|x_1 - x_2|)$ for some $\phi: \mathbb{R} \rightarrow \mathbb{R}$. Then the learning problems for ρ and $\tilde{\rho}$ as defined above are equivalent in the following sense:*

- (i) $f_\rho = f_{\tilde{\rho}} \circ A$
- (ii) $f_{\lambda, \rho} = f_{\lambda, \tilde{\rho}} \circ A$
- (iii) *The parameters α and β from Assumption 1 can always be chosen equally for the learning problems with ρ and $\tilde{\rho}$.*

The result comes as no surprise, since both kernel matrices and output variables of the learning problems with ρ and $\tilde{\rho}$ can be transformed into each other by invertible maps. Nevertheless, the result is important in the sense that it perfectly fits the situation of PCA. For this setting, the result formalizes the intuition that it does not matter whether one uses the data as D -dimensional points lying on a linear plane or the respective d -dimensional parametrization thereof.

The next result emphasizes the potential benefit of going from a true high-dimensional space to a lower-dimensional one. At least in the extreme case, where all one does is filter out independent noise, this can only improve the eigenvalue decay speed of the kernel integral operator and thus improve on the parameter α . Formally, we choose some noise distribution κ on \mathcal{X} and denote by $\rho_X * \kappa$ the convolution of ρ_X with κ . This means, if $X \sim \rho_X$ and $\varepsilon \sim \kappa$ are independent, then $X + \varepsilon \sim \rho_X * \kappa$.

Lemma 4 *Grant Assumption 2. Assume ρ_X and κ have finite first moment. Denote the sequences of decreasing eigenvalues for L_{K, ρ_X} , $L_{K, \kappa}$ and $L_{K, \rho_X * \kappa}$ by $(\sigma_n)_{n \in \mathbb{N}}$, $(\gamma_n)_{n \in \mathbb{N}}$ and $(\hat{\sigma}_n)_{n \in \mathbb{N}}$, respectively. If $\sigma_n \sim n^{-a}$ and $\gamma_n \sim n^{-b}$ for $a, b > 0$. Then, if $\hat{\sigma}_n \lesssim n^{-c}$, we have $c \leq \min\{a, b\}$.*

Lemma 4 implies that the decay rate of eigenvalues for the parameter α in Assumption 1 is faster for ρ_X compared to $\rho_X * \kappa$, and hence when adding noise, the influence of α goes in the direction of a slower learning rate. In terms of α , filtering out noise can thus only improve the learning rate. In many situations, the improvement implied by Lemma 4 is strict as the following remark illustrates.

Remark 5 *Let X be supported on a d -dimensional subspace of \mathbb{R}^D and ε has D -dimensional support (e.g., κ is the uniform distribution on $[-\delta, \delta]^D$). For standard choices of kernels (cf. Remark 1, with $D < 2m$, where m is the regularity of the kernel), this leads to $\sigma_n \sim n^{-m/d}$ while $\gamma_n \sim n^{-m/D}$. Thus, by Lemma 4,*

$$\hat{\sigma}_n \gtrsim \max\{\sigma_n, \gamma_n\} = \gamma_n > \sigma_n$$

for n large enough. Comparing the case with noise ($X + \varepsilon$ with eigenvalues $\hat{\sigma}_n$) to the case without noise (X with eigenvalues σ_n), the choice of α in Assumption 1 (i) differs by a gap of at least $\frac{D-d}{m}$.

While Lemma 4 only deals with the parameter α , a classical model where the influence of the observation error $\varepsilon \sim \kappa$ on the regularity parameter β can be determined is the errors-in-variables model (see, e.g., Meister, 2009). The following example combines the dimension reduction perspective from above with these errors-in-variables models.

Example 1 *We consider the model*

$$Z = A^\top X + \varepsilon \in \mathcal{X} \quad \text{and} \quad Y = f(X) + \delta \in \mathcal{Y}$$

with $X \in \mathbb{R}^d$, centered observation errors $\varepsilon \in \mathcal{X} \subseteq \mathbb{R}^D$, $\delta \in \mathcal{Y} \subseteq [-M, M]$ where X, δ and ε are independent and with a matrix $A \in \mathbb{R}^{d \times D}$ such that $AA^\top = E_d$. An i.i.d. training sample (Z_i, Y_i) distributed as (Z, Y) is observed. Let ε admit a density $\kappa: \mathbb{R}^D \rightarrow \mathbb{R}_+$ and X has a density $\rho_X: \mathbb{R}^d \rightarrow \mathbb{R}_+$. In this setting the regression function f_ρ is given by

$$\begin{aligned} f_\rho(z) &= \mathbb{E}[Y|Z = z] = \mathbb{E}[Y|A^\top X + \varepsilon = z] = \int f(x) \mathbb{P}(X = dx | A^\top X + \varepsilon = z) \\ &= \frac{\int f(x) \rho_X(x) \kappa(z - A^\top x) dx}{\int \rho_X(x) \kappa(z - A^\top x) dx}. \end{aligned}$$

A regression with the projected data $AZ = X + A\varepsilon$ as discussed in the context of Lemma 3 leads to

$$f_{\tilde{\rho}}(z^*) = \frac{\int f(x) \rho_X(x) \kappa(A^\top(z^* - x)) dx}{\int \rho_X(x) \kappa(A^\top(z^* - x)) dx} = \frac{((f \cdot \rho_X) * \kappa(A^\top \cdot))(z^*)}{(\rho_X * \kappa(A^\top \cdot))(z^*)}$$

for any z^* in the range of A . In particular, $f_{\tilde{\rho}} = f_\rho(A^\top \cdot)$ inherits the regularity of $f\rho_X$ and κ thanks to the convolution structure. In contrast, f_ρ only achieves the regularity of κ in directions which are not aligned with A^\top such that $f_{\tilde{\rho}}$ might be much more regular than f_ρ , see Figure 2. This exemplifies a case where the regularity parameter β from Assumption 1 may improve through dimensionality reduction.

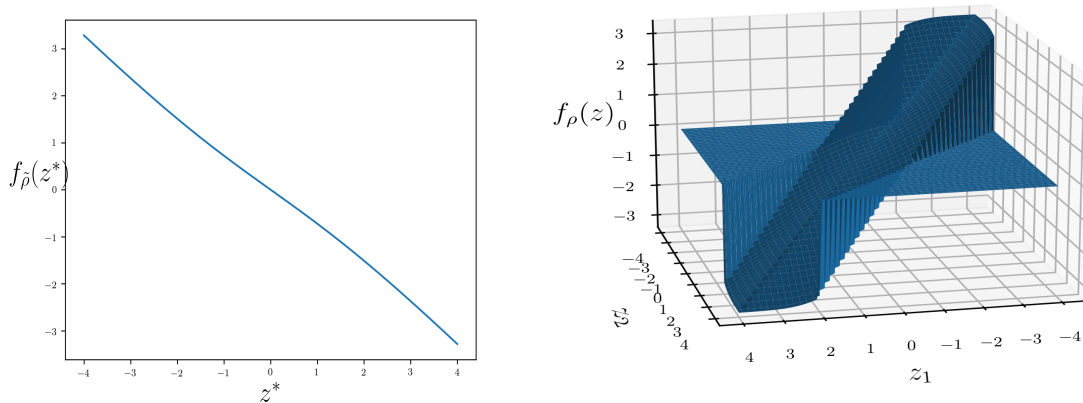


Figure 2: An illustration of the errors-in-variables model from Example 1 where the low-dimensional regression function (left) is much smoother than the high-dimensional regression function (right). The plotted functions arise in the errors-in-variables model by setting $X \sim \mathcal{N}(0, 1)$, $A = (1, 1)$, $\varepsilon \sim \mathcal{U}([-1, 1]^2)$ and $f(x) = x$.

2.3 Stability Result

This section presents a general stability result for kernel regression with respect to the Wasserstein distance. In Section 3, this result will be applied to the case where the error in the data arises from a dimension reduction preprocessing step. We first state the general result in Theorem 7, and discuss some of its aspects subsequently.

Definition 6 For two probability measures ρ_1 and ρ_2 on $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ we define the Wasserstein distance $W_1(\rho_1, \rho_2)$ between ρ_1 and ρ_2 by

$$W_1(\rho_1, \rho_2) := \inf_{\pi \in \Pi(\rho_1, \rho_2)} \int (|x_1 - x_2| + |y_1 - y_2|) \pi(d(x_1, y_1), d(x_2, y_2))$$

where $\Pi(\rho_1, \rho_2)$ is the set of all probability measures π on \mathcal{Z}^2 with first marginal distribution ρ_1 and second marginal distribution ρ_2 .

Theorem 7 Grant Assumption 2. If three distributions ρ_1, ρ_2, ρ_3 on \mathcal{Z} satisfy the relation

$$W_1(\rho_1, \rho_3) \leq \frac{\lambda}{4L} \tag{5}$$

for $\lambda > 0$, then it holds

$$\|f_{\lambda, \rho_1} - f_{\lambda, \rho_2}\|_{L^2(\rho_3)} \leq (\max\{1, LM\} + L\|f_{\lambda, \rho_2}\|_\infty + \|f_{\lambda, \rho_2}\|_{\text{Lip}}) \frac{W_1(\rho_1, \rho_2)}{\sqrt{\lambda}}.$$

The conclusion remains true if we replace Condition (5) by

$$\|(T_{\rho_3} - T_{\rho_1})(T_{\rho_3} + \lambda I)^{-1}\|_{\text{op}} \leq \frac{1}{2}. \tag{6}$$

In Theorem 7 the assumption on the relation between ρ_1 and ρ_3 is obviously satisfied for $\rho_1 = \rho_3$. Further, as will later be used, it is satisfied when ρ_1 is an empirical measure of ρ_3 of sufficiently large sample size, relative to λ . The assumption of Lipschitz continuity of the kernel is made so that it fits the standard definition of the Wasserstein distance. One may, for instance, weaken the assumption of Lipschitz continuity to Hölder continuity with exponent $a \in (0, 1)$ (that is, weaken the growth condition in Assumption 2 to $\phi(0) - \phi(r) \leq r^{2a}$) while simultaneously adjusting the cost function of the Wasserstein distance to $|x_1 - x_2|^a + |y_1 - y_2|$.

Remark 8 *Due to Assumption 2 the norms $\|f_{\lambda, \rho_2}\|_\infty$ and $\|f_{\lambda, \rho_2}\|_{\text{Lip}}$ on the right-hand side of the inequality in Theorem 7 can be bounded by $\|f_{\lambda, \rho_2}\|_{\mathcal{H}}$, which can be controlled in various ways. Notably, it always holds $\|f_{\lambda, \rho_2}\|_{\mathcal{H}} \leq \frac{M}{\sqrt{\lambda}}$ and more generally, stronger bounds on the \mathcal{H} -norm are for instance given in Wu et al. (2006, Lemma 5.2) and indirectly by Cucker and Smale (2002, Theorem 3). In particular, if $f_{\rho_2} \in \mathcal{H}$, then $\|f_{\lambda, \rho_2}\|_{\mathcal{H}}$ is uniformly bounded for all λ , see also Lemma 13 below.*

Example 2 *We exemplify Theorem 12 in the setting of Lemma 4. Say $(X, Y) \sim \rho$ and $\varepsilon \sim \kappa$, where (X, Y) and ε are independent. Assume*

$$\int |x| \kappa(dx) = \delta.$$

Define $(X + \varepsilon, Y) \sim \tilde{\rho}$. Then it is straightforward to see that $W_1(\rho, \tilde{\rho}) = \delta$ and thus for some constant B_λ only depending on λ but independent of δ , Theorem 7 yields that

$$\|f_{\lambda, \rho} - f_{\lambda, \tilde{\rho}}\|_{L^2(\rho)} \leq B_\lambda \delta.$$

As mentioned in Remark 8, both the Lipschitz constant and the L^∞ bound on $f_{\lambda, \tilde{\rho}}$ which are contained in the constant B_λ can be controlled independently of $\tilde{\rho}$ by $\|f_{\lambda, \tilde{\rho}}\|_\infty + \|f_{\lambda, \tilde{\rho}}\|_{\text{Lip}} \leq (1 + L)\|f_{\lambda, \tilde{\rho}}\|_{\mathcal{H}} \leq \frac{M}{\sqrt{\lambda}}$.

We discuss the order of magnitude of the relation implied by Theorem 7 for the difference $\|f_{\lambda, \rho_1} - f_{\lambda, \rho_2}\|_{L^2(\rho_3)}$ in terms of both $W_1(\rho_1, \rho_2)$ and λ in the following example.

Example 3 *Let $\rho_1 = \rho_3 = \frac{1}{2}\delta_{(0,0)} + \frac{1}{2}\delta_{(b,1)}$ and $\rho_2 = \frac{1}{2}\delta_{(a,0)} + \frac{1}{2}\delta_{(b+a,1)}$ for $a, b \in (0, 1)$ and $\lambda \in (0, 1]$. In particular, we have $W_1(\rho_1, \rho_2) = a$. By the representer theorem we have*

$$\begin{aligned} f_{\lambda, \rho_1}(x) &= \hat{w}_1 K(0, x) + \hat{w}_2 K(b, x) & \text{and} \\ f_{\lambda, \rho_2}(x) &= \hat{w}_1 K(a, x) + \hat{w}_2 K(b + a, x) \end{aligned}$$

where $\hat{w} = (\hat{w}_1, \hat{w}_2)^\top$ is given by

$$\hat{w} = \arg \min_{w \in \mathbb{R}^2} \frac{1}{2} \left| (0, 1)^\top - \mathbf{K}w \right|^2 + \lambda w^\top \mathbf{K}w, \quad \mathbf{K} = \begin{pmatrix} 1 & K(0, b) \\ K(0, b) & 1 \end{pmatrix}.$$

We choose b such that $K(0, b) = 1 - \lambda$. Due to Assumption 2, this implies that b is of order $\sqrt{\lambda}$. Basic linear algebra yields $\|f_{\lambda, \rho_1} - f_{\lambda, \rho_2}\|_{L^2(\rho_1)} \geq c \frac{ab}{\lambda}$ and $\|f_{\lambda, \rho_2}\|_{\mathcal{H}} \leq C \lambda^{-1/2}$ for some constants $C, c > 0$ which do not depend on λ and a . We conclude from Theorem 7 that

$$c \frac{a}{\sqrt{\lambda}} \leq \|f_{\lambda, \rho_1} - f_{\lambda, \rho_2}\|_{L^2(\rho_3)} \leq C' \frac{1 + \|f_{\lambda, \rho_2}\|_\infty + \|f_{\lambda, \rho_2}\|_{\text{Lip}}}{\sqrt{\lambda}} W(\rho_1, \rho_2) \leq C'' \frac{a}{\lambda}$$

with $C', C'' > 0$. We recover the linear dependence on $W_1(\rho_1, \rho_2)$. Although the order in λ is not sharp in this simple example, it already reveals that the distance $\|f_{\lambda, \rho_1} - f_{\lambda, \rho_2}\|_{L^2(\rho_3)}$ might indeed explode for $\lambda \rightarrow 0$.

3. Dimensionality Reduction

This section studies the combined approach of dimensionality reduction and function estimation. First, we derive an error estimate for the two-step procedure allowing for a general dimension reduction method. Subsequently, this result is detailed out for principal component analysis.

3.1 General Dimensionality Reduction Estimates

As before, the training data is given by pairs

$$(X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n) \in \mathcal{Z}$$

which are independent and identically distributed according to ρ . We first consider the general situation where we have some optimal dimensionality reduction map

$$P: \mathbb{R}^D \rightarrow \mathbb{R}^d$$

with $d \leq D$ and some estimator \hat{P}_n . Hereby, P is a dimensionality reduction for our explanatory data assuming full knowledge of the distribution of $X_1 \sim \rho_X$, while \hat{P}_n is an estimator depending on X_1, \dots, X_n . In this subsection, we will not specify P and \hat{P}_n any further, and derive general error estimates depending on the difference $\|\hat{P}_n - P\|_{\text{op}}$. In particular, the map P does not need be linear in this subsection. More specific estimates for the case of PCA will be derived in Sections 3.2 and 3.3. For further examples of dimensionality reduction methods we refer to Lee and Verleysen (2007), noting in particular that kernel PCA methods may yield suitable estimates as well, cf. Schölkopf et al. (1997) or Reiß and Wahl (2020).

We define $X_i^* := P(X_i)$ and the estimated counterpart as $\hat{X}_i := \hat{P}_n(X_i)$. We are looking for the optimal regression function given the dimension-reduced input data. Denoting the underlying distribution of the reduced data (X_1^*, Y_1) by $\tilde{\rho}$, the best predictor is given by

$$f_{\tilde{\rho}}(x) = \int y \tilde{\rho}(dy|x).$$

Since we lack knowledge of both ρ and P , we do not know $\tilde{\rho}$ either. To estimate $f_{\tilde{\rho}}$ with our finite data $(\hat{X}_1, Y_1), \dots, (\hat{X}_n, Y_n)$, we define the regularized least square kernel fit

$$\hat{f}_{\lambda, n} := \arg \min_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n (f(\hat{X}_i) - Y_i)^2 + \lambda \|f\|_{\mathcal{H}}^2.$$

For our error analysis we will later also require

$$f_{\lambda, n}^* := \arg \min_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n (f(X_i^*) - Y_i)^2 + \lambda \|f\|_{\mathcal{H}}^2.$$

As before, $\hat{\mathbf{f}}_{\lambda,n}$ and $\mathbf{f}_{\lambda,n}^*$ are the clipped versions.

The mappings \hat{P}_n and $\hat{f}_{\lambda,n}$ can be calculated given the sample data, while P and $f_{\hat{\rho}}$ are the respective best possible fits that can result from the employed procedure. Of main interest in our analysis is thus the estimation error

$$\|\hat{f}_{\lambda,n} \circ \hat{P}_n - f_{\hat{\rho}} \circ P\|_{L^2(\rho_X)}. \quad (7)$$

In addition there is the structural error arising from the dimensionality reduction, i.e., $\|f_{\hat{\rho}} \circ P - f_{\rho}\|_{L^2(\rho_X)}$. This error of course highly depends on the structure of the data X and the suitability of the applied dimensionality reduction P . A bound for the structural error in case of PCA depending on the linear structure of X will be given in Proposition 16.

First, we give the specification of Theorem 7 to the case relevant for bounding the error in Equation (7). To this end, Proposition 9 bounds the error between the kernel regression problems using differing input data determined by either P or \hat{P}_n . We recall that $\mathcal{Y} \subseteq [-M, M]$ and the constant L coming from Assumption 2.

Proposition 9 *Let $\eta \in (0, 1)$ and grant Assumption 2. For $\mathcal{N}(\lambda) := \text{Tr}((T_{\hat{\rho}} + \lambda)^{-1}T_{\hat{\rho}})$, assume that*

$$\lambda \leq \|T_{\hat{\rho}}\|_{\text{op}} \quad \text{and} \quad n \geq 64 \frac{\log^2(6\eta^{-1})\mathcal{N}(\lambda)}{\lambda} \quad (8)$$

and set $S_n := \frac{1}{n} \sum_{i=1}^n |X_i|$. Then

$$\|\hat{f}_{\lambda,n} \circ P - \mathbf{f}_{\lambda,n}^* \circ P\|_{L^2(\rho_X)} \leq \frac{S_n(1 + LM + \|\hat{f}_{\lambda,n}\|_{\text{Lip}} + \|\hat{f}_{\lambda,n}\|_{\infty})}{\sqrt{\lambda}} \|\hat{P}_n - P\|_{\text{op}}$$

holds with probability at least $1 - \eta/3$ with respect to $\rho^{\otimes n}$.

In the above result, the term S_n is of course either bounded (if X_1 is bounded) or at least bounded with high-probability by the law of large numbers, since we will always assume that X_1 has a finite second moment. Further, while $\|\hat{f}_{\lambda,n}\|_{\text{Lip}} + \|\hat{f}_{\lambda,n}\|_{\infty}$ can be unbounded in general, in various situations one may reasonably assume that this term is (uniformly across n) bounded with high probability. One case may be if $\hat{f}_{\lambda,n}$ converges in a sufficiently strong sense to a bounded and Lipschitz continuous function. Another case may be if all estimators are uniformly bounded with respect to $\|\cdot\|_{\mathcal{H}}$ (cf. Remark 8). Finally, it is worth mentioning that the assumption in (8) may be restrictive in general, although in the setting of Caponnetto and De Vito (2007) the assumption is shown to be satisfied asymptotically for a suitable choice of λ leading to optimal rates (cf. the related discussion in Wendland, 2004, Chapter 13).

Below, we state the main result of this section, which shows how to reduce the error of the joint procedure (estimation of both dimensionality reduction and kernel regression) to each individual procedure, where only the results for the kernel regression on the low-dimensional explanatory data (given by the push-forward distribution $\rho_X \circ P^{-1}$) is needed.

Theorem 10 *Let $\lambda \in (0, 1]$ and $\eta \in (0, 1)$ such that (8) holds. Define $S^2 = \mathbb{E}[|X_1|^2]$ and assume $S < \infty$. Then*

$$\begin{aligned} & \frac{1}{3} \|\hat{\mathbf{f}}_{\lambda,n} \circ \hat{P}_n - f_{\hat{\rho}} \circ P\|_{L^2(\rho_X)}^2 \\ & \leq \frac{(S_n^2 + S^2)(1 + LM + \|\hat{f}_{\lambda,n}\|_{\text{Lip}} + \|\hat{f}_{\lambda,n}\|_{\infty})^2}{\lambda} \|\hat{P}_n - P\|_{\text{op}}^2 + \|\mathbf{f}_{\lambda,n}^* - f_{\hat{\rho}}\|_{L^2(\rho_X \circ P^{-1})}^2 \end{aligned}$$

holds with probability at least $1 - \eta/3$ with respect to $\rho^{\otimes n}$.

To prove this theorem, we decompose

$$\begin{aligned} \frac{1}{3} \|\hat{\mathbf{f}}_{\lambda,n} \circ \hat{P}_n - f_{\hat{\rho}} \circ P\|_{L^2(\rho_X)}^2 &\leq \|(\hat{f}_{\lambda,n} \circ \hat{P}_n - \hat{f}_{\lambda,n} \circ P)\|_{L^2(\rho_X)}^2 \\ &\quad + \|(\hat{f}_{\lambda,n} \circ P - f_{\lambda,n}^* \circ P)\|_{L^2(\rho_X)}^2 \\ &\quad + \|(f_{\lambda,n}^* \circ P - f_{\hat{\rho}} \circ P)\|_{L^2(\rho_X)}^2, \end{aligned} \quad (9)$$

where the first term can be estimated with the Lipschitz constant of $\hat{f}_{\lambda,n}$ combined with the reconstruction error $\|\hat{P}_n - P\|_{\text{op}}$ and the second term is bounded by Proposition 9.

Theorem 10 implies that as soon as we can control $\|\hat{P}_n - P\|_{\text{op}}$, we can reduce the estimation error $\|\hat{\mathbf{f}}_{\lambda,n} \circ \hat{P}_n - f_{\hat{\rho}} \circ P\|_{L^2(\rho_X)}$ of the combined learning problem to the estimation error which would result if we just had to estimate $f_{\hat{\rho}}$ while assuming perfect knowledge of P , i.e., to $\|f_{\lambda,n}^* - f_{\hat{\rho}}\|_{L^2(\rho_X \circ P^{-1})}$.

3.2 Principal Component Analysis

Next, we state results on PCA, which are mainly taken from Reiß and Wahl (2020). For a more general treatment of PCA in the context of dimensionality reduction we refer for instance to Jolliffe (2002); Lee and Verleysen (2007). The results by Reiß and Wahl (2020) will later be used in our estimates for the combined study of PCA with regularized kernel regression. The goal is to reduce the dimension of the input data from dimension D to $d < D$. The reduced data will still be regarded as points in \mathbb{R}^D , albeit Lemma 3 shows that it does not matter whether one uses the points as a linear subset in \mathbb{R}^D or a d -dimensional representation thereof.

For the PCA procedure, we are in the same setting as introduced in Section 2.1. Let

$$\mathcal{P}_d := \{P \in \mathbb{R}^{D \times D} : P \text{ is orthogonal projection of rank } d\}. \quad (10)$$

This means $P \in \mathcal{P}_d$ can be written as $P = A^\top A$ for $A \in \mathbb{R}^{d \times D}$ having orthonormal rows. Define

$$P := \arg \min_{\tilde{P} \in \mathcal{P}_d} \int |x - \tilde{P}x|^2 \rho_X(dx) \quad \text{and} \quad \hat{P}_n := \arg \min_{\tilde{P} \in \mathcal{P}_d} \frac{1}{n} \sum_{i=1}^n |X_i - \tilde{P}X_i|^2.$$

The following assumption will be used to apply the results on principal component analysis. Note that part (i) implies that X has finite second moment.

Assumption 3 *Let $X \sim \rho_X$.*

- (i) X is sub-Gaussian, i.e., $\sup_{u \in \mathbb{R}^D: \mathbb{E}[\langle X, u \rangle^2] \leq 1} \sup_{k \geq 1} k^{-1/2} \mathbb{E}[|\langle X, u \rangle|^k]^{1/k} < \infty$.
- (ii) $\sigma_d^X - \sigma_{d+1}^X > 0$, where $\sigma_1^X, \sigma_2^X, \dots, \sigma_D^X$ is the sequence of decreasing eigenvalues of the covariance matrix of X .
- (iii) X is centered: $\mathbb{E}[X] = 0$.

The following is a corollary of Reiß and Wahl (2020, Proposition 2) and the main reconstruction result for PCA which will be used in Section 3.3.

Lemma 11 (Reiß and Wahl (2020)) *Grant Assumption 3. There exists a constant C_{pca} such that for all $n \in \mathbb{N}$, $\eta \in (0, 1)$, with probability $(1 - \eta)$, it holds*

$$\|\hat{P}_n - P\|_{\text{op}} \leq \|\hat{P}_n - P\|_2 \leq \frac{\Gamma}{3\eta\sqrt{n}} \quad \text{with} \quad \Gamma := \frac{3C_{pca}}{(\sigma_{d+1}^X - \sigma_d^X)}.$$

The PCA projections we work with defined in Equation (10) are regarded as mappings from \mathbb{R}^D onto itself. In terms of matrices, this means the projection mapping $P = A^\top A$ is used to get the reduced data $PX \in \mathbb{R}^D$ lying on a d -dimensional plane. An alternative route would be to work with $A \in \mathbb{R}^{d \times D}$ directly instead and obtain reduced data $AX \in \mathbb{R}^d$. Similarly to Lemma 11 one can obtain error estimates in this setting using the sin θ -Theorem (see, e.g., Yu et al., 2015) combined with optimal rates on covariance matrix estimation (see, e.g., Cai et al., 2010).

3.3 Estimation Error for PCA and Kernel Regression

We can now bound the statistical error from (7) for PCA and kernel regression. Recall the constants C_l and Γ from Lemma 2 and Lemma 11, respectively. The parameter β is as in Assumption 1 for the measure $\tilde{\rho}$ instead of ρ . As before, $S_n := \frac{1}{n} \sum_{i=1}^n |X_i|$ and $S^2 = \mathbb{E}[|X_1|^2]$, and we assume $S < \infty$ throughout.

Theorem 12 *Grant Assumptions 2 and 3 and let $\tilde{\rho}$ satisfy (3) for $\beta \in (0, 1]$. Let $\lambda \in (0, 1]$ and $\eta \in (0, 1)$ such that (8) holds. Then there is a constant $\bar{C} > 0$ depending on M, L and Γ such that*

$$\begin{aligned} \|\hat{f}_{\lambda,n} \circ \hat{P}_n - f_{\tilde{\rho}} \circ P\|_{L^2(\rho_X)}^2 &\leq \bar{C} \frac{(S^2 + S_n^2)(\|\hat{f}_{\lambda,n}\|_{\text{Lip}}^2 + \|\hat{f}_{\lambda,n}\|_\infty^2 + 1)}{\eta^2 n \lambda} \\ &\quad + 3C_l \log\left(\frac{9}{\eta}\right) \left(\lambda^\beta + \frac{1}{\lambda n}\right). \end{aligned} \quad (11)$$

with probability at least $1 - \eta$ according to the n -fold product measure $\rho^{\otimes n}$.

To show this result, we apply Theorem 10 in combination with the PCA error bound from Lemma 11 and the kernel regression error bound on the low-dimensional space from Lemma 2.

In the following, we specify the above Theorem 12 by optimizing for λ under the simplifying assumption that $\|\hat{f}_{\lambda,n}\|_{\text{Lip}} + \|\hat{f}_{\lambda,n}\|_\infty$ is bounded. Since both norms are bounded by $\|\hat{f}_{\lambda,n}\|_{\mathcal{H}}$, the following lemma provides a high probability bound in the regular case $\beta = 1$:

Lemma 13 *If $f_{\tilde{\rho}} \in \mathcal{H}$ and $\lambda = n^{-1/2}$, then there is a constant $B > 0$ such that*

$$\|\hat{f}_{\lambda,n}\|_{\mathcal{H}}^2 \leq B\eta^{-3/2} + 4\|f_{\tilde{\rho}}\|_{\mathcal{H}}^2$$

is satisfied with probability of at least $1 - 5\eta$.

On the event that $\|\hat{f}_{\lambda,n}\|_{\text{Lip}} + \|\hat{f}_{\lambda,n}\|_{\infty}$ is bounded, the dominating terms in the bound of Theorem 12 for $n \rightarrow \infty$ and $\lambda \rightarrow 0$ on the right hand side of Equation (11) are $\frac{1}{n\lambda}$ and λ^{β} . Note that the resulting optimal λ satisfies (8) as soon as $\beta > \alpha$.

Corollary 14 *We are in the setting of Theorem 12. Optimizing the right hand side of Equation (11) for λ yields*

$$\lambda = n^{-\frac{1}{\beta+1}}.$$

Assuming (8), there are high-probability events $\mathcal{A}_n \subseteq \Omega, n \in \mathbb{N}$, with $\mathbb{P}(\mathcal{A}_n) \geq 1 - \eta_n$ such that we have an overall learning rate of the form

$$\begin{aligned} \|\hat{\mathbf{f}}_{\lambda,n} \circ \hat{P}_n - f_{\bar{\rho}} \circ P\|_{L^2(\rho_X)}^2 &\leq \frac{C}{\eta_n^2} n^{-\frac{\beta}{\beta+1}} \\ &\text{on } \mathcal{A}_n \cap \{\|\hat{f}_{\lambda,n}\|_{\text{Lip}} + \|\hat{f}_{\lambda,n}\|_{\infty} \leq \mathcal{L}\} \end{aligned} \quad (12)$$

with some constant $C > 0$ (which depends on \mathcal{L}).

Instead of the procedure introduced in this section (first PCA, then kernel regression on d -dimensional subspace), consider the alternative method to fit f_{ρ} directly using kernel regression on \mathbb{R}^D . Let us denote the clipped estimated kernel functions by $\mathbf{g}_{\lambda,n}$ and $\bar{\alpha}, \bar{\beta}$ be the parameters for Assumption 1 for ρ . By Lemma 2, the resulting learning rate is

$$\|f_{\rho} - \mathbf{g}_{\lambda,n}\|_{L^2(\rho_X)}^2 \leq C n^{-\frac{\bar{\beta}}{\bar{\beta}+\alpha}}.$$

Compared with the rate in (12), we find that it is not immediate which procedure converges faster. Indeed, only if

$$\frac{\bar{\beta}}{\bar{\alpha}} < \beta,$$

then the two-step approach converges faster than a direct high-dimensional estimation. Note that it is natural to assume $\beta \geq \bar{\beta}$, i.e. the lower dimensional regression function is at least as regular as the high-dimensional regression function. In view of Example 1 we may indeed gain in the regularity, i.e. $\beta > \bar{\beta}$.

In Corollary 14, the influence of α from Assumption 1 is lost in the overall rate. More precisely, the rate behaves as if α takes the (worst possible) value of $\alpha = 1$. In a sense, this means that the rate is no longer adaptive with respect to the complexity of ρ_X measured by the RKHS \mathcal{H} . This makes intuitive sense, since now the complexity of ρ_X also influences the estimation procedure of the PCA map, which is independent of \mathcal{H} .

On the other hand, in the setting of Remark 1 but for the low dimensional space, we can adapt the regularity m of the kernel to the dimension d which has two advantages: First, we can allow for much smaller regularities since d might be considerably smaller than D and thus the numerical stability of the kernel method can be improved (cf. Wendland, 2004, Section 12). Second, the resulting $\alpha = \frac{d}{2m}$ can be chosen close to one and the rate in Corollary 14 is only slightly worse than the best possible rate $n^{-2k/(2k+d)}$ for k -Sobolev regular regression functions $f_{\bar{\rho}}$.

Moreover, this comparison certainly motivates that dimensionality reduction is particularly suitable to apply if one does not expect small values of $\bar{\alpha}$ anyways. Recall that $\bar{\alpha}$ is

usually small if the kernel K is very smooth and the dimension is not too large. Especially for large dimensions, one could expect high values of $\bar{\alpha}$, and thus not much is lost when applying dimensionality reduction, which basically leads to $\alpha = 1$. This intuition is consistent with the numerical results in Section 4, where Figure 4 shows that utilizing PCA for dimensionality reduction leads to notable improvements for C^0 or C^2 kernels, while this is not (or less so) the case for the Gaussian kernel, which is C^∞ .

The two-step procedure has another interesting advantage. Consider the case where more X -sample points are given compared to (X, Y) sample pairs, which is a situation frequently observed in semi-supervised learning settings. Then, the error bounds for the PCA map can use all X samples, even though the function estimation can only use the pairs. This can improve the overall estimates of the whole procedure.

Let $m > n$ and additional $m - n$ many X -samples X_{n+1}, \dots, X_m be given and we consider the case where the PCA map is estimated using all m many X -samples.

Theorem 15 *With the setting and notation of Theorem 12, granting additionally that Assumption 1 holds for $\tilde{\rho}$, we have with probability $(1 - \eta)$ with respect to $\rho^{\otimes m}$ that*

$$\begin{aligned} & \|\hat{f}_{\lambda, n} \circ \hat{P}_m - f_{\tilde{\rho}} \circ P\|_{L^2(\rho_X)}^2 \\ & \leq \bar{C} \frac{(1 + S_n^2)(\|\hat{f}_{\lambda, n}\|_{\text{Lip}}^2 + \|\hat{f}_{\lambda, n}\|_\infty^2)}{\eta^2 m \lambda} + C_l \log(9/\eta) \left(\lambda^\beta + \frac{1}{\lambda^{\alpha n}} \right). \end{aligned} \quad (13)$$

We have $\|\hat{f}_{\lambda, n}\|_{\text{Lip}}^2 + \|\hat{f}_{\lambda, n}\|_\infty^2 \leq (1 + L^2)M^2/\lambda$. Assuming (8), the right hand side of (13) is of order $\frac{1}{\lambda^2 m} + \lambda^\beta + \frac{1}{\lambda^{\alpha n}}$ for large n, m and small λ . If $m > n^{\frac{2+\beta}{\alpha+\beta}}$ the optimized choice $\lambda = n^{-\frac{1}{\alpha+\beta}}$ yields the rate of convergence

$$\|\hat{f}_{\lambda, n} \circ \hat{P}_m - f_{\tilde{\rho}} \circ P\|_{L^2(\rho_X)}^2 \leq C n^{-\frac{\beta}{\beta+\alpha}}.$$

Otherwise we obtain the rate $m^{-\frac{\beta}{\beta+2}}$ for $\lambda = m^{-\frac{1}{2+\beta}}$.

For sufficiently large m , we thus recover the rate from Lemma 2 in the low-dimensional space.

To complete the picture, Proposition 16 gives a direct result on controlling the structural error $\|f_\rho - f_{\tilde{\rho}} \circ P\|_{L^2(\rho_X)}$ of the procedure under a Lipschitz condition for f_ρ . It shows that the magnitude of the error is controlled by the smallest eigenvalues of the covariance matrix of $X \sim \rho_X$.

Proposition 16 *Let f_ρ be Lipschitz continuous with constant L . Then it holds*

$$\|f_\rho - f_{\tilde{\rho}} \circ P\|_{L^2(\rho_X)} \leq 2L \left(\sum_{i=d+1}^D \sigma_i^X \right)^{1/2},$$

where $\sigma_1^X, \dots, \sigma_D^X$ are the eigenvalues of the covariance matrix of $X \sim \rho_X$ in decreasing order.




Figure 3: Rate for the excess reconstruction error (the error occurring from sample estimation of the PCA map) for the PCA map in the Example given in Section 4. To estimate the expectation, we used the Monte-Carlo method with 10000 sample points for \hat{P}_n for each n .

4. Numerical Examples

In this section, the overall procedure of dimensionality reduction and kernel regression is illustrated in a simple setting. We mainly aim to shed some light on the comparison of the occurring errors for direct regression versus inclusion of dimensionality reduction, as discussed above. Due to computational constraints, the insight of these examples towards true asymptotic behavior is of course limited, but we believe the illustration for smaller sample sizes and visualization of the absolute errors involved can nevertheless be insightful.

Let $d = 2$ and $D = 10$. The X -data is generated as follows: Consider \tilde{X} to be uniformly distributed on $[-1, 1]^d \times [-\varepsilon, \varepsilon]^{D-d}$, where we set $\varepsilon = 0.1$. We define X to be some rotation (with an orthogonal transformation) of \tilde{X} , say $X = A\tilde{X}$. The first two eigenvalues of the covariance matrix of X are thus $1/3$ and the remaining ones are $1/30$.

As in Section 3.2, we denote by P the PCA map and \hat{P}_n the estimated counterpart. The excess reconstruction error for the PCA procedure depending on the number of sample points used, which is studied in Lemma 11, is showcased in Figure 3. The shown error, after an application of Markov's inequality, yields Lemma 11 which governs the term ε in the stability result in Proposition 9. We see that the (predicted) linear behavior in n arises.

For the Y -data, we treat two different cases. Roughly speaking, the cases differ in the sense of whether the dependence structure of (X, Y) includes the noise inherent in X or not. For both cases, we fix a function $f^{(1)}: \mathbb{R}^D \rightarrow \mathbb{R}$, which we (rather arbitrarily) define by $f^{(1)}(x) := \sin(\sum_{i=1}^D x_i)$.

For the first case, we define $Y := f^{(1)}(P(X)) + U$, and for the second case $Y := f^{(1)}(X) + U$, where U is uniformly distributed on $[-0.1, 0.1]$ and independent of all other variables. We see that in the first case, the dependence structure between X and Y is completely mediated by the first two principal components of X , while in the second case




Figure 4: Overall estimation error $\int (\hat{f}(x) - y)^2 \rho(dx, dy)$ for different estimators \hat{f} and different generation of Y data. The different estimators are the asymptotic estimates f_ρ and $f_\rho \circ P$, the direct kernel estimator $g_{\lambda,n}$, and the estimator arising from the PCA and subsequent kernel regression procedure $\hat{f}_{\lambda,n} \circ \hat{P}_n$. The respective kernels are described in (14) and λ is set via cross-validation. Each reported value is an average over 50 independent runs of generating the respective sample.

all principal components of X are relevant. We thus expect the dimensionality reduction method to be particularly useful in the first case.

We use cross-validation for the choice of λ as described in Steinwart et al. (2009), which is therein shown to be optimal asymptotically and thus leads to a regime of λ where the error bound from Corollary 14 applies. The results are presented for different kernel functions varying in their regularity. Each kernel is of the form $K(x, y) = \phi(|x - y|)$, where we use

$$\begin{aligned} \phi_\infty(r) &= \exp(-r^2) && (C^\infty, \text{Gaussian kernel}), \\ \phi_2(r) &= \max\{0, 1 - r\}^8 (8r + 1) && (C^2, \text{cf. Zhu, 2012, Table 4.1}), \\ \phi_0(r) &= \max\{0, 1 - r\}^6 (8r + 1) && (C^0, \text{cf. Zhu, 2012, Table 4.1}). \end{aligned} \tag{14}$$

We observe the estimation errors for the first case and different choices of kernel functions in Figure 4. As expected by the data generating process for the first case, the procedure including the dimensionality reduction step performs very well. However, the dimensionality reduction also mostly reduces the total error in the second case (for the C^0 and C^2 kernels), while only in the Gaussian case (where the absolute error is already very small), the asymptotic error dominates and the procedure involving dimensionality reduction produces larger absolute error, see also the discussion after Corollary 14, which gives an explanation for this behavior.¹

Since the function $f^{(1)}(x) := \sin(\sum_{i=1}^D x_i)$ defining the dependence structure of (X, Y) is C^∞ , it is natural that the Gaussian kernel performs well. It is interesting to observe what happens when a less regular function is used. To this end, we define $f^{(2)}(x) := |\sin(2 \sum_{i=1}^D x_i)|$ and (X, Y) for the two different cases analogously to the above. The top two graphs of Figure 5 report the results. Obviously, all occurring errors are larger and the observed rate of convergence is slower compared to the data generated via $f^{(1)}$, compatible with Lemma 2 and Corollary 14. Since the Gaussian kernel does not quickly convergence to the best estimator, the effect of the dimensionality reduction is visible even for Case 2. A plausible reason is that the irregular relation of (X, Y) is smoothed via the dimensionality reduction, and thus the parameter β of Assumption 1 is larger (leading to faster convergence) for the two-step procedure (c.f. the discussion after Corollary 14). In Figure 5, we further showcase the behavior that occurs when we use cross-validation for the bandwidth parameter for the Gaussian kernel (the two graphs in the middle) and when generating the X -data using a nonlinear transformation (the bottom two graphs in Figure 5). While cross-validation of the bandwidth is an important practical tool, in the literature on learning rates of kernel regression, one usually fixes the kernel, which is also the case in this paper. Nevertheless, we see in Figure 5 that the observed patterns still appear similar to the previous cases even if cross-validation for the bandwidth is used. Finally, as expected, when the X -data was generated via a nonlinear transformation, the method using PCA performs generally worse compared to previous cases. Overall, we believe the study of nonlinear manifolds using also nonlinear dimensionality reduction methods is an interesting avenue for further research.

To emphasize the effect of estimating the dimensionality reduction map instead of starting in the low-dimensional space directly, Figure 6 shows the difference between the estimator $\hat{f}_{\lambda,n}$ which uses the estimated PCA map \hat{P}_n and the estimator $f_{\lambda,n}^*$ using the true

1. We note another possible reason for the difference in absolute errors between the choices of kernel can be the bandwidth. This means, the absolute errors may be more similar if we tried to optimize the bandwidth suitably (i.e., work with kernels $r \mapsto \phi(\gamma r)$ for suitable choices of $\gamma > 0$).

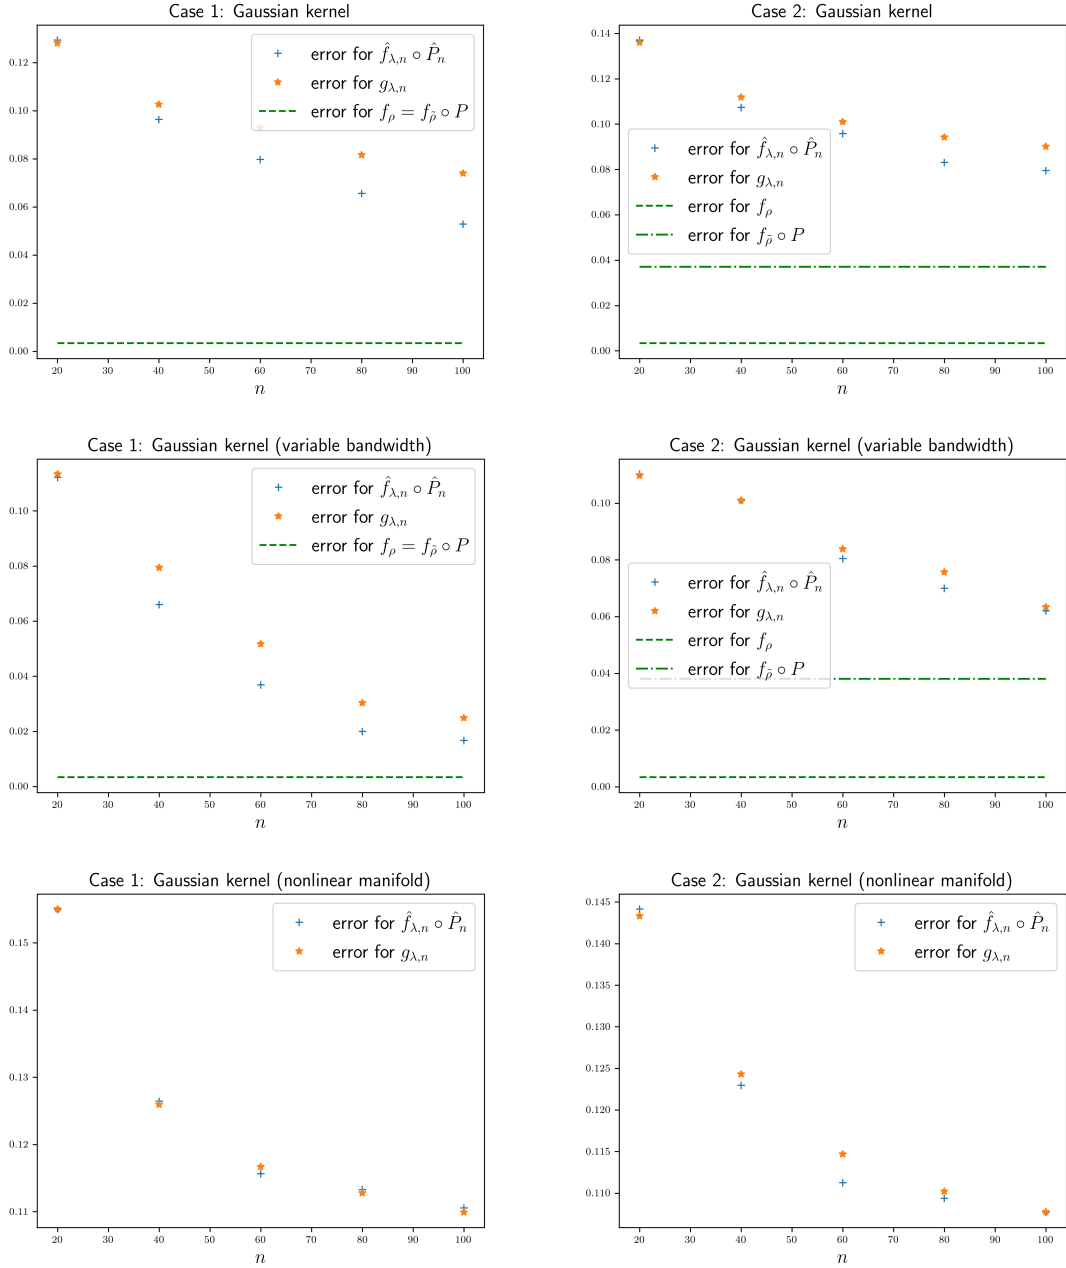


Figure 5: Overall estimation error $\int (\hat{f}(x) - y)^2 \rho(dx, dy)$ for different estimators \hat{f} where the data is determined via the non-differentiable function $f^{(2)}$. Each reported value is an average over 50 independent runs of generating the respective sample. The different estimators are the asymptotic estimates f_ρ and $f_{\hat{\rho}} \circ P$, the direct kernel estimator $g_{\lambda,n}$, and the estimator arising from the PCA and subsequent kernel regression procedure $\hat{f}_{\lambda,n} \circ \hat{P}_n$. *First line:* Gaussian kernel $\phi(\lambda r) = \exp(-hr^2)$ with fixed bandwidth $h = 1$ and cross-validated λ . *Second line:* The bandwidth $h \in [10^{-3}, 10]$ is chosen with cross-validation, too. *Third line:* Features are generated via $X = A\tilde{X} + 0.25 \sin(\sum_{i=1}^{10} \tilde{X}_i)$ (instead of $X = A\tilde{X}$).

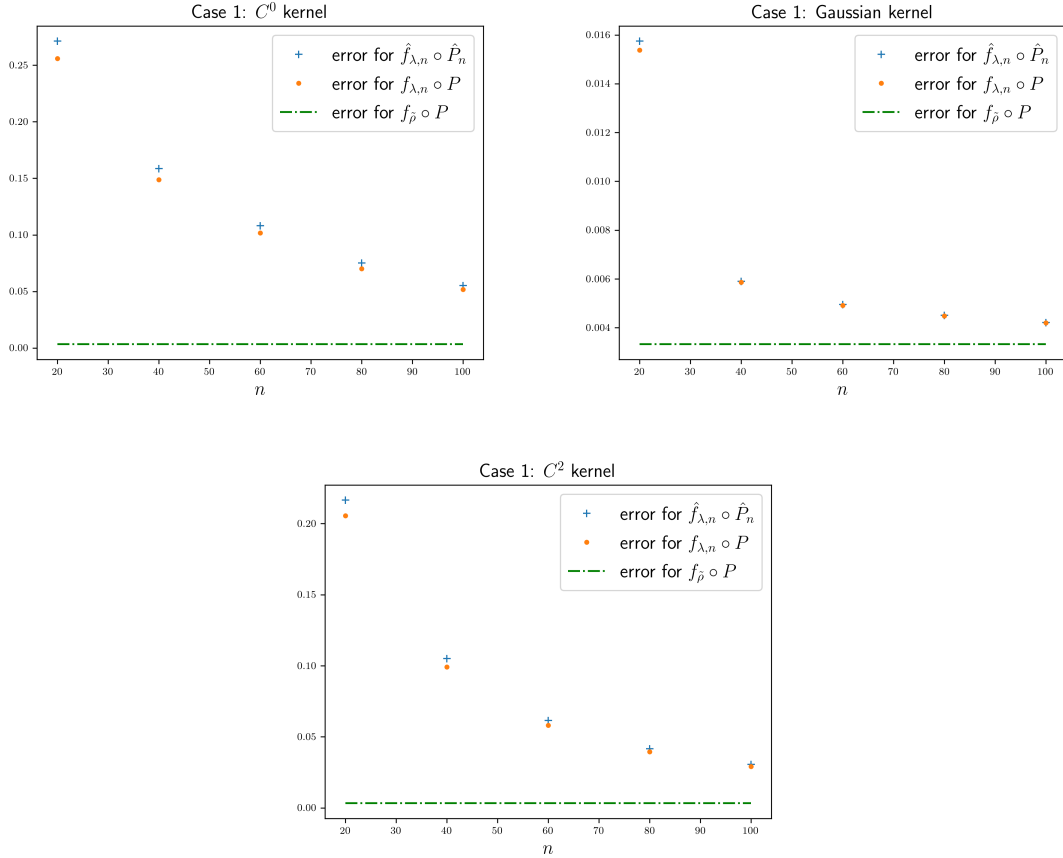


Figure 6: Overall estimation error $\int (f(x) - y)^2 \rho(dx, dy)$ for different estimators f . The different estimators are the asymptotic estimator $f_{\hat{\rho}} \circ P$, the estimator arising from the PCA and subsequent kernel regression procedure $\hat{f}_{\lambda,n} \circ \hat{P}_n$, as well as the estimator $f_{\lambda,n}^* \circ P$ which uses samples but assumes knowledge of the true PCA map P . The respective kernels are described in (14) and λ is set via cross-validation. Each reported value is an average over 50 independent runs of generating the respective sample.

PCA map P (for the first case described above, where $Y = f(P(X)) + U$). While we see that the estimation of the PCA map does influence the total error in a non-negligible way, it is not clear whether this influence is really relevant asymptotically. This suggests that the estimates of Theorems 7 and Proposition 9 may be improved upon for situations as the one treated in this example, perhaps under the assumption of higher order differentiability of the kernel function, instead of just the Lipschitz condition required in Assumption 2.

5. Proofs

5.1 Proofs for Section 2

Proof of Lemma 3 Before showing the individual claims, we first make the following observation: Take $\lambda > 0$ and $\varphi: \mathcal{X} \rightarrow \mathbb{R}$. Then (λ, φ) is an eigenvalue/-function pair (λ, φ) of L_{K, ρ_X} if and only if $(\lambda, \varphi \circ A_{\text{inv}})$ is one of $L_{K, \tilde{\rho}_{\tilde{X}}}$. This follows because A_{inv} has orthonormal columns and thus $|\tilde{x}| = |A_{\text{inv}}(\tilde{x})|$ for all $\tilde{x} \in \tilde{\mathcal{X}}$. This implies

$$\begin{aligned} L_{K, \rho_X} \varphi(x) &= \int \phi(|x - u|) \varphi(u) \rho_X(du) \\ &= \int \phi(|A_{\text{inv}}(A(x)) - \tilde{u}|) \varphi \circ A_{\text{inv}}(\tilde{u}) \tilde{\rho}_{\tilde{X}}(d\tilde{u}) \\ &= \int \phi(|A(x) - \tilde{u}|) \varphi \circ A_{\text{inv}}(\tilde{u}) \tilde{\rho}_{\tilde{X}}(d\tilde{u}) = L_{K, \tilde{\rho}_{\tilde{X}}} \varphi \circ A_{\text{inv}}(A(x)) \end{aligned}$$

and hence if $L_{K, \rho_X} \varphi(x) = \lambda \varphi(x)$ then

$$L_{K, \tilde{\rho}_{\tilde{X}}}(\varphi \circ A_{\text{inv}})(A(x)) = L_{K, \rho_X} \varphi(x) = \lambda \varphi(x) = \lambda(\varphi \circ A_{\text{inv}})(A(x))$$

and vice versa.

(i): We first show that $f_\rho = f_{\tilde{\rho}} \circ A$. To this end, we show that $\rho(\cdot|x) = \tilde{\rho}(\cdot|A(x))$. First note that $\tilde{\rho}_X = \rho_X \circ A^{-1}$ ($\tilde{\rho}_X$ is the pushforward measure of ρ_X under the map A) and $\rho_X = \tilde{\rho}_X \circ A_{\text{inv}}^{-1}$. For a continuous and bounded function $g: \mathcal{Z} \rightarrow \mathbb{R}$, it holds

$$\begin{aligned} \int g d\rho &= \int g(A_{\text{inv}}(A(x)), y) \rho(dx, dy) \\ &= \int \int g(A_{\text{inv}}(\tilde{x}), y) \tilde{\rho}(dy|\tilde{x}) \tilde{\rho}_{\tilde{X}}(d\tilde{x}) \\ &= \int \int g(A_{\text{inv}}(\tilde{x}), y) \tilde{\rho}(dy|A(A_{\text{inv}}(\tilde{x}))) \tilde{\rho}_{\tilde{X}}(d\tilde{x}) \\ &= \int \int g(x, y) \tilde{\rho}(dy|A(x)) \rho_X(dx), \end{aligned}$$

which shows $\rho(\cdot|x) = \tilde{\rho}(\cdot|A(x))$.

(ii): We show $f_{\lambda, \rho} = f_{\lambda, \tilde{\rho}} \circ A$. By the relation of the eigenvalue/-function pairs of L_{K, ρ_X} and $L_{K, \tilde{\rho}_{\tilde{X}}}$ as shown at the beginning of the proof, it follows from Mercer's theorem (see, e.g., Steinwart and Scovel, 2012) that $\|f\|_{\mathcal{H}} = \|f \circ A\|_{\mathcal{H}}$ holds for all $f \in \mathcal{H}$ mapping from $\tilde{\mathcal{X}}$ to \mathbb{R} . Further,

$$\int (f(\tilde{x}) - y)^2 \tilde{\rho}(d\tilde{x}, dy) = \int (f(A(x)) - y)^2 \rho(dx, dy)$$

and thus $f_{\lambda, \tilde{\rho}} \circ A$ is the unique optimizer of

$$\arg \min_{f \in \mathcal{H}} \int (f(x) - y) \rho(dx, dy) + \lambda \|f\|_{\mathcal{H}}^2$$

which yields $f_{\lambda, \tilde{\rho}} \circ A = f_{\lambda, \rho}$.

(iii): The argument given at the beginning of the proof immediately implies the same eigenvalue decay of the operators L_{K, ρ_X} and $L_{K, \tilde{\rho}_{\tilde{X}}}$ and thus the parameter α can be chosen equally. Using (i) and (ii), we have

$$\int (f_{\lambda, \rho} - f_{\rho})^2 d\rho = \int (f_{\lambda, \tilde{\rho}} \circ A - f_{\rho} \circ A)^2 d\rho = \int (f_{\lambda, \tilde{\rho}} - f_{\tilde{\rho}})^2 d\tilde{\rho}.$$

Therefore, the parameter β can be chosen equally, too. \blacksquare

Proof of Lemma 4 We show that $c \leq a$, while $c \leq b$ will follow by symmetry. For $z \in \mathbb{R}^D$, define ρ_X^z as the distribution ρ_X shifted by z , i.e., $\rho_X^z(A) = \rho_X(A + z)$ for all Borel sets A . An elementary calculation yields

$$L_{K, \rho_X * \kappa}(f) = \int L_{K, \rho_X^z}(f) \kappa(dz).$$

Note that for all $z \in \mathbb{R}^D$ the operators L_{K, ρ_X^z} have the same eigenvalues as L_{K, ρ_X} , namely $\sigma_1, \sigma_2, \dots$. By linearity of the trace this shows

$$\sum_{i=1}^{\infty} \hat{\sigma}_i = \sum_{i=1}^{\infty} \sigma_i. \quad (15)$$

Consider now an empirical approximation $\kappa^m = \frac{1}{m} \sum_{i=1}^m \delta_{z_i}$ for certain $z_i \in \mathbb{R}^D$ and let $L_m := \int L_{K, \rho_X^z} \kappa^m(dz)$. Denote the eigenvalues of L_m in decreasing order by u_1^m, u_2^m, \dots . Using (Wielandt, 1955, Theorem 2)², we obtain

$$\sum_{i=1}^r u_i^m \leq \sum_{i=1}^r \sigma_i. \quad (16)$$

Consider $\Delta_m = L_{K, \rho_X * \kappa} - L_m$. Since κ has finite first moment, we can choose suitable z_1, \dots, z_m such that $W_1(\kappa, \kappa^m) \rightarrow 0$ for $m \rightarrow \infty$. We consider Δ_m as an operator from \mathcal{H} to \mathcal{H} , since this does not change its eigenvalue behavior, c.f. (Caponnetto and De Vito, 2007, Remark 2). Since K is bounded by 1 and Lipschitz continuous with constant L , we obtain for $f \in \mathcal{H}$, $\|f\|_{\mathcal{H}} = 1$ that

$$\|L_{K, \rho_X^z}(f) - L_{K, \rho_X^{\tilde{z}}}(f)\|_{\mathcal{H}} \leq 2L|z - \tilde{z}|$$

for all $z, \tilde{z} \in \mathbb{R}^D$. Hence, for an optimal coupling $\pi \in \Pi(\kappa, \kappa^m)$ attaining $W_1(\kappa, \kappa^m)$, we get

$$\|\Delta_m(f)\|_{\mathcal{H}} \leq \int \|L_{K, \rho_X^z}(f) - L_{K, \rho_X^{\tilde{z}}}(f)\|_{\mathcal{H}} \pi(dz, d\tilde{z}) \leq 2LW_1(\kappa, \kappa^m),$$

2. Note while the Theorem is only stated for finite dimensional operators, it is mentioned in the introduction that the techniques extend to self-adjoint compact operators on a Hilbert space, which fits our setting. See also the related discussion in Fulton (2000).

which implies $\|\Delta_m\|_{\text{op}} \rightarrow 0$ for $m \rightarrow \infty$. Thus we can apply stability results for the eigenvalues of $L_{K, \rho_X^{**\kappa}}$ (e.g., Chiappinelli (2000, Theorem 1)), showing that $u_i^m \rightarrow \hat{\sigma}_i$ for $i = 1, \dots, r$ and $m \rightarrow \infty$. Continuing from (16), we conclude

$$\sum_{i=1}^r \hat{\sigma}_i \leq \sum_{i=1}^r \sigma_i.$$

Together with (15), this yields

$$\sum_{i=r+1}^{\infty} \hat{\sigma}_i \geq \sum_{i=r+1}^{\infty} \sigma_i.$$

Since $\sigma_n \sim n^{-a}$, assuming $\hat{\sigma}_n \lesssim n^{-c}$, this inequality yields $c \leq a$, completing the proof. \blacksquare

Proof of Theorem 7 First, since the embedding $E: \mathcal{H} \rightarrow L^2(\rho_3)$ satisfies $E^*E = T_{\rho_3}$, using the polar decomposition for E one obtains

$$\|f_{\lambda, \rho_1} - f_{\lambda, \rho_2}\|_{L^2(\rho_3)} = \|\sqrt{T_{\rho_3}}(f_{\lambda, \rho_1} - f_{\lambda, \rho_2})\|_{\mathcal{H}},$$

which is shown for instance in De Vito and Caponnetto (2005, Proof of Proposition 2). Caponnetto and De Vito (2007, Proof of Theorem 4, Step 2.1) have shown by a Neumann series argument, that the assumption from Equation (6) implies

$$\|\sqrt{T_{\rho_3}}(T_{\rho_1} + \lambda I)^{-1}\|_{\text{op}} \leq \frac{1}{\sqrt{\lambda}}.$$

To bound $\|f_{\lambda, \rho_1} - f_{\lambda, \rho_2}\|_{L^2(\rho_3)}$, we use the decomposition

$$f_{\lambda, \rho_1} - f_{\lambda, \rho_2} = (T_{\rho_1} + \lambda I)^{-1}(g_{\rho_1} - g_{\rho_2}) + ((T_{\rho_1} + \lambda I)^{-1} - (T_{\rho_2} + \lambda I)^{-1})g_{\rho_2}$$

and obtain

$$\|f_{\lambda, \rho_1} - f_{\lambda, \rho_2}\|_{L^2(\rho_3)} \leq R_1 + R_2,$$

where

$$\begin{aligned} R_1 &= \|\sqrt{T_{\rho_3}}(T_{\rho_1} + \lambda I)^{-1}(g_{\rho_1} - g_{\rho_2})\|_{\mathcal{H}}, \\ R_2 &= \|\sqrt{T_{\rho_3}}((T_{\rho_1} + \lambda I)^{-1} - (T_{\rho_2} + \lambda I)^{-1})g_{\rho_2}\|_{\mathcal{H}}. \end{aligned}$$

To bound R_1 , we find

$$R_1 \leq \|\sqrt{T_{\rho_3}}(T_{\rho_1} + \lambda I)^{-1}\|_{\text{op}} \|g_{\rho_1} - g_{\rho_2}\|_{\mathcal{H}} \leq \frac{1}{\sqrt{\lambda}} \|g_{\rho_1} - g_{\rho_2}\|_{\mathcal{H}}$$

and for an optimal coupling $\pi \in \Pi(\rho_1, \rho_2)$ that attains $W_1(\rho_1, \rho_2)$ (such an optimal coupling exists by standard results, cf. (Villani et al., 2009, Theorem 4.1)) and abbreviating

$\pi(dx, dy) = \pi(d(x_1, y_1), d(x_2, y_2))$, we get

$$\begin{aligned}
 \|g_{\rho_1} - g_{\rho_2}\|_{\mathcal{H}} &= \left\| \int K(x_1, \cdot) y_1 \rho_1(dx_1, dy_1) - \int K(x_2, \cdot) y_2 \rho_2(dx_2, dy_2) \right\|_{\mathcal{H}} \\
 &= \left\| \int K(x_1, \cdot) y_1 - K(x_2, \cdot) y_2 \pi(dx, dy) \right\|_{\mathcal{H}} \\
 &\leq \int \|K(x_1, \cdot) y_1 - K(x_2, \cdot) y_2\|_{\mathcal{H}} \pi(dx, dy) \\
 &\leq \int |y_1| \|K(x_1, \cdot) - K(x_2, \cdot)\|_{\mathcal{H}} + |y_1 - y_2| \|K(x_1, \cdot)\|_{\mathcal{H}} \pi(dx, dy) \\
 &\leq \int ML|x_1 - x_2| + |y_1 - y_2| \pi(dx, dy) \\
 &\leq \max\{ML, 1\} W_1(\rho_1, \rho_2),
 \end{aligned}$$

which completes the bound for R_1 .

For R_2 , we calculate

$$\begin{aligned}
 R_2 &= \|\sqrt{T_{\rho_3}} \left((T_{\rho_1} + \lambda I)^{-1} (T_{\rho_1} - T_{\rho_2}) (T_{\rho_2} + \lambda I)^{-1} \right) g_{\rho_2}\|_{\mathcal{H}} \\
 &\leq \|\sqrt{T_{\rho_3}} (T_{\rho_1} + \lambda I)^{-1}\|_{\text{op}} \| (T_{\rho_1} - T_{\rho_2}) f_{\lambda, \rho_2} \|_{\mathcal{H}} \\
 &\leq \frac{1}{\sqrt{\lambda}} \| (T_{\rho_1} - T_{\rho_2}) f_{\lambda, \rho_2} \|_{\mathcal{H}}.
 \end{aligned}$$

We choose an optimal coupling $\pi \in \Pi(\rho_1, \rho_2)$ attaining $W_1(\rho_1, \rho_2)$ and denote by π_X its marginal distribution on the space $\mathcal{X} \times \mathcal{X}$. We get

$$\begin{aligned}
 \|(T_{\rho_1} - T_{\rho_2}) f_{\lambda, \rho_2}\|_{\mathcal{H}} &= \left\| \int K(x_1, \cdot) f_{\lambda, \rho_2}(x_1) - K(x_2, \cdot) f_{\lambda, \rho_2}(x_2) \pi_X(dx) \right\|_{\mathcal{H}} \\
 &\leq \int \|(K(x_1, \cdot) - K(x_2, \cdot)) f_{\lambda, \rho_2}(x_1)\|_{\mathcal{H}} \\
 &\quad + \|(f_{\lambda, \rho_2}(x_1) - f_{\lambda, \rho_2}(x_2)) K(x_2, \cdot)\|_{\mathcal{H}} \pi_X(dx) \\
 &\leq \int \|K(x_1, \cdot) - K(x_2, \cdot)\|_{\mathcal{H}} |f_{\lambda, \rho_2}(x_1)| \\
 &\quad + |f_{\lambda, \rho_2}(x_1) - f_{\lambda, \rho_2}(x_2)| \|K(x_2, \cdot)\|_{\mathcal{H}} \pi_X(dx) \tag{17} \\
 &\leq \int L \|f_{\lambda, \rho_2}\|_{\infty} |x_1 - x_2| + \|f_{\lambda, \rho_2}\|_{\text{Lip}} |x_1 - x_2| \pi_X(dx) \\
 &\leq (L \|f_{\lambda, \rho_2}\|_{\infty} + \|f_{\lambda, \rho_2}\|_{\text{Lip}}) W_1(\rho_1, \rho_2).
 \end{aligned}$$

It remains to verify condition $\|(T_{\rho_3} - T_{\rho_1})(T_{\rho_3} + \lambda I)^{-1}\|_{\text{op}} \leq \frac{1}{2}$ from Equation (6) based on (5). We have

$$\|(T_{\rho_3} - T_{\rho_1})(T_{\rho_3} + \lambda I)^{-1}\|_{\text{op}} \leq \|T_{\rho_3} - T_{\rho_1}\|_{\text{op}} \|(T_{\rho_3} + \lambda I)^{-1}\|_{\text{op}} \leq \lambda^{-1} \|T_{\rho_3} - T_{\rho_1}\|_{\text{op}}.$$

To bound the latter operator norm we use (17) and (4) to estimate for any $f \in \mathcal{H}$ and an optimal coupling $\pi \in \Pi(\rho_1, \rho_3)$ attaining $W_1(\rho_1, \rho_3)$

$$\begin{aligned} \|(T_{\rho_3} - T_{\rho_1})f\|_{\mathcal{H}} &\leq \int \|K(x_1, \cdot) - K(x_2, \cdot)\|_{\mathcal{H}} |\langle K(x_1, \cdot), f \rangle_{\mathcal{H}}| \\ &\quad + |\langle K(x_1, \cdot) - K(x_2, \cdot), f \rangle_{\mathcal{H}}| \|K(x_2, \cdot)\|_{\mathcal{H}} \pi_X(dx) \\ &\leq 2\|f\|_{\mathcal{H}} \int \|K(x_1, \cdot) - K(x_2, \cdot)\|_{\mathcal{H}} \pi_X(dx) \\ &\leq 2L\|f\|_{\mathcal{H}} W_1(\rho_1, \rho_3). \end{aligned}$$

Therefore, $\|T_{\rho_3} - T_{\rho_1}\|_{\text{op}} \leq 2LW_1(\rho_1, \rho_3)$ and the required condition is satisfied if we have $W_1(\rho_1, \rho_3) \leq \frac{\lambda}{4L}$. \blacksquare

5.2 Proofs of Section 3

Proof of Proposition 9 We apply Theorem 7 with the assumption given by Equation (6) for $\rho_3 := \tilde{\rho}$ and

$$\rho_1 := \frac{1}{n} \sum_{i=1}^n \delta_{(X_i^*, Y_i)}, \quad \rho_2 := \frac{1}{n} \sum_{i=1}^n \delta_{(\hat{X}_i, Y_i)}.$$

Hence, $W_1(\rho_1, \rho_2) \leq S_n \|\hat{P}_n - P\|_{\text{op}}$. It remains to verify that assumption (8) implies condition (5) with high probability. To this end, we proceed as Caponnetto and De Vito (2007, Proof of Theorem 4, Step 2.1). We can write

$$(T_{\rho_3} + \lambda)^{-1}(T_{\rho_3} - T_{\rho_1}) = -\frac{1}{n} \sum_{i=1}^n (\Xi_i - \mathbb{E}[\Xi_i])$$

for random Hilbert-Schmidt operators

$$\Xi_i : \mathcal{H} \rightarrow \mathcal{H} \quad f \mapsto \Xi_i f = (T_{\tilde{\rho}} + \lambda I)^{-1} T_{\delta_{X_i}}$$

where δ_{X_i} denotes the Dirac measure in X_i . Since $\|T_{\tilde{\rho}}\|_{\text{op}} \leq \int_{\mathcal{X}} \|T_{\delta_x}\|_{\text{op}} \tilde{\rho}_{X^*}(dx) \leq 1$ and thus $\|(T_{\tilde{\rho}} + \lambda I)^{-1}\|_{\text{op}} \leq \lambda^{-1}$ we indeed have a bounded Hilbert-Schmidt norm:

$$\|\Xi_i\|_{\text{HS}} \leq \|(T_{\tilde{\rho}} + \lambda I)^{-1}\|_{\text{op}} \|T_{\delta_{X_i}}\|_{\text{HS}} \leq \frac{1}{\lambda} \|K(X_i, \cdot)\|_{\mathcal{H}}^2 = \frac{1}{\lambda}.$$

Moreover,

$$\begin{aligned} \mathbb{E}[\|\Xi_i\|_{\text{HS}}^2] &\leq \int_{\mathcal{X}} \|T_{\delta_x}\|_{\text{op}} \text{Tr}((T_{\tilde{\rho}} + \lambda I)^{-2} T_{\delta_x}) \tilde{\rho}_{X^*}(dx) \\ &\leq \text{Tr}((T_{\tilde{\rho}} + \lambda I)^{-2} T_{\tilde{\rho}}) \\ &\leq \|(T_{\tilde{\rho}} + \lambda I)^{-1}\|_{\text{op}} \text{Tr}((T_{\tilde{\rho}} + \lambda I)^{-1} T_{\tilde{\rho}}) \leq \frac{\mathcal{N}(\lambda)}{\lambda}. \end{aligned}$$

Based on these estimates and independence of Ξ_i , we can apply a Bernstein inequality for Hilbert space valued random variables (Caponnetto and De Vito, 2007, Proposition 2) which yields

$$\|(T_{\rho_3} + \lambda)^{-1}(T_{\rho_3} - T_{\rho_1})\|_{\text{HS}} \leq 2 \log\left(\frac{6}{\eta}\right) \left(\frac{2}{n\lambda} + \sqrt{\frac{\mathcal{N}(\lambda)}{\lambda n}}\right)$$

with probability greater than $1 - \eta/3$. The condition on $\mathcal{N}(\lambda)$ together with $\|T_{\tilde{\rho}}\|_{\text{op}} \geq \lambda$ finally imply that the previous upper bound not larger than $1/2$. \blacksquare

Proof of Theorem 10 We use the error decomposition (9). The first term can be bounded by

$$\begin{aligned} \|(\hat{f}_{\lambda,n} \circ \hat{P}_n - \hat{f}_{\lambda,n} \circ P)\|_{L^2(\rho_X)}^2 &\leq \|\hat{f}_{\lambda,n}\|_{\text{Lip}}^2 \|\hat{P}_n - P\|_{L^2(\rho_X)}^2 \\ &\leq \|\hat{f}_{\lambda,n}\|_{\text{Lip}}^2 \|\hat{P}_n - P\|_{\text{op}}^2 \int x^2 \rho_X(dx). \end{aligned}$$

The second term can be treated with Proposition 9:

$$\|(\hat{f}_{\lambda,n} \circ P - f_{\lambda,n}^* \circ P)\|_{L^2(\rho_X)}^2 \leq \frac{S_n^2(1 + LM + \|\hat{f}_{\lambda,n}\|_{\text{Lip}} + \|\hat{f}_{\lambda,n}\|_{\infty})^2}{\lambda} \|\hat{P}_n - P\|_{\text{op}}^2.$$

while the last term $\|(\mathbf{f}_{\lambda,n}^* \circ P - f_{\tilde{\rho}} \circ P)\|_{L^2(\rho_X)}^2 = \|\mathbf{f}_{\lambda,n}^* - f_{\tilde{\rho}}\|_{L^2(\rho_X \circ P^{-1})}^2$ is the error term for the kernel learning problem for the distribution $\tilde{\rho}$. \blacksquare

Proof of Lemma 11 We require some additional notation. Let

$$R(P) := \int |x - P(x)|^2 \rho_X(dx) \quad \text{and} \quad R(\hat{P}_n) := \int |x - \hat{P}_n(x)|^2 \rho_X(dx).$$

Define $\mathcal{E}_d^{\text{PCA}} := R(\hat{P}_n) - R(P)$. Of course, \hat{P}_n is a random variable depending on (X_1, \dots, X_n) . By Reiß and Wahl (2020, Equation 2.21) it holds

$$\mathbb{E}[\|\hat{P}_n - P\|_2^2] \leq \frac{2\mathbb{E}[\mathcal{E}_d^{\text{PCA}}]}{\sigma_{d+1}^X - \sigma_d^X}.$$

In Reiß and Wahl (2020, Proposition 2, Equation 2.6) it is shown that there exists a constant C_{pca} which is independent of n such that

$$\mathbb{E}[\mathcal{E}_d^{\text{PCA}}] \leq \frac{C_{\text{pca}}^2}{n(\sigma_{d+1}^X - \sigma_d^X)}$$

and by an application of Markov's inequality, the claim follows. \blacksquare

Proof of Theorem 12 To apply Theorem 10, we use Steinwart et al. (2009, Theorem 1) (c.f. Lemma 2) to obtain

$$\|(\mathbf{f}_{\lambda,n}^* \circ P - f_{\tilde{\rho}} \circ P)\|_{L^2(\rho_X)}^2 \leq C_l \log(9/\eta) \left(\lambda^\beta + \frac{1}{\lambda n}\right)$$

with probably $1 - \eta/3$, where we note that part (i) of Assumption 1 always holds with $\alpha = 1$ (c.f. Remark 1). We further apply Lemma 11 with $\eta/3$ which yields

$$\|\hat{P}_n - P\|_{\text{op}} \leq \frac{\Gamma}{\eta\sqrt{n}}.$$

with probability $1 - \eta/3$. With a constant \bar{C} such that

$$\Gamma(1 + LM + \|\hat{f}_{\lambda,n}\|_{\text{Lip}} + \|\hat{f}_{\lambda,n}\|_{\infty})^2 \leq \bar{C}(\|\hat{f}_{\lambda,n}\|_{\text{Lip}}^2 + \|\hat{f}_{\lambda,n}\|_{\infty}^2 + 1)$$

Theorem 10 yields the claim. ■

Proof of Lemma 13 By construction of $\hat{f}_{\lambda,n}$ we have for any $f \in \mathcal{H}$

$$\frac{1}{n} \sum_{i=1}^n (\hat{f}_{\lambda,n}(\hat{X}_i) - Y_i)^2 + \lambda \|\hat{f}_{\lambda,n}\|_{\mathcal{H}}^2 \leq \frac{1}{n} \sum_{i=1}^n (f(\hat{X}_i) - Y_i)^2 + \lambda \|f\|_{\mathcal{H}}^2.$$

This gives the basic oracle inequality

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n (\hat{f}_{\lambda,n}(\hat{X}_i) - f_{\hat{\rho}}(\hat{X}_i))^2 + \lambda \|\hat{f}_{\lambda,n}\|_{\mathcal{H}}^2 &\leq \frac{1}{n} \sum_{i=1}^n (f(\hat{X}_i) - f_{\hat{\rho}}(\hat{X}_i))^2 + \lambda \|f\|_{\mathcal{H}}^2 \\ &\quad - \frac{2}{n} \sum_{i=1}^n (Y_i - f_{\hat{\rho}}(\hat{X}_i))(f(\hat{X}_i) - \hat{f}_{\lambda,n}(\hat{X}_i)). \end{aligned}$$

On the event

$$\mathcal{A}_n := \left\{ \left\| \frac{1}{n} \sum_{i=1}^n (Y_i - f_{\hat{\rho}}(\hat{X}_i)) K(\hat{X}_i, \cdot) \right\|_{\mathcal{H}} \leq \frac{\kappa}{\sqrt{n}} \right\}$$

we have

$$\begin{aligned} \left| \frac{2}{n} \sum_{i=1}^n (Y_i - f_{\hat{\rho}}(\hat{X}_i))(f(\hat{X}_i) - \hat{f}_{\lambda,n}(\hat{X}_i)) \right| &= \left| \frac{2}{n} \sum_{i=1}^n (Y_i - f_{\hat{\rho}}(\hat{X}_i)) \langle K(\hat{X}_i, \cdot), f - \hat{f}_{\lambda,n} \rangle_{\mathcal{H}} \right| \\ &\leq \|f - \hat{f}_{\lambda,n}\|_{\mathcal{H}} \left\| \frac{2}{n} \sum_{i=1}^n (Y_i - f_{\hat{\rho}}(\hat{X}_i)) K(\hat{X}_i, \cdot) \right\|_{\mathcal{H}} \\ &= \frac{2\kappa}{\sqrt{n}} \|f - \hat{f}_{\lambda,n}\|_{\mathcal{H}} \\ &\leq \frac{2\kappa}{\sqrt{n}} \|f\|_{\mathcal{H}} + \frac{2\kappa}{\sqrt{n}} \|\hat{f}_{\lambda,n}\|_{\mathcal{H}} \\ &\leq \frac{3\kappa}{n\lambda} + \lambda \|f\|_{\mathcal{H}}^2 + \frac{\lambda}{2} \|\hat{f}_{\lambda,n}\|_{\mathcal{H}}^2. \end{aligned}$$

Therefore, we have on \mathcal{A}_n for any $f \in \mathcal{H}$

$$\frac{1}{n} \sum_{i=1}^n (\hat{f}_{\lambda,n}(\hat{X}_i) - f_{\hat{\rho}}(\hat{X}_i))^2 + \frac{\lambda}{2} \|\hat{f}_{\lambda,n}\|_{\mathcal{H}}^2 \leq \frac{1}{n} \sum_{i=1}^n (f(\hat{X}_i) - f_{\hat{\rho}}(\hat{X}_i))^2 + \frac{3\kappa}{n\lambda} + 2\lambda \|f\|_{\mathcal{H}}^2.$$

If $f_{\hat{\rho}} \in \mathcal{H}$, we can choose $f = f_{\hat{\rho}}$ and obtain

$$\|\hat{f}_{\lambda,n}\|_{\mathcal{H}}^2 \leq \frac{2}{n\lambda} \sum_{i=1}^n (\hat{f}_{\lambda,n}(\hat{X}_i) - f_{\hat{\rho}}(\hat{X}_i))^2 + \|\hat{f}_{\lambda,n}\|_{\mathcal{H}}^2 \leq \frac{6\kappa}{n\lambda^2} + 4\|f_{\hat{\rho}}\|_{\mathcal{H}}^2.$$

For $\lambda = n^{-1/2}$ (which is the optimal choice for $\alpha = \beta = 1$) we conclude $\|\hat{f}_{\lambda,n}\|_{\mathcal{H}}^2 \leq 6\kappa + 4\|f_{\hat{\rho}}\|_{\mathcal{H}}^2$ on the event \mathcal{A}_n .

It remains to verify $\mathbb{P}(\mathcal{A}_n^c) \rightarrow 0$. To this end, we exploit $f_{\hat{\rho}}(X_i^*) = \mathbb{E}[Y_i|X_i^*]$ such that we can write $Y_i = f_{\hat{\rho}}(X_i^*) + \delta_i$ with $\mathbb{E}[\delta_i|X_i^*] = 0$ and $|\delta_i| \leq 2M$. The decomposition $Y_i - f_{\hat{\rho}}(\hat{X}_i) = \delta_i + f_{\hat{\rho}}(X_i^*) - f_{\hat{\rho}}(\hat{X}_i)$ yields

$$\begin{aligned} \left\| \frac{1}{n} \sum_{i=1}^n (Y_i - f_{\hat{\rho}}(\hat{X}_i)) K(\hat{X}_i, \cdot) \right\|_{\mathcal{H}} &\leq \left\| \frac{1}{n} \sum_{i=1}^n \delta_i K(X_i^*, \cdot) \right\|_{\mathcal{H}} + \left\| \frac{1}{n} \sum_{i=1}^n \delta_i (K(\hat{X}_i, \cdot) - K(X_i^*, \cdot)) \right\|_{\mathcal{H}} \\ &\quad + \left\| \frac{1}{n} \sum_{i=1}^n (f_{\hat{\rho}}(X_i^*) - f_{\hat{\rho}}(\hat{X}_i)) K(\hat{X}_i, \cdot) \right\|_{\mathcal{H}} \\ &=: T_1 + T_2 + T_3. \end{aligned}$$

In particular, we have $\mathbb{P}(\mathcal{A}_n^c) \leq \mathbb{P}(T_1 \geq \frac{\kappa}{3\sqrt{n}}) + \mathbb{P}(T_2 \geq \frac{\kappa}{3\sqrt{n}}) + \mathbb{P}(T_3 \geq \frac{\kappa}{3\sqrt{n}})$.

The first probability can be bounded with Markov's inequality and using RKHS properties:

$$\begin{aligned} \mathbb{P}(T_1 \geq \frac{\kappa}{2\sqrt{n}}) &\leq \frac{4n}{\kappa^2} \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n \delta_i K(X_i^*, \cdot) \right\|_{\mathcal{H}}^2 \right] \\ &= \frac{4n}{\kappa^2} \mathbb{E} \left[\frac{1}{n^2} \sum_{i,j=1}^n \delta_i \delta_j \langle K(X_i^*, \cdot), K(X_j^*, \cdot) \rangle_{\mathcal{H}} \right] = \frac{4}{\kappa^2 n} \sum_{i,j=1}^n \mathbb{E} [\delta_i \delta_j K(X_i^*, X_j^*)]. \end{aligned}$$

Since (δ_i) are independent and satisfy $\mathbb{E}[\delta_i|X_i^*] = 0$, $\mathbb{E}[\delta_i^2|X_i^*] =: \sigma^2 \leq 4M^2$, we obtain

$$\mathbb{P}(T_1 \geq \frac{\kappa}{2\sqrt{n}}) \leq \frac{9\sigma^2}{\kappa^2 n} \sum_{i=1}^n \mathbb{E}[K(X_i^*, X_i^*)] = \frac{9\sigma^2}{\kappa^2} \mathbb{E}[K(X_1^*, X_1^*)] \leq \eta$$

for $\kappa^2 > \frac{9\sigma^2}{\eta} \mathbb{E}[K(X_1^*, X_1^*)]$.

For the second probability, we use $|\delta_i| \leq 2M$ and (4) to bound

$$\begin{aligned} T_2 &\leq \frac{1}{n} \sum_{i=1}^n |\delta_i| \|K(\hat{X}_i, \cdot) - K(X_i^*, \cdot)\|_{\mathcal{H}} \leq \frac{2ML}{n} \sum_{i=1}^n |\hat{X}_i - X_i^*| \\ &\leq \frac{2ML}{n} \sum_{i=1}^n |(\hat{P} - P)X_i| \\ &\leq 2ML \|\hat{P} - P\|_{\text{op}} \left(\frac{1}{n} \sum_{i=1}^n |X_i|^2 \right)^{1/2}. \end{aligned}$$

By Lemma 11 we have $\mathbb{P}\left(\|\hat{P}_n - P\|_{\text{op}} \geq \frac{\Gamma}{3\eta\sqrt{n}}\right) \leq \eta$ with Γ from Theorem 12. Hence,

$$\begin{aligned} \mathbb{P}(T_2 \geq \frac{\kappa}{3\sqrt{n}}) &\leq \mathbb{P}\left(\frac{2ML\Gamma}{\eta} \left(\frac{1}{n} \sum_{i=1}^n |X_i|^2\right)^{1/2} \geq \kappa\right) + \eta \\ &\leq \frac{4(ML\Gamma)^2}{\eta^2\kappa^2} \mathbb{E}[|X_1|^2] + \eta \leq 2\eta \end{aligned}$$

for $\kappa^2 > \eta^{-3}4(ML\Gamma)^2\mathbb{E}[|X_1|^2]$.

For the third probability, we estimate

$$\begin{aligned} T_3 &\leq \frac{1}{n} \sum_{i=1}^n |f_{\hat{\rho}}(X_i^*) - f_{\hat{\rho}}(\hat{X}_i)| \|K(\hat{X}_i, \cdot)\|_{\mathcal{H}} \\ &\leq \|f_{\hat{\rho}}\|_{\text{Lip}} \sqrt{K(\hat{X}_i, \hat{X}_i)} \frac{1}{n} \sum_{i=1}^n |X_i^* - \hat{X}_i| = \|f_{\hat{\rho}}\|_{\text{Lip}} \frac{1}{n} \sum_{i=1}^n |X_i^* - \hat{X}_i|. \end{aligned}$$

From here we can proceed as with T_2 and conclude

$$\begin{aligned} \mathbb{P}(T_3 \geq \frac{\kappa}{3\sqrt{n}}) &\leq \mathbb{P}\left(\frac{\|f_{\hat{\rho}}\|_{\text{Lip}}\Gamma}{\eta} \left(\frac{1}{n} \sum_{i=1}^n |X_i|^2\right)^{1/2} \geq \kappa\right) + \eta \\ &\leq \frac{\|f_{\hat{\rho}}\|_{\text{Lip}}^2\Gamma^2}{\eta^2\kappa^2} \mathbb{E}[|X_1|^2] + \eta \leq 2\eta \end{aligned}$$

for $\kappa^2 > \eta^{-3}(\|f_{\hat{\rho}}\|_{\text{Lip}}\Gamma)^2\mathbb{E}[|X_1|^2]$. Therefore, there is a constant $D > 0$ such that $\mathbb{P}(\mathcal{A}_n^c) \leq 1 - 5\eta$ for $\kappa = D\eta^{-3/2}$. \blacksquare

Proof of Proposition 15 The main inequality (13) follows completely analogously to Theorem 12. The optimal choices of λ in the given regimes follows directly since for $n > m^{\alpha/(\beta+1)}$ and $\lambda = m^{-\frac{1}{1+\beta}}$, it holds $\frac{1}{\lambda m} > \frac{1}{\lambda^{\alpha}n}$. And on the other hand for $m > n^{\frac{1+\beta}{\beta+\alpha}}$ and $\lambda = n^{-\frac{1}{\beta+\alpha}}$ it holds $\frac{1}{\lambda m} < \frac{1}{\lambda^{\alpha}n}$. \blacksquare

Proof of Theorem 16 We decompose

$$\|f_{\rho} - f_{\hat{\rho}} \circ P\|_{L^2(\rho_X)} \leq \|f_{\rho} - f_{\rho} \circ P\|_{L^2(\rho_X)} + \|f_{\rho} \circ P - f_{\hat{\rho}} \circ P\|_{L^2(\rho_X)}.$$

For the second term, we note that

$$f_{\hat{\rho}}(P(X)) = \mathbb{E}[Y|P(X)] = \mathbb{E}[\mathbb{E}[Y|X]|P(X)] = \mathbb{E}[f_{\rho}(X)|P(X)].$$

Therefore,

$$\begin{aligned} \|f_{\rho} \circ P - f_{\hat{\rho}} \circ P\|_{L^2(\rho_X)}^2 &= \mathbb{E}\left[\left(f_{\rho}(P(X)) - \mathbb{E}[f_{\rho}(X)|P(X)]\right)^2\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[f_{\rho}(P(X)) - f_{\rho}(X)|P(X)\right]^2\right] \\ &\leq \mathbb{E}\left[\left(f_{\rho}(P(X)) - f_{\rho}(X)\right)^2\right] \\ &= \|f_{\rho} - f_{\rho} \circ P\|_{L^2(\rho_X)}^2. \end{aligned}$$

We obtain

$$\begin{aligned} \|f_\rho - f_{\hat{\rho}} \circ P\|_{L^2(\rho_X)} &\leq 2\|f_\rho - f_\rho \circ P\|_{L^2(\rho_X)} \\ &\leq 2L\mathbb{E}[|X - P(X)|^2]^{1/2} \leq 2L\left(\sum_{i=d+1}^D \sigma_i^X\right)^{1/2}. \end{aligned}$$

where the last inequality is the well known reconstruction error for PCA, see, e.g., Reiß and Wahl (2020). ■

Acknowledgments

Research supported by the Landesforschungsförderung Hamburg under the project LD-SODA. MT acknowledges financial support by the DFG through the grant TR 1349-3.

References

- Sabah A Abdul-Wahab, Charles S Bakheit, and Saleh M Al-Alawi. Principal component and multiple regression analysis in modelling of ground-level ozone and factors affecting its concentrations. *Environmental Modelling & Software*, 20(10):1263–1271, 2005.
- Kofi P Adragani and R Dennis Cook. Sufficient dimension reduction and prediction in regression. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1906):4385–4405, 2009.
- Pierre Alquier and Gérard Biau. Sparse single-index model. *Journal of Machine Learning Research*, 14(1), 2013.
- Peter Andras. High-dimensional function approximation with neural networks for large volumes of data. *IEEE transactions on neural networks and learning systems*, 29(2): 500–508, 2017.
- Ella Bingham and Heikki Mannila. Random projection in dimensionality reduction: applications to image and text data. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 245–250, 2001.
- Olivier Bousquet and André Elisseeff. Stability and generalization. *The Journal of Machine Learning Research*, 2:499–526, 2002.
- T Tony Cai, Cun-Hui Zhang, and Harrison H Zhou. Optimal rates of convergence for covariance matrix estimation. *The Annals of Statistics*, 38(4):2118–2144, 2010.
- Abdulkadir Canatar, Blake Bordelon, and Cengiz Pehlevan. Out-of-distribution generalization in kernel regression. *arXiv preprint arXiv:2106.02261*, 2021.
- Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.

- Raffaele Chiappinelli. Nonlinear stability of eigenvalues of compact self-adjoint operators. *Recent Trends in Nonlinear Analysis: Festschrift Dedicated to Alfonso Vignoli on the Occasion of His Sixtieth Birthday*, pages 93–103, 2000.
- Hugh A Chipman and Hong Gu. Interpretable dimension reduction. *Journal of applied statistics*, 32(9):969–987, 2005.
- Andreas Christmann and Ingo Steinwart. Consistency and robustness of kernel-based regression in convex risk minimization. *Bernoulli*, 13(3):799–819, 2007.
- Andreas Christmann, Daohong Xiang, and Ding-Xuan Zhou. Total stability of kernel methods. *Neurocomputing*, 289:101–118, 2018.
- Michael AA Cox and Trevor F Cox. Multidimensional scaling. In *Handbook of data visualization*, pages 315–347. Springer, 2008.
- Felipe Cucker and Steve Smale. On the mathematical foundations of learning. *American Mathematical Society*, 39(1):1–49, 2002.
- Hugo Cui, Bruno Loureiro, Florent Krzakala, and Lenka Zdeborová. Generalization error rates in kernel regression: The crossover from the noiseless to noisy regime. *arXiv preprint arXiv:2105.15004*, 2021.
- Ernesto De Vito and Andrea Caponnetto. Risk bounds for regularized least-squares algorithm with operator-value kernels. Technical report, 2005.
- Ronald DeVore, Guergana Petrova, and Przemyslaw Wojtaszczyk. Approximation of functions of few variables in high dimensions. *Constructive Approximation*, 33(1):125–143, 2011.
- Simon Fischer and Ingo Steinwart. Sobolev norm learning rates for regularized least-squares algorithms. *The Journal of Machine Learning Research*, 21(1):8464–8501, 2020.
- Dmitriy Fradkin and David Madigan. Experiments with random projections for machine learning. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 517–522, 2003.
- William Fulton. Eigenvalues, invariant factors, highest weights, and schubert calculus. *Bulletin of the American Mathematical Society*, 37(3):209–249, 2000.
- Stéphane Gaïffas and Guillaume Lecué. Optimal rates and adaptation in the single-index model using aggregation. *Electronic journal of statistics*, 1:538–573, 2007.
- Chong Gu. *Smoothing spline ANOVA models*, volume 297. Springer Science & Business Media, 2013.
- M Guillemard and A Iske. Interactions between kernels, frames, and persistent homology. In *Recent Applications of Harmonic Analysis to Function Spaces, Differential Equations, and Data Science. Volume 2: Novel Methods in Harmonic Analysis*, pages 861–888. Birkhäuser, 2017.

- Thomas Hamm and Ingo Steinwart. Adaptive learning rates for support vector machines working on data with low intrinsic dimension. *The Annals of Statistics*, 49(6):3153–3180, 2021.
- Thomas Hamm and Ingo Steinwart. Intrinsic dimension adaptive partitioning for kernel methods. *SIAM Journal on Mathematics of Data Science*, 4(2):721–749, 2022.
- XG Hua, YQ Ni, JM Ko, and KY Wong. Modeling of temperature–frequency correlation using combined principal component analysis and support vector regression technique. *Journal of Computing in Civil Engineering*, 21(2):122–135, 2007.
- Armin Iske. *Approximation Theory and Algorithms for Data Analysis*, volume 68 of *Texts in Applied Mathematics*. Springer, 2018.
- Ian T Jolliffe. *Principal component analysis*. Springer, New York, second edition, 2002.
- Hyunsoo Kim, Peg Howland, Haesun Park, and Nello Christianini. Dimension reduction in text classification with support vector machines. *Journal of machine learning research*, 6(1), 2005.
- John A Lee and Michel Verleysen. *Nonlinear dimensionality reduction*, volume 1. Springer, 2007.
- Ker-Chau Li. Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327, 1991.
- Junhong Lin, Alessandro Rudi, Lorenzo Rosasco, and Volkan Cevher. Optimal rates for spectral algorithms with least-squares regression over hilbert spaces. *Applied and Computational Harmonic Analysis*, 48(3):868–890, 2020.
- Yi Lin and Hao Helen Zhang. Component selection and smoothing in multivariate non-parametric regression. *The Annals of Statistics*, 34(5):2272–2297, 2006.
- Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. Umap: Uniform manifold approximation and projection. *Journal of Open Source Software*, 3(29):861, 2018.
- Alexander Meister. *Deconvolution problems in nonparametric statistics*, volume 193 of *Lecture Notes in Statistics*. Springer-Verlag, Berlin, 2009.
- Shahar Mendelson and Joseph Neeman. Regularization in kernel learning. *The Annals of Statistics*, 38(1):526–565, 2010.
- Andrew Ng. Sparse autoencoder. *CS294A Lecture notes*, 72(2011):1–19, 2011.
- Garvesh Raskutti, Martin J Wainwright, and Bin Yu. Minimax-optimal rates for sparse additive models over kernel classes via convex programming. *Journal of Machine Learning Research*, 13(2), 2012.
- Markus Reiß and Martin Wahl. Nonasymptotic upper bounds for the reconstruction error of PCA. *Annals of Statistics*, 48(2):1098–1123, 2020.

- Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Kernel principal component analysis. In *International conference on artificial neural networks*, pages 583–588. Springer, 1997.
- Ingo Steinwart and Andreas Christmann. *Support vector machines*. Information Science and Statistics. Springer, New York, 2008. ISBN 978-0-387-77241-7.
- Ingo Steinwart and Clint Scovel. Mercer’s theorem on general domains: On the interaction between measures, kernels, and RKHSs. *Constructive Approximation*, 35(3):363–417, 2012.
- Ingo Steinwart, Don R Hush, and Clint Scovel. Optimal rates for regularized least squares regression. In *COLT*, pages 79–93, 2009.
- Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schoelkopf. Wasserstein auto-encoders. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=HkL7n1-0b>.
- Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
- Yasi Wang, Hongxun Yao, and Sicheng Zhao. Auto-encoder based dimensionality reduction. *Neurocomputing*, 184:232–242, 2016.
- Holger Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.
- Helmut Wielandt. An extremum property of sums of eigenvalues. *Proceedings of the American Mathematical Society*, 6(1):106–110, 1955.
- Svante Wold, Kim Esbensen, and Paul Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52, 1987.
- Qiang Wu, Yiming Ying, and Ding-Xuan Zhou. Learning rates of least-square regularized regression. *Foundations of computational mathematics*, 6(2):171–192, 2006.
- Zongmin Wu. Compactly supported positive definite radial functions. *Advances in computational mathematics*, 4:283–292, 1995.
- Huan Xu, Constantine Caramanis, and Shie Mannor. Robustness and regularization of support vector machines. *Journal of machine learning research*, 10(7), 2009.
- Yi Yu, Tengyao Wang, and Richard J Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2015.
- Shengxin Zhu. Compactly supported radial basis functions: how and why? *OCCAM Preprint Number 12/57*, 2012.
- Xiaojin Zhu and Andrew B Goldberg. Introduction to semi-supervised learning. *Synthesis lectures on artificial intelligence and machine learning*, 3(1):1–130, 2009.
- Xiaojin Jerry Zhu. Semi-supervised learning literature survey. *Technical Report*, 2005.