
PSD Representations for Effective Probability Models

Alessandro Rudi

Inria, École normale supérieure,
CNRS, PSL Research University, Paris, France
alessandro.rudi@inria.fr

Carlo Ciliberto

Department of Computer Science
University College London, London, UK
c.ciliberto@ucl.ac.uk

Abstract

Finding a good way to model probability densities is key to probabilistic inference. An ideal model should be able to concisely approximate any probability while being also compatible with two main operations: multiplications of two models (product rule) and marginalization with respect to a subset of the random variables (sum rule). In this work, we show that a recently proposed class of positive semi-definite (PSD) models for non-negative functions is particularly suited to this end. In particular, we characterize both approximation and generalization capabilities of PSD models, showing that they enjoy strong theoretical guarantees. Moreover, we show that we can perform efficiently both sum and product rule in closed form via matrix operations, enjoying the same versatility of mixture models. Our results open the way to applications of PSD models to density estimation, decision theory and inference.

1 Introduction

Modeling probability distributions is a key task for many applications in machine learning [17, 5]. To this end, several strategies have been proposed in the literature, such as adopting mixture models (e.g. Gaussian mixtures) [5], exponential models [32], implicit generative models [11, 13] or kernel (conditional) mean embeddings [16]. An ideal probabilistic model should have two main features: *i*) efficiently perform key operations for probabilistic inference, such as sum rule (i.e. marginalization) and product rule [5] and, *ii*) concisely approximate a large class of probabilities. Finding models that satisfy these two conditions is challenging and current methods tend to tackle only one of the two. Exponential and implicit generative models have typically strong approximation properties (see e.g. [32, 29]) but cannot easily perform operations such as marginalization. On the contrary, mixture models are designed to efficiently integrate and multiply probabilities, but tend to require a large number of components to approximate complicated distributions.

In principle, mixture models would offer an appealing strategy to model probability densities since they allow for efficient computations when performing key operations such as sum and product rule. However, these advantages in terms of computations, come as a disadvantage in terms of expressiveness: even though mixture models are universal approximators (namely they can approximate arbitrarily well any probability density), they require a significant number n of observations and of components to do that. Indeed it is known that models that are non-negative mixtures of non-negative components lead to learning rates that suffer from the curse of dimensionality. For example, when approximating probabilities on \mathbb{R}^d , kernel density estimation (KDE) [21] with non-negative components has rates as slow as $n^{-2/(4+d)}$ (see, e.g. [38, page 100]), and cannot improve with the regularity (smoothness) of the target probability (see e.g. [9, Thm. 16.1] for impossibility results in the one dimensional case).

In the past decades, this limitation has been overcome by removing either the non-negativity of the weights of the mixture (leading to RBF networks [27][35]), or the non-negativity of the components

	Non-negative	Sum Rule	Product Rule	Concise Approximation	Optimal Learning	Efficient Sampling
Linear Models	✗	✓	✓	✓	✓	✗
Mean Embeddings	✗	✓	✓	✓	✓	✗
Mixture Models	✓	✓	✓	✗	✗	✓
Exponential Models	✓	✗	✓	✓	✓	✗
PSD Models	✓ (see [14])	✓ (Prop. 1)	✓ (Prop. 2)	✓ (Thm. 6)	✓ (Thm. 7)	✓ (see [15])

Table 1: Summary of the main desirable properties for a probability model.

used in the mixture (leading to KDE with oscillating kernel [35]), or both. On the positive side this allows to achieve a learning rate of $n^{-\beta/(2\beta+d)}$, for β -times differentiable densities on \mathbb{R}^d . Note that such rate is *minimax optimal* [10] and overcomes the curse of dimensionality when $\beta \geq d$. Additionally the resulting model is very *concise*, since only $m = O(n^{d/\beta})$ centers are necessary to achieve the optimal rate [36]. However, on the negative side, *the resulting model is not a probability*, since it may attain negative values, and so it cannot be used where a proper probability is required.

In this paper, we show that the positive semidefinite (PSD) models, recently proposed in [14], offer a way out of this dichotomy. By construction PSD models generalize mixture models by allowing also for negative weights, while still guaranteeing that the resulting function is non-negative everywhere. Here, we prove that they get the best of both worlds: expressivity (optimal learning and approximation rates with concise models) and flexibility (exact and efficient sum and product rule).

Contributions. The main contributions of this paper are:

- i) Showing that PSD models can perform exact sum and product rules in terms of efficient matrix operations (Sec. 2.1).
- ii) Characterizing the approximation and learning properties of PSD models with respect to a large family of probability densities (Sec. 3).
- iii) Providing a “compression” method to control the number of components of a PSD model (introduced in Sec. 2.2 and analyzed in Sec. 3.3).
- iv) Discussing a number of possible applications of PSD models (Sec. 4).

Summary of the results. Table 1 summarizes all the desirable properties that a probability model should have and which of them are satisfied by state-of-the-art estimators. Among these we have: *non-negativity*, the model should be point-wise non-negative; *sum and product rule*, the model should allow for efficient computation of such operations between probabilities; *concise approximation*, a model with the optimal number of components $m = O(\varepsilon^{-d/\beta})$ is enough to approximate with error ε a non-negative function in a wide family of smooth β -times differentiable functions [36]; *optimal learning*, the model achieves optimal learning rates of order $n^{-\beta/(2\beta+d)}$ when learning the unknown target probability from i.i.d. samples [35]; *efficient sampling* the model allows to extract i.i.d. samples without incurring in the curse of dimensionality in terms of computational complexity. We note that for PSD models, this last property has been recently investigated in [15]. We consider under the umbrella of *linear models* all the models which corresponds to a mixture $\sum_{j=1}^m w_j f_j(x)$ of functions $f_j : \mathcal{X} \rightarrow \mathbb{R}$, with either $w_j \geq 0, \forall j = 1, \dots, m$ or $f_j \geq 0, \forall j = 1, \dots, m$ or both. We denote by *mixture models* the mixture models defined as above, where $w_j \geq 0, f_j \geq 0, \forall j = 1, \dots, m$. By *mean embeddings*, we denote the kernel mean embedding estimator framework from [30] see also [16]). With *exponential models* we refer to models of the form $\exp(\sum_{j=1}^m w_j f_j(x))$ with $w_j \in \mathbb{R}, f_j : X \rightarrow \mathbb{R}$ [32]. With *PSD models* we refer to the framework introduced in [14] and studied in this paper. We note that this work contributes in showing that the latter are the only probability models to satisfy all requirements (we recall that non-negativity and efficient sampling have been shown respectively in [14] and [15]).

Notation. We denote by \mathbb{R}_{++}^d the space vectors in \mathbb{R}^d with positive entries, $\mathbb{R}^{n \times d}$ the space of $n \times d$ matrices, $\mathbb{S}_+^n = \mathbb{S}_+(\mathbb{R}^n)$ the space of positive semidefinite $n \times n$ matrices. Given a vector $\eta \in \mathbb{R}^d$, we denote $\text{diag}(\eta) \in \mathbb{R}^{d \times d}$ the diagonal matrix associated to η . We denote by $A \circ B$

and $A \otimes B$ respectively the entry-wise and Kronecker product between two matrices A and B . We denote by $\|A\|_F$, $\|A\|$, $\det(A)$, $\text{vec}(A)$ and A^\top respectively the Frobenius norm, the operator norm (i.e. maximum singular value), the determinant, the (column-wise) vectorization of a matrix and the (conjugate) transpose of A . With some abuse of notation, where clear from context we write element-wise products and division of vectors $u, v \in \mathbb{R}^d$ as $uv, u/v$. Given two matrices $X \in \mathbb{R}^{n \times d_1}$, $Y \in \mathbb{R}^{n \times d_2}$ with same number of rows, we denote by $[X, Y] \in \mathbb{R}^{n \times (d_1 + d_2)}$ their concatenation row-wise. The term $\mathbf{1}_n \in \mathbb{R}^n$ denotes the vector with all entries equal to 1.

2 PSD Models

Following [14], in this work we consider the family of positive semi-definite (PSD) models, namely non-negative functions parametrized by a feature map $\phi : \mathcal{X} \rightarrow \mathcal{H}$ from an input space \mathcal{X} to a suitable feature space \mathcal{H} (a separable Hilbert space e.g. \mathbb{R}^q) and a linear operator $M \in \mathbb{S}_+(\mathcal{H})$, of the form

$$f(x; M, \phi) = \phi(x)^\top M \phi(x). \quad (1)$$

PSD models offer a general way to parametrize non-negative functions (since M is positive semidefinite, $f(x; M, \phi) \geq 0$ for any $x \in \mathcal{X}$) and enjoy several additional appealing properties discussed in the following. In this work, we will focus on a special family of models of the form (1) to parametrize probability densities over $\mathcal{X} = \mathbb{R}^d$. In particular, we will consider the case where: *i*) $\phi = \phi_\eta : \mathbb{R}^d \rightarrow \mathcal{H}_\eta$ is a feature map associated to the Gaussian kernel [26] $k_\eta(x, x') = \phi_\eta(x)^\top \phi_\eta(x') = e^{-(x-x')^\top \text{diag}(\eta)(x-x')}$, with $\eta \in \mathbb{R}_{++}^d$ and, *ii*) the operator M lives in the span of $\phi(x_1), \dots, \phi(x_n)$ for a given set of points $(x_i)_{i=1}^n$, namely there exists $A \in \mathbb{S}_+^n$ such that $M = \sum_{i,j} A_{ij} \phi(x_i) \phi(x_j)^\top$. We define a *Gaussian PSD model* by specializing the definition in (1) as

$$f(x; A, X, \eta) = \sum_{i,j=1}^n A_{i,j} k_\eta(x_i, x) k_\eta(x, x_j), \quad \forall x \in \mathbb{R}^d \quad (2)$$

in terms of the coefficient matrix $A \in \mathbb{S}_+^n$, the base points matrix $X \in \mathbb{R}^{n \times d}$, whose i -th row corresponds to the point x_i for each $i = 1, \dots, n$ and kernel parameter η . In the following, given two base point matrices $X \in \mathbb{R}^{n \times d}$ and $X' \in \mathbb{R}^{m \times d}$, we denote by $K_{X, X', \eta} \in \mathbb{R}^{n \times m}$ the kernel matrix with entries $(K_{X, X', \eta})_{ij} = k_\eta(x_i, x'_j)$ where x_i, x'_j are the i -th and j -th rows of X, X' respectively. When clear from context, in the following we will refer to Gaussian PSD models as PSD models.

Remark 1 (PSD models generalize Mixture models). *Mixture models (a mixture of Gaussian distributions) are a special case of PSD models. Let $A = \text{diag}(a)$ be a diagonal matrix of n positive weights $a \in \mathbb{R}_{++}^n$. We have $f(x; A, X, \eta/2) = \sum_{i=1}^n A_{ii} k_{\eta/2}(x_i, x)^2 = \sum_{i=1}^n a_i k_\eta(x_i, x)$.*

Remark 2 (PSD models allow negative weights). *From (2), we immediately see that PSD models generalize mixture models by allowing also for negative weights: e.g., $f(\cdot; A, X, \eta)$ with $A = (1, -\frac{1}{2}; -\frac{1}{2}, \frac{1}{4}) \in \mathbb{S}_+^2, \eta = 1, X = (x_1; x_2)$ and $x_1 = 0, x_2 = 2$, corresponding to $f(x; A, X, \eta) = e^{-2x^2} + \frac{1}{4}e^{-2(x-2)^2} - \frac{1}{2}e^{-2(x-1)^2}$, i.e. a mixture of Gaussians with also negative weights.*

2.1 Operations with PSD models

In Sec. 3 we will show that PSD models can approximate a wide class of probability densities, significantly outperforming mixture models. Here we show that this improvement does not come at the expenses of computations. In particular, we show that PSD models enjoy the same flexibility of mixture models: *i*) they are closed with respect to key operations such as marginalization and multiplication and *ii*) these operations can be performed efficiently in terms of matrix sums/products. The derivation of the results reported in the following is provided in Appendix F. They follow from well-known properties of the Gaussian function.

Evaluation. Evaluating a PSD model in a point $x_0 \in \mathcal{X}$ corresponds to $f(x = x_0; A, X, \eta) = K_{X, x_0, \eta}^\top A K_{X, x_0, \eta}$. Moreover, partially evaluating a PSD in one variable yields

$$f(x, y = y_0; A, [X, Y], (\eta_1, \eta_2)) = f(x; B, X, \eta_1) \quad \text{with} \quad B = A \circ (K_{Y, y_0, \eta_2} K_{Y, y_0, \eta_2}^\top). \quad (3)$$

Note that $f(x; B, X, \eta_1)$ is still a PSD model since B is positive semidefinite.

Sum Rule (Marginalization and Integration). The integral of a PSD model can be computed as

$$\int f(x; A, X, \eta) dx = c_{2\eta} \text{Tr}(A K_{X, X, \frac{\eta}{2}}) \quad \text{with} \quad c_\eta = \int k_\eta(0, x) dx, \quad (4)$$

where $c_\eta = \pi^{d/2} \det(\text{diag}(\eta))^{-1/2}$. This is particularly useful to model probability densities with PSD models. Let $Z = \int f(x; A, X, \eta) dx$, then the function $f(x; A/Z, X, \eta) = \frac{1}{Z} f(x; A, X, \eta)$ is a probability density. Integrating only one variable of a PSD model we obtain the sum rule.

Proposition 1 (Sum Rule – Marginalization). *Let $X \in \mathbb{R}^{n \times d}$, $Y \in \mathbb{R}^{n \times d'}$, $A \in \mathbb{S}_+(\mathbb{R}^n)$ and $\eta \in \mathbb{R}_{++}^d, \eta' \in \mathbb{R}_{++}^{d'}$. Then, the following integral is a PSD model*

$$\int f(x, y; A, [X, Y], (\eta, \eta')) dx = f(y; B, Y, \eta'), \quad \text{with} \quad B = c_{2\eta} A \circ K_{X, X, \frac{\eta}{2}}, \quad (5)$$

The result above shows that we can efficiently marginalize a PSD model with respect to one variable by means of an entry-wise multiplication between two matrices.

Remark 3 (Integration and marginalization on the hypercube). *The integrals in (4) and (5) can be performed also on when \mathcal{X} is a hypercube $H = \prod_{t=1}^d [a_t, b_t]$ rather than the entire space \mathbb{R}^d . This leads to a closed form, where the matrix $K_{X, X, \frac{\eta}{2}}$ is replaced by a suitable $K_{X, X, \frac{\eta}{2}, H}$ that can be computed with same number of operations (the full form of such matrix is reported in [Appendix F](#)).*

Product Rule (Multiplication). Multiplying two probabilities is key to several applications in probabilistic inference [5]. The family of PSD models is closed with respect to this operation.

Proposition 2 (Multiplication). *Let $X \in \mathbb{R}^{n \times d_1}$, $Y \in \mathbb{R}^{n \times d_2}$, $Y' \in \mathbb{R}^{m \times d_2}$, $Z \in \mathbb{R}^{m \times d_3}$, $A \in \mathbb{S}_+^n$, $B \in \mathbb{S}_+^m$ and $\eta_1 \in \mathbb{R}_{++}^{d_1}$, $\eta_2, \eta'_2 \in \mathbb{R}_{++}^{d_2}$, $\eta_3 \in \mathbb{R}_{++}^{d_3}$. Then*

$$f(x, y; A, [X, Y], (\eta_1, \eta_2)) f(y, z; B, [Y', Z], (\eta'_2, \eta_3)) = f(x, y, z; C, W, \eta), \quad (6)$$

is a PSD model, where $C = (A \otimes B) \circ (\text{vec}(K_{Y, Y', \tilde{\eta}_2}) \text{vec}(K_{Y, Y', \tilde{\eta}_2})^\top)$, with $\tilde{\eta}_2 = \frac{\eta_2 \eta'_2}{\eta_2 + \eta'_2}$, base matrix $W = [X \otimes \mathbf{1}_m, \frac{\eta_2}{\eta_2 + \eta'_2} Y \otimes \mathbf{1}_m + \frac{\eta'_2}{\eta_2 + \eta'_2} \mathbf{1}_n \otimes Y', \mathbf{1}_n \otimes Z]$ and $\eta = (\eta_1, \eta_2 + \eta'_2, \eta_3)$.

We note that, despite the heavy notation, multiplying two PSD is performed via simple operations such as tensor and entry-wise product between matrices. In particular, we note that $X \otimes \mathbf{1}_m \in \mathbb{R}^{nm \times d_1}$ and $\mathbf{1}_n \otimes Z \in \mathbb{R}^{nm \times d_3}$ correspond respectively to the $nm \times d_1$ -matrix containing m copies of each row of X and the $nm \times d_3$ -matrix containing n copies of Z . Finally, $\eta Y \otimes \mathbf{1}_m + \eta' \mathbf{1}_n \otimes Y$ is the $nm \times d_2$ base point matrix containing all possible sums $(\eta y_i + \eta' y'_j)_{i,j=1}^{n,m}$ of points from Y and Y' .

Reduction. As observed above, when performing operations with PSD models, the resulting base point matrix might be of the form $X \otimes \mathbf{1}_m$ (e.g. if we couple the multiplication in (6) with marginalization as in a Markov transition, see the corollary below). In these cases we can reduce the PSD model to have only n base points (rather than nm), as follows

$$f(x; A, X \otimes \mathbf{1}_m, \eta) = f(x; B, X, \eta) \quad \text{with} \quad B = (I_n \otimes \mathbf{1}_m^\top) A (I_n \otimes \mathbf{1}_m), \quad (7)$$

where $A \in \mathbb{S}_+^{nm}$ and $I_n \in \mathbb{R}^{n \times n}$ is the $n \times n$ identity matrix. The reduction operation is useful to avoid the dimensionality of the PSD model grow unnecessarily. This is for instance the case of a Markov transition.

Corollary 3 (Markov Transition). *Let $X \in \mathbb{R}^{n \times d_1}$, $Y \in \mathbb{R}^{n \times d_2}$, $Y' \in \mathbb{R}^{m \times d_2}$, $A \in \mathbb{S}_+^n$, $B \in \mathbb{S}_+^m$ and $\eta_1 \in \mathbb{R}_{++}^{d_1}$, $\eta_2, \eta'_2 \in \mathbb{R}_{++}^{d_2}$. Then*

$$\int f(x, y; A, [X, Y], (\eta_1, \eta_2)) f(y; B, Y', \eta'_2) dy = f(x; C, X, \eta_1), \quad (8)$$

with $C \in \mathbb{S}_+^n$ obtained by applying in order, [Prop. 2](#), [Prop. 1](#) and reduction (7).

We remark that the result of a Markov transition retains the same base point matrix X and parameters η of the transition kernel. This is thanks to the reduction operation in (7), which avoids the resulting matrix C to be $nm \times nm$. This fact is particularly useful in applications that require multiple Markov transitions, such as in hidden Markov models (see also [Sec. 4](#)).

2.2 Compression of a PSD model

From Prop. 2 we note that the multiplication operation can rapidly yield a large number of base points and thus incur in high computational complexity in some settings. It might therefore be useful to have a method to reduce the number of base points while retaining essentially the same model. To this purpose, here we propose a dimensionality reduction strategy. In particular, given a set of points $\tilde{x}_1, \dots, \tilde{x}_m \in \mathbb{R}^d$, we leverage the representation of the PSD model in terms of reproducing kernel Hilbert spaces [3] to use powerful sketching techniques as *Nyström projection* (see, e.g., [40]), to project the PSD model on a new PSD model now based only on the new points. Given $A \in \mathbb{S}_+^n$, $X \in \mathbb{R}^{n \times d}$, $\eta \in \mathbb{R}_{++}^d$. Let $\tilde{X} \in \mathbb{R}^{m \times d}$ be the base point matrix whose j -rows corresponds to the point \tilde{x}_j . The *compression* of $f(\cdot; A, X, \eta)$ corresponds to

$$f(x; \tilde{A}, \tilde{X}, \eta) \quad \text{with} \quad \tilde{A} = BAB^\top \in \mathbb{S}_+^m, \quad B = K_{\tilde{X}, \tilde{X}, \eta}^{-1} K_{\tilde{X}, X, \eta} \in \mathbb{R}^{m \times n}, \quad (9)$$

which it is still a PSD model since the matrix $BAB^\top \in \mathbb{S}_+^m$. We note that even with a rather simple strategy to choose the new base points $\tilde{x}_1, \dots, \tilde{x}_m$ – such as uniform sampling – compression is an effective tool to reduce the computational complexity of costly operations. In particular, in Sec. 3.3 we show that a compressed model with only $O(t \text{ polylog}(1/\varepsilon))$ centers (instead of t^2) can ε -approximate the product of two PSD models with t points each.

3 Representation power of PSD models

In this section we study the theoretical properties of Gaussian PSD models. We start by showing that they admit concise approximations of the target density (in the sense discussed in the introduction to this paper and of Table 1). We then proceed to studying the setting in which we aim to learn an unknown probability from i.i.d. samples. We conclude the section by characterizing the approximation properties of the compression operation introduced in Sec. 2.2.

3.1 Approximation properties of Gaussian PSD models

We start the section recalling that Gaussian PSD models are universal approximators for probability densities. In particular, the following result restates [Thm. 2 14] for the case of probabilities.

Proposition 4 (Universal consistency – Thm. 2 in [14]). *The Gaussian PSD family is a universal approximator for probabilities that admit a density.*

The result above is not surprising since Gaussian PSD models generalize classical Gaussian mixtures (see Remark 1), which are known to be universal [5]. We now introduce a mild assumption that will enable us to complement Prop. 4 with approximation and learning results. In the rest of the section we assume that $\mathcal{X} = (-1, 1)^d$ (or more generally an open bounded subset of \mathbb{R}^d with Lipschitz boundary). Here $L^\infty(\mathcal{X})$ and $L^2(\mathcal{X})$ denote respectively the space of essentially bounded and square-integrable functions over \mathcal{X} , while $W_2^\beta(\mathcal{X})$ denotes the *Sobolev space* of functions whose weak derivatives up to order β are square-integrable on \mathcal{X} (see [1] or Appendix A for more details).

Assumption 1. *Let $\beta > 0, q \in \mathbb{N}$. There exists $f_1, \dots, f_q \in W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$, such that the density $p: \mathcal{X} \rightarrow \mathbb{R}$ satisfies*

$$p(x) = \sum_{j=1}^q f_j(x)^2, \quad \forall x \in \mathcal{X}. \quad (10)$$

The assumption above is quite general and satisfied by a wide family of probabilities, as discussed in the following proposition. The proof is reported in Appendix D.1

Proposition 5 (Generality of Assumption 1). *The assumption above is satisfied by*

- (a) *any probability density $p \in W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$ and strictly positive on $[-1, 1]^d$,*
- (b) *any exponential model $p(x) = e^{-v(x)}$ with $v \in W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$,*
- (c) *any mixture model of Gaussians or, more generally, of exponential models from (b),*
- (d) *any p that is $\beta + 2$ -times differentiable on $[-1, 1]^d$, with a finite set of zeros, all in $(-1, 1)^d$, and with positive definite Hessian in each zero. E.g. $p(x) \propto x^2 e^{-x^2}$.*

Moreover if p is β -times differentiable over $[-1, 1]^d$, then it belongs to $W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$.

We note that in principle the class $\mathcal{C}_{\beta,d} = W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$ is larger than the Nikol'skii class usually considered in density estimation [10] when $\beta \in \mathbb{N}$, but [Assumption 1](#) imposes a restriction on it. However, [Prop. 5](#) shows that the restriction imposed by [Assumption 1](#) is quite mild, since it the resulting space includes all probabilities that are β -times differentiable and strictly positive (it also contains probabilities that are not strictly positive but have some zeros, see [24]). We can now proceed to the main result of this work, which characterizes the approximation capabilities of PSD models.

Theorem 6 (Conciseness of PSD Approximation). *Let p satisfy [Assumption 1](#). Let $\varepsilon > 0$. There exists a Gaussian PSD model of dimension $m \in \mathbb{N}$, i.e., $\hat{p}_m(x) = f(x; A_m, X_m, \eta_m)$, with $A_m \in \mathbb{S}_+^m$ and $X_m \in \mathbb{R}^{m \times d}$ and $\eta_m \in \mathbb{R}_{++}^d$, such that*

$$\|p - \hat{p}_m\|_{L^2(\mathcal{X})} \leq \varepsilon, \quad \text{with } m = O(\varepsilon^{-d/\beta} (\log \frac{1}{\varepsilon})^{d/2}). \quad (11)$$

The proof of [Thm. 6](#) is reported in [Appendix D.3](#). According to the result, the number of base points needed for a PSD model to approximate a density up to precision ε depends on its smoothness (the smoother the better) and matches the bound $m = \varepsilon^{-d/\beta}$ that is optimal for function interpolation [18], corresponding to models allowing for negative weights, and is also optimal for convex combinations of oscillating kernels [36].

3.2 Learning a density with PSD models

In this section we study the capabilities of PSD models to estimate a density from n samples. Let $\mathcal{X} = (-1, 1)^d$ and let p be a probability on \mathcal{X} . Denote by x_1, \dots, x_n the samples independently and identically distributed according to p , with $n \in \mathbb{N}$. We consider a Gaussian PSD estimator $\hat{p}_{n,m} = f(x; \hat{A}, \tilde{X}, \eta)$ that is built on top of m additional points $\tilde{x}_1, \dots, \tilde{x}_m$, sampled independently and uniformly at random in \mathcal{X} . In particular, $\eta \in \mathbb{R}_{++}^d$, $\tilde{X} \in \mathbb{R}^{m \times d}$ is the base point matrix whose j -th row corresponds to the point \tilde{x}_j and $\hat{A} \in \mathbb{S}_+^m$ is trained as follows

$$\hat{A} = \operatorname{argmin}_{A \in \mathbb{S}_+^m} \int_{\mathcal{X}} f(x; A, \tilde{X}, \eta)^2 dx - \frac{2}{n} \sum_{i=1}^n f(x_i; A, \tilde{X}, \eta) + \lambda \|K^{1/2} A K^{1/2}\|_F^2, \quad (12)$$

where $K = K_{\tilde{X}, \tilde{X}, \eta}$. Note that the functional is constituted by two parts. The first two elements are an empirical version of $\|f - p\|_{L^2(\mathcal{X})}^2$ modulo a constant independent of f (and so not affecting the optimization problem), since x_i are identically distributed according to p and so $\frac{1}{n} \sum_{i=1}^n f(x_i) \approx \int f(x)p(x)dx$. The last term is a regularizer and corresponds to $\|K^{1/2} A K^{1/2}\|_F^2 = \operatorname{Tr}(A K A K)$, i.e. the Frobenius norm of $M = \sum_{i,j=1}^m A_{ij} \phi_\eta(\tilde{x}_i) \phi_\eta(\tilde{x}_j)$. The problem in (12) corresponds to a quadratic problem with a semidefinite constraint and can be solved using techniques such as Newton method [24] or first order dual methods [14]. We are now ready to state our result.

Theorem 7. *Let $n, m \in \mathbb{N}$, $\lambda > 0$, $\eta \in \mathbb{R}_{++}^d$ and p be a density satisfying [Assumption 1](#). With the definitions above, let $\hat{p}_{n,m}$ be the model $\hat{p}_{n,m}(x) = f(x; \hat{A}, \tilde{X}, \eta)$, with \hat{A} the minimizer of (12). Let $\eta = n^{\frac{2}{2\beta+d}} \mathbf{1}_d$ and $\lambda = n^{-\frac{2\beta+2d}{2\beta+d}}$. When $m \geq C' n^{\frac{d}{2\beta+d}} (\log n)^d \log(C'' n (\log n))$, the following holds with probability at least $1 - \delta$,*

$$\|p - \hat{p}_{n,m}\|_{L^2(\mathcal{X})} \leq C n^{-\frac{\beta}{2\beta+d}} (\log n)^{d/2}, \quad (13)$$

where constant C depends only on β, d and p and the constants C', C'' depend only on β, d .

The proof of [Thm. 7](#) is reported in [Appendix E.2](#). The theorem guarantees that under [Assumption 1](#), Gaussian PSD models can achieve the rate $O(n^{-\beta/(2\beta+d)})$ – that is optimal for the β -times differentiable densities – while admitting a concise representation. Indeed, it needs a number $m = O(n^{d/(2\beta+d)})$ of base points, matching the optimal rate in [36]. When $\beta \geq d$, a model with $m = O(n^{1/3})$ centers achieves optimal learning rates.

3.3 The Effect of compression

We have seen in the previous section that Gaussian PSD models achieve the optimal learning rates, with concise models. However, we have seen in the operations section that multiplying two PSD

models of m centers leads to a PSD model with m^2 centers. Here we study the effect of compression, to show that it is possible to obtain an ε -approximation of the product via a compressed model with $O(m \text{ polylog}(1/\varepsilon))$ centers. In the following theorem we analyze the effect in terms of the L^∞ distance on a domain $[-1, 1]^d$, induced by the compression, when using points taken independently and uniformly at random from the same domain.

Let $A \in \mathbb{S}_+^n$, $X \in \mathbb{R}^{n \times d}$, $\eta \in \mathbb{R}_{++}^d$, we want to study the compressibility of the PSD model $p(x) = f(x; A, X, \eta)$. Let $\tilde{X} \in \mathbb{R}^{m \times d}$ be the base point matrix whose j -rows corresponds to the point \tilde{x}_j with $\tilde{x}_1, \dots, \tilde{x}_m$ be sampled independently and uniformly at random from $[-1, 1]^d$. Denote by $\tilde{p}_m(x)$ the PSD model $\tilde{p}_m(x) = f(x; \tilde{A}, \tilde{X}, \eta)$ where \tilde{A} is the compression of A via (9). We have the following theorem.

Theorem 8 (Compression of Gaussian PSD models). *Let $\delta \in (0, 1]$, $\eta_+ = \max(1, \max_{i=1, \dots, d} \eta_i)$. When m satisfies*

$$m \geq O\left(\left(\eta_+^{1/2} \log \frac{\|A\|n}{\varepsilon}\right)^d \log \frac{1}{\delta}\right), \quad (14)$$

then the following holds with probability at least $1 - \delta$,

$$|p(x) - \tilde{p}_m(x)| \leq \varepsilon^2 + \varepsilon \sqrt{p(x)}, \quad \forall x \in [-1, 1]^d, \quad (15)$$

The proof of the theorem above is in [Appendix C.1](#). To understand its relevance, let \hat{p}_1 be a PSD model trained via (12) on n points sampled from p_1 and \hat{p}_2 trained from n points sampled from p_2 , where both p_1, p_2 satisfy [Assumption 1](#) for the same β and m, λ, η are chosen as [Thm. 7](#), in particular $m = n^{d/(2\beta+d)}$ and $\eta = \eta_+ \mathbf{1}_d$, $\eta_+ = n^{2/(2\beta+d)}$. Consider the model $\hat{p} = \hat{p}_1 \cdot \hat{p}_2$. By construction \hat{p} has $m^2 = n^{2d/(2\beta+d)}$ centers, since it is the pointwise product of \hat{p}_1, \hat{p}_2 (see [Prop. 2](#)) and approximates $p_1 \cdot p_2$ with error $\varepsilon = n^{-\beta/(2\beta+d)} \text{polylog}(n)$, since both \hat{p}_1, \hat{p}_2 are ε -approximators of p_1, p_2 . Instead, by compressing \hat{p} , we obtain an estimator \bar{p} , that according to [Thm. 8](#), achieves error ε with a number of center

$$m' = O(\eta_+^{d/2} \text{polylog}(1/\varepsilon)) = O(n^{d/(2\beta+d)} \text{polylog}(1/\varepsilon)) = O(m \text{polylog}(1/\varepsilon)). \quad (16)$$

Then \bar{p} approximates $p_1 \cdot p_2$ at the optimal rate $n^{-\beta/(2\beta+d)}$, but with a number of centers m' that is only $O(m \text{polylog}(n))$, instead of m^2 . This means that \bar{p} is essentially as good as if we learned it from n samples taken directly from $p_1 \cdot p_2$. This renders compression a suitable method to reduce the computational complexity of costly inference operations as the product rule.

4 Applications

PSD models are a strong candidate in a variety of probabilistic settings. On the one hand, they are computationally amenable to performing key operations such as sum and product rules, similarly to mixture models ([Sec. 2.1](#)). On the other hand, they are remarkably flexible and can approximate/learn (coincisely) a wide family of target probability densities ([Sec. 3](#)). Building on these properties, in this section we consider different possible applications of PSD models in practice.

4.1 PSD Models for Decision Theory

Decision theory problems (see e.g. [\[5\]](#) and references therein) can be formulated as a minimization

$$\theta_* = \operatorname{argmin}_{\theta \in \Theta} L(\theta) = \mathbb{E}_{x \sim p} \ell(\theta, x), \quad (17)$$

where ℓ is a loss function, Θ is the space of target parameters (decisions) and p is the underlying data distribution. When we can sample directly from p – e.g. in supervised or unsupervised learning settings – we can apply methods such as stochastic gradient descent to efficiently solve (17). However, in many applications, sampling from p is challenging or computationally unfeasible. This is for instance the case when p has been obtained via inference (e.g. it is the t -th estimate in a hidden Markov model, see [Sec. 4.3](#)) or it is fully known but has a highly complex form (e.g. the dynamics of a physical system). In contexts where sampling cannot be performed efficiently, it is advisable to consider alternative approaches. Here we propose a strategy to tackle (17) when p can be modeled (or well-approximated) by a PSD model. Our method hinges on the following result.

Proposition 9. Let $p(x) = f(x; A, X, \eta)$ with $X \in \mathbb{R}^{n \times d}$, $A \in \mathbb{S}_+^n$, $\eta \in \mathbb{R}_{++}^d$. Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ and define $c_{g,\eta}(z) = \int g(x) e^{-\eta \|x-z\|^2} dx$ for any $z \in \mathbb{R}^d$. Then

$$\mathbb{E}_{x \sim p} g(x) = \text{Tr}((A \circ K_{X,X,\eta/2})G) \quad \text{with} \quad G_{ij} = c_{g,2\eta}\left(\frac{x_i + x_j}{2}\right). \quad (18)$$

Thanks to [Prop. 9](#) we can readily compute several quantities related to a PSD model such as its mean $\mathbb{E}_p[x] = X^\top b$ with $b = (A \circ K_{X,\eta/2})\mathbf{1}_n$, its covariance or its characteristic function (see [Appendix F](#) for the explicit formulas and derivations). However, the result above is particularly useful to tackle the minimization in [\(17\)](#). In particular, since $\nabla L(\theta) = \mathbb{E}_{x \sim p} \nabla_\theta \ell(\theta, x)$, we can use [Prop. 9](#) to directly compute the gradient of the objective function: it is sufficient to know how to evaluate (or approximate) the integral $c_{\nabla_\theta \ell, \eta}(z) = \int \nabla_\theta \ell(\theta, x) e^{-\eta \|x-z\|^2} dx$ for any $\theta \in \Theta$ and $z \in \mathbb{R}^d$. Then, we can use first order optimization methods, such as gradient descent, to efficiently solve [\(17\)](#). Remarkably, this approach works well also when we approximate p with a PSD model \hat{p} . If ℓ is convex, since \hat{p} is non-negative, the resulting $\hat{L}(\theta) = \mathbb{E}_{x \sim \hat{p}} \ell(\theta, x)$ is still a convex functional (see also the discussion on structured prediction in [Sec. 4.2](#)). This is not the case if we use more general estimators of p that do not preserve non-negativity.

4.2 PSD Models for Estimating Conditional Probabilities

In supervised learning settings, one is typically interested in solving decision problems of the form $\min_{\theta \in \Theta} \mathbb{E}_{(x,y) \sim p} \ell(h_\theta(x), y)$ where p is a probability over the joint input-output space $\mathcal{X} \times \mathcal{Y}$ and $h_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ is a function parameterized by θ . It is well-known [see e.g. [34](#)] that the ideal solution of this problem is the θ_* such that for any $x \in \mathcal{X}$ the function $h_{\theta_*}(x) = \arg\min_{z \in \mathcal{Y}} \mathbb{E}_{y \sim p(\cdot|x)} \ell(z, y)$ is the minimizer with respect to $z \in \mathcal{Y}$ of the conditional expectation of $\ell(z, y)$ given x . This leads to target functions that capture specific properties of p , such as moments. For instance, when ℓ is the squared loss, $h_{\theta_*}(x) = \mathbb{E}_{y \sim p(\cdot|x)} y$ corresponds to the conditional expectation of y given x , while for ℓ the absolute value loss, h_{θ_*} recovers the conditional median of p .

In several applications, associating an input x to a single quantity $h_\theta(x)$ in output is not necessarily ideal. For instance, when $p(y|x)$ is multi-modal, estimating the mean or median might not yield useful predictions for the given task. Moreover, estimators of the form h_θ require access to the full input x to return a prediction, and therefore cannot be used when some features are missing (e.g. due to data corruption). In these contexts, an alternative viable strategy is to directly model the conditional probability. When using PSD models, conditional estimation can be performed in two steps, by first modeling the joint distribution $p(y, x) = f(y, x; A, [Y, X], (\eta, \eta'))$ (e.g. by learning it as suggested [Sec. 3](#)) and then use the operations in [Sec. 2.1](#) to condition it with respect to $x_0 \in \mathcal{X}$ as

$$p(y|x_0) = \frac{p(y, x_0)}{p(x_0)} = f(y; B, Y, \eta) \quad \text{with} \quad B = \frac{A \circ (K_{X,x_0,\eta'} K_{X,x_0,\eta'}^\top)}{c_{2\eta} \text{Tr}(A \circ (K_{X,x_0,\eta'} K_{X,x_0,\eta'}^\top) K_{Y,Y,\eta})}. \quad (19)$$

In case of data corruption, it is sufficient to first marginalize $p(y, x)$ on the missing variables and then apply [\(19\)](#). Below we discuss a few applications of the conditional estimator.

Conditional Expectation. Conditional mean embeddings [\[31\]](#) are a well-established tool to efficiently compute the conditional expectation $\mathbb{E}_{y \sim p(\cdot|x_0)} g(y)$ of a function $g : \mathcal{Y} \rightarrow \mathbb{R}$. However, although they enjoy good approximation properties [\[16\]](#), they do not guarantee the resulting estimator to take only non-negative values. In contrast, when p is a PSD model (or an approximation), we can apply [Prop. 9](#) to $p(\cdot|x_0)$ in [\(19\)](#) and evaluate the conditional expectation of any g for which we know how to compute (or approximate) the integral $c_{g,\eta}(z) = \int g(y) e^{-\eta \|y-z\|^2} dy$. In particular we have $\mathbb{E}_{y \sim p(\cdot|x)} g(y) = \text{Tr}((B \circ K_{Y,Y,\eta/2})G)$ with B as in [\(19\)](#) and G the matrix with entries $G_{ij} = c_{g,2\eta}\left(\frac{y_i + y_j}{2}\right)$. Remarkably, differently from conditional mean embeddings estimators, this strategy allows us to compute the conditional expectations also of functions g not in \mathcal{H}_η .

Structured Prediction. Structured prediction identifies supervised learning problems where the output space \mathcal{Y} has complex structures so that it is challenging to find good parametrizations for $h_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ [\[4, 19\]](#). In [\[7\]](#), a strategy was proposed to tackle these settings by first learning an approximation $\psi_\theta(z, x) \approx \mathbb{E}_{y \sim p(\cdot|x)} \ell(z, y)$ and then model $h_\theta(x) = \arg\min_{z \in \mathcal{Y}} \psi_\theta(z, x)$. However, the resulting function $\psi_\theta(\cdot, x)$ is not guaranteed to be convex, even when ℓ is convex. In contrast, by combining the conditional PSD estimator in [\(19\)](#) with the reasoning in [Sec. 4.1](#), we have a strategy that overcomes this issue: when $p(\cdot|x)$ is a PSD model approximating p , we can compute its gradient

Algorithm 1 PSD Hidden Markov Model

Input: Transition $\hat{\tau}(x_+, x) = f(x_+, x; B, [X_+, X], (\eta_+, \eta))$, initial $\hat{p}(x_0) = f(x_0; A_0, X_0, \eta_0)$ and observation $\hat{\omega}(y, x) = f(y, x; C, [Y, X'], (\eta_{obs}, \eta'))$ distributions. \tilde{X} as in [Prop. 10](#). $\tilde{\eta} = \frac{\eta(\eta' + \eta_+)}{\eta + \eta' + \eta_+}$.

$$\tilde{X}' = \frac{\eta}{\eta + \eta' + \eta_+} X \otimes \mathbf{1}_{nm} + \frac{\eta' + \eta_+}{\eta + \eta' + \eta_+} \mathbf{1}_n \otimes \tilde{X}, \quad \text{and} \quad \tilde{\eta}' = \frac{\eta' \eta_+}{\eta' + \eta_+}$$

For any new observation y_t :

$$\begin{aligned} C_t &= C \circ (K_{Y, y_t, \eta_{obs}}^\top K_{Y, y_t, \eta_{obs}}) && // \text{Partial evaluation } \hat{\omega}_t(x_+) = \hat{\omega}(y = y_t, x_+) \\ B_t &= (B \otimes A_{t-1}) \circ (\text{vec}(K_{X, \tilde{X}, \tilde{\eta}}) \text{vec}(K_{X, \tilde{X}, \tilde{\eta}})^\top) && // \text{Product } \hat{\beta}_t(x_+, x) = \hat{\tau}(x_+, x) \hat{p}(x|y_{1:t-1}) \\ D_t &= (I_n \otimes \mathbf{1}_{nm}^\top) (B_t \circ K_{\tilde{X}', \tilde{X}', \frac{\tilde{\eta}}{2}}) (I_n \otimes \mathbf{1}_{nm}) && // \text{Marginalization } \hat{\beta}_t(x_+) = \int \beta_t(x_+, x) dx \\ E_t &= (C_t \otimes D_t) \circ (\text{vec}(K_{X_+, X', \frac{\tilde{\eta}'}{2}}) \text{vec}(K_{X_+, X', \frac{\tilde{\eta}'}{2}})^\top) && // \text{Product } \hat{\pi}_t(x_+) = \hat{\omega}_t(x_+) \hat{\beta}_t(x_+) \\ A_t &= E_t / c_t, \text{ with } c_t = c_{2(\eta' + \eta_+)} \text{Tr}(E_t K_{\tilde{X}, \tilde{X}, \frac{\eta' + \eta_+}{2}}) && // \text{Normalization } \hat{p}(x_+ | y_{1:t}) = \frac{\hat{\pi}_t(x_+)}{\int \hat{\pi}_t(x_+) dx} \end{aligned}$$

Return $f(x_t; A_t, \tilde{X}, \eta' + \eta_+)$

$\mathbb{E}_{y \sim p(\cdot|x)} \nabla_z \ell(z, y)$ as mentioned in [Sec. 4.1](#) using [Prop. 9](#). Moreover, if $\ell(\cdot, y)$ is convex, the term $\mathbb{E}_{y \sim p(\cdot|x)} \ell(\cdot, y)$ is also convex, and we can use methods such as gradient descent to find $h(x)$ exactly.

Mode Estimation. When the output distribution $p(y|x)$ is multimodal, having access to an explicit form for the conditional density can be useful to estimate its modes. This problem is typically non-convex, yet, when the output y belongs to a small dimensional space (e.g. in classification or scalar-valued regression settings), efficient approximations exist (e.g. bisection).

4.3 Inference on Hidden Markov Models

We consider the problem of performing inference on hidden Markov models (HMM) using PSD models. Let $(x_t)_{t \in \mathbb{N}}$ and $(y_t)_{t \in \mathbb{N}}$ denote two sequences of states and observations respectively. For each $t \geq 1$, we denote by $x_{0:t} = x_0, \dots, x_t$ and $y_{1:t} = y_1, \dots, y_t$ and we assume that $p(x_t | x_{0:t-1}, y_{1:t-1}) = p(x_t | x_{t-1}) = \tau(x_t, x_{t-1})$ and $p(y_t | x_{0:t}, y_{1:t-1}) = p(y_t | x_t) = \omega(y_t, x_t)$ with $\tau : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_+$ and $\omega : \mathcal{Y} \times \mathcal{X} \rightarrow \mathbb{R}_+$ respectively the transition and observation functions.

Our goal is to infer the distribution of possible states x_t at time t , given all the observations $y_{1:t}$ and a probability $p(x_0)$ on the possible initial states. We focus on this goal for simplicity, but other forms of inferences are possible (e.g. estimating x_{m+t} , namely m steps into the future or the past). We assume that the transition and observation functions can be well approximated by PSD models $\hat{\tau}$ and $\hat{\omega}$ (e.g. by learning them or known a-priori for the problem's dynamics). Then, given a PSD model estimate $\hat{p}(x_0)$ of the initial state $p(x_0)$, we can recursively define the sequence of estimates

$$\hat{p}(x_t | y_{t:1}) = \frac{\hat{\tau}(y_t, x_t) \int \omega(x_t, x_{t-1}) \hat{p}(x_t | y_{1:t-1}) dx_{t-1}}{\int \hat{\tau}(y_t, x_t) \omega(x_t, x_{t-1}) \hat{p}(x_t | y_{1:t-1}) dx_t dx_{t-1}}. \quad (20)$$

Note that when $\hat{\tau}, \hat{\omega}, \hat{p}(x_0)$ correspond to the real transition, observation and initial state probability, the formula above yields the exact distribution $p(x_t | y_{1:t})$ over the states at time t (this follows directly by subsequent applications of Bayes' rule. See also e.g. [\[5\]](#)). If $\hat{\tau}, \hat{\omega}, \hat{p}(x_0)$ are PSD models, then each of the $\hat{p}(x_t | y_{t:1})$ is a PSD model recursively defined only in terms of the previous estimate and the operations introduced in [Sec. 2.1](#). In particular, we have the following result.

Proposition 10 (PSD Hidden Markov Models (HMM)). *Let $X_0 \in \mathbb{R}^{n_0 \times d}$, $X_+, X \in \mathbb{R}^{n \times d}$, $X' \in \mathbb{R}^{m \times d}$, $Y \in \mathbb{R}^{m \times d'}$, $A_0 \in \mathbb{S}_+^{n_0}$, $A \in \mathbb{S}_+^n$, $B \in \mathbb{S}_+^m$ and $\eta_0, \eta, \eta', \eta_+ \in \mathbb{R}_{++}^d$, $\eta_{obs} \in \mathbb{R}_{++}^{d'}$. Let*

$$\hat{\tau}(x_+, x) = f(x_+, x; B, [X_+, X], (\eta_+, \eta)), \quad \hat{\omega}(y, x) = f(y, x; C, [Y, X'], (\eta_{obs}, \eta')), \quad (21)$$

be approximate transition and observation functions. Then, given the initial state probability $\hat{p}(x_0) = f(x_0; A_0, X_0, \eta_0)$, for any $t \geq 1$, the estimate \hat{p} in [\(20\)](#) is a PSD model of the form

$$\hat{p}(x_t | y_{t:1}) = f(x_t; A_t, \tilde{X}, \eta' + \eta_+), \quad (22)$$

where $\tilde{X} = \frac{\eta'}{\eta' + \eta_+} X' \otimes \mathbf{1}_n + \frac{\eta_+}{\eta' + \eta_+} \mathbf{1}_m \otimes X_+$ and A_t is recursively obtained from A_{t-1} as in [Alg. 1](#).

Remark 4 (Sum-product Algorithm). *Eq. (20) is an instance of the so-called sum-product algorithm, a standard inference method for graphical models [5] (of which HMMs are a special case). The application of the sum-product algorithm relies mainly on sum and product rules for probabilities (as is the case for HMMs in (20)). Hence, according to Sec. 2.1, it is highly compatible with PSD models.*

5 Discussion

In this work we have shown that PSD models are a strong candidate in practical application related to probabilistic inference. They satisfy both requirements for an ideal probabilistic model: *i*) they perform exact sum and product rule in terms of efficient matrix operations; *ii*) we proved that they can concisely approximate a wide range of probabilities.

Future Directions. We identify three main directions for future work: *i*) when performing inference on large graphical models (see Remark 4) the multiplication of PSD models might lead to an inflation in the number of base points. Building on our compression strategy, we plan to further investigate low-rank approximations to mitigate this issue. *ii*) An interesting problem is to understand how to efficiently sample from a PSD model. A first answer to this open question was recently given in [15]. *iii*) The current paper has a purely theoretical and algorithmic focus. In the future, we plan to investigate the empirical behavior of PSD models on the applications introduced in Sec. 4. Related to this, we plan to develop a library for operations with PSD models and make it available to the community.

Acknowledgments. A.R. acknowledges the support of the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute) and the support of the European Research Council (grant REAL 947908). C.C. acknowledges the support of the Royal Society (grant SPREM RGS\R1\201149) and Amazon.com Inc. (Amazon Research Award – ARA 2020)

References

- [1] Robert A Adams and John JF Fournier. *Sobolev spaces*. Elsevier, 2003.
- [2] Jason Altschuler, Francis Bach, Alessandro Rudi, and Jonathan Niles-Weed. Massively scalable sinkhorn distances via the nyström method. *arXiv preprint arXiv:1812.05189*, 2018.
- [3] Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, 1950.
- [4] G. H. Bakir, T. Hofmann, B. Schölkopf, A. J. Smola, B. Taskar, and S. V. N. Vishwanathan. *Predicting Structured Data*. MIT Press, 2007.
- [5] Christopher M Bishop. *Machine learning and pattern recognition. Information Science and Statistics*. Springer, Heidelberg, 2006.
- [6] René Erlín Castillo and Humberto Rafeiro. *An introductory course in Lebesgue spaces*. Springer, 2016.
- [7] Carlo Ciliberto, Lorenzo Rosasco, and Alessandro Rudi. A general framework for consistent structured prediction with implicit loss embeddings. *Journal of Machine Learning Research*, 21(98):1–67, 2020.
- [8] Felipe Cucker and Ding Xuan Zhou. *Learning theory: an approximation theory viewpoint*, volume 24. Cambridge University Press, 2007.
- [9] Luc Devroye and Gábor Lugosi. *Combinatorial methods in density estimation*. Springer Science & Business Media, 2012.
- [10] Alexander Goldenshluger and Oleg Lepski. On adaptive minimax density estimation on r^d . *Probability Theory and Related Fields*, 159(3-4):479–543, 2014.
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

- [12] Lars Hörmander. *The analysis of linear partial differential operators II: Differential operators with constant coefficients*. Springer, 1990.
- [13] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *International Conference on Learning Representation (ICLR)*, 2014.
- [14] Ulysse Marteau-Ferey, Francis Bach, and Alessandro Rudi. Non-parametric models for non-negative functions. In *Advances in neural information processing systems*, 2020.
- [15] Ulysse Marteau-Ferey, Francis Bach, and Alessandro Rudi. Sampling from arbitrary functions via psd models. *arXiv preprint arXiv:2110.10527*, 2021.
- [16] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends(r) in Machine Learning Series*, 2017.
- [17] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [18] Erich Novak. *Deterministic and stochastic error bounds in numerical analysis*. Springer, 2006.
- [19] Sebastian Nowozin, Christoph H Lampert, et al. Structured learning and prediction in computer vision. *Foundations and Trends in Computer Graphics and Vision*, 2011.
- [20] Nicolò Pagliana, Alessandro Rudi, Ernesto De Vito, and Lorenzo Rosasco. Interpolation and learning with scale dependent kernels. *arXiv preprint arXiv:2006.09984*, 2020.
- [21] Emanuel Parzen. On estimation of a probability density function and mode. *The annals of mathematical statistics*, 33(3):1065–1076, 1962.
- [22] Christian Rieger and Barbara Zwicknagl. Sampling inequalities for infinitely smooth functions, with applications to interpolation and machine learning. *Advances in Computational Mathematics*, 32(1):103, 2010.
- [23] Luke Gervase Rogers. *A degree-independent Sobolev extension operator*. Yale University, 2004.
- [24] Alessandro Rudi, Ulysse Marteau-Ferey, and Francis Bach. Finding global minima via kernel approximations. *arXiv preprint arXiv:2012.11978*, 2020.
- [25] Alessandro Rudi and Lorenzo Rosasco. Generalization properties of learning with random features. In *NIPS*, pages 3215–3225, 2017.
- [26] Bernhard Schölkopf and Alexander J Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- [27] John Shawe-Taylor and N Cristianini. *An introduction to support vector machines and other kernel-based learning methods*, volume 204. Volume, 2000.
- [28] Winfried Sickel. Composition operators acting on sobolev spaces of fractional order—a survey on sufficient and necessary conditions. *Function spaces, differential operators and nonlinear analysis (Paseky nad Jizerou, 1995)*, pages 159–182, 1996.
- [29] Shashank Singh, Ananya Uppal, Boyue Li, Chun-Liang Li, Manzil Zaheer, and Barnabás Póczos. Nonparametric density estimation with adversarial losses. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 10246–10257. Curran Associates Inc., 2018.
- [30] Alex Smola, Arthur Gretton, Le Song, and Bernhard Schölkopf. A hilbert space embedding for distributions. In *International Conference on Algorithmic Learning Theory*, pages 13–31. Springer, 2007.
- [31] Le Song, Kenji Fukumizu, and Arthur Gretton. Kernel embeddings of conditional distributions: A unified kernel framework for nonparametric inference in graphical models. *IEEE Signal Processing Magazine*, 30(4):98–111, 2013.

- [32] Bharath Sriperumbudur, Kenji Fukumizu, Arthur Gretton, Aapo Hyvärinen, and Revant Kumar. Density estimation in infinite dimensional exponential families. *Journal of Machine Learning Research*, 18, 2017.
- [33] Elias M Stein and Guido Weiss. *Introduction to Fourier Analysis on Euclidean Spaces (PMS-32), Volume 32*. Princeton university press, 1971.
- [34] Ingo Steinwart and Andreas Christmann. *Support Vector Machines*. Information Science and Statistics. Springer New York, 2008.
- [35] Alexandre B Tsybakov. *Introduction to nonparametric estimation*. Springer Science & Business Media, 2008.
- [36] Paxton Turner, Jingbo Liu, and Philippe Rigollet. A statistical perspective on coresets for density estimation. In *International Conference on Artificial Intelligence and Statistics*, pages 2512–2520. PMLR, 2021.
- [37] Adrien Vacher, Boris Muzellec, Alessandro Rudi, Francis Bach, and Francois-Xavier Vialard. A dimension-free computational upper-bound for smooth optimal transport estimation. *arXiv preprint arXiv:2101.05380*, 2021.
- [38] Matt P Wand and M Chris Jones. *Kernel smoothing*. CRC press, 1994.
- [39] Holger Wendland. *Scattered Data Approximation*, volume 17. Cambridge University Press, 2004.
- [40] Christopher Williams and Matthias Seeger. Using the nyström method to speed up kernel machines. In *Proceedings of the 14th annual conference on neural information processing systems*, pages 682–688, 2001.
- [41] Vadim Yurinsky. *Sums and Gaussian vectors*. Springer, 1995.

Appendix

The appendix is organized as follows:

- [Appendix A](#) introduces notation and some key definitions and results that will be useful to prove the results in this work.
- [Appendix B](#) provides basic notation and definitions for working with linear operators between reproducing kernel Hilbert spaces.
- [Appendix C](#) discusses in detail the effects of compression introduced in [Sec. 2.2](#) in the main paper. In particular we study the approximation error incurred by a compressed model as a function of the number of base points used.
- [Appendix D](#) reports the proofs of the results in [Sec. 3.1](#) regarding the approximation properties of PSD models.
- [Appendix E](#) reports the proof of [Thm. 7](#) characterizing the learning capabilities of PSD models.
- [Appendix F](#) provides the derivations for the PSD models operations discussed in [Sec. 2.1](#) as well as some results directly related, namely [Prop. 9](#) and [Prop. 10](#).

A Notation and definitions

We introduce basic notation and review results that will be useful in the following.

Multi-index notation. Let $\alpha \in \mathbb{N}_0^d$, $x \in \mathbb{R}^d$ and f be an infinitely differentiable function on \mathbb{R}^d , we introduce the following notation

$$|\alpha| = \sum_{j=1}^d \alpha_j, \quad \alpha! = \prod_{j=1}^d \alpha_j!, \quad x^\alpha = \prod_{j=1}^d x_j^{\alpha_j}, \quad \partial^\alpha f = \frac{\partial^{|\alpha|} f}{\partial x_1^{\alpha_1} \cdots \partial x_d^{\alpha_d}}.$$

We introduce also the notation D^α that corresponds to the multivariate distributional derivative of order α and such that

$$D^\alpha f = \partial^\alpha f$$

for functions that are differentiable at least $|\alpha|$ times [1].

Fourier Transform. Given two functions $f, g : \mathbb{R}^d \rightarrow \mathbb{R}$ on some set \mathbb{R}^d , we denote by $f \cdot g$ the function corresponding to *pointwise product* of f, g , i.e.,

$$(f \cdot g)(x) = f(x)g(x), \quad \forall x \in \mathbb{R}^d.$$

Let $f, g \in L^1(\mathbb{R}^d)$ we denote the *convolution* by $f \star g$

$$(f \star g)(x) = \int_{\mathbb{R}^d} f(y)g(x-y)dy.$$

We now recall some basic properties, that will be used in the rest of the appendix.

Proposition A.1 (Basic properties of the Fourier transform [39], Chapter 5.2.).

- (a) *There exists a linear isometry $\mathcal{F} : L^2(\mathbb{R}^d) \rightarrow L^2(\mathbb{R}^d)$ satisfying*

$$\mathcal{F}[f] = \int_{\mathbb{R}^d} e^{-2\pi i \omega^\top x} f(x) dx \quad \forall f \in L^1(\mathbb{R}^d) \cap L^2(\mathbb{R}^d),$$

where $i = \sqrt{-1}$. The isometry is uniquely determined by the property in the equation above.

- (b) *Let $f \in L^2(\mathbb{R}^d)$, then $\|\mathcal{F}[f]\|_{L^2(\mathbb{R}^d)} = \|f\|_{L^2(\mathbb{R}^d)}$.*
- (c) *Let $f \in L^2(\mathbb{R}^d)$, $r > 0$ and define $f_r(x) = f(\frac{x}{r})$, $\forall x \in \mathbb{R}^d$, then $\mathcal{F}[f_r](\omega) = r^d \mathcal{F}[f](r\omega)$.*
- (d) *Let $f, g \in L^1(\mathbb{R}^d)$, then $\mathcal{F}[f \cdot g] = \mathcal{F}[f] \star \mathcal{F}[g]$.*

- (e) Let $\alpha \in \mathbb{N}_0^d$, $f, D^\alpha f \in L^2(\mathbb{R}^d)$, then $\mathcal{F}[D^\alpha f](\omega) = (2\pi i)^{|\alpha|} \omega^\alpha \mathcal{F}[f](\omega)$, $\forall \omega \in \mathbb{R}^d$.
- (f) Let $f \in L^1(\mathbb{R}^d) \cap L^2(\mathbb{R}^d)$, then $\|\mathcal{F}[f]\|_{L^\infty(\mathbb{R}^d)} \leq \|f\|_{L^1(\mathbb{R}^d)}$.
- (g) Let $f \in L^\infty(\mathbb{R}^d) \cap L^2(\mathbb{R}^d)$, then $\|f\|_{L^\infty(\mathbb{R}^d)} \leq \|\mathcal{F}[f]\|_{L^1(\mathbb{R}^d)}$.

Reproducing kernel Hilbert spaces for translation invariant kernels. We now list some important facts about reproducing kernel Hilbert spaces in the case of translation invariant kernels on \mathbb{R}^d . For this paragraph, we refer to [34, 39]. For the general treatment of positive kernels and Reproducing kernel Hilbert spaces, see [3, 34]. Let $v : \mathbb{R}^d \rightarrow \mathbb{R}$ such that its Fourier transform $\mathcal{F}[v] \in L^1(\mathbb{R}^d)$ and satisfies $\mathcal{F}[v](\omega) \geq 0$ for all $\omega \in \mathbb{R}^d$. Then, the following hold.

- (a) The function $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ defined as $k(x, x') = v(x - x')$ for any $x, x' \in \mathbb{R}^d$ is a positive kernel and is called *translation invariant kernel*.
- (b) The *reproducing kernel Hilbert space* (RKHS) \mathcal{H} and its norm $\|\cdot\|_{\mathcal{H}}$ are characterized by

$$\mathcal{H} = \{f \in L^2(\mathbb{R}^d) \mid \|f\|_{\mathcal{H}} < \infty\}, \quad \|f\|_{\mathcal{H}}^2 = \int_{\mathbb{R}^d} \frac{|\mathcal{F}[f](\omega)|^2}{\mathcal{F}[v](\omega)} d\omega, \quad (\text{A.1})$$

- (c) \mathcal{H} is a separable Hilbert space, whose inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ is characterized by

$$\langle f, g \rangle_{\mathcal{H}} = \int_{\mathbb{R}^d} \frac{\mathcal{F}[f](\omega) \overline{\mathcal{F}[g](\omega)}}{\mathcal{F}[v](\omega)} d\omega.$$

In the rest of the paper, when clear from the context we will simplify the notation of the inner product, by using $f^\top g$ for $f, g \in \mathcal{H}$, instead of the more cumbersome $\langle f, g \rangle_{\mathcal{H}}$.

- (d) The feature map $\phi : \mathbb{R}^d \rightarrow \mathcal{H}$ is defined as $\phi(x) = k(x - \cdot) \in \mathcal{H}$ for any $x \in \mathbb{R}^d$.
- (e) The functions in \mathcal{H} have the *reproducing property*, i.e.,

$$f(x) = \langle f, \phi(x) \rangle_{\mathcal{H}}, \quad \forall f \in \mathcal{H}, x \in \mathbb{R}^d, \quad (\text{A.2})$$

in particular $k(x', x) = \langle \phi(x'), \phi(x) \rangle_{\mathcal{H}}$ for any $x', x \in \mathbb{R}^d$.

We now introduce an important example of translation invariant kernel and the associated RKHS, that will be useful in our analysis.

Example 1 (Gaussian Reproducing kernel Hilbert space). Let $\eta \in \mathbb{R}_{++}^d$ and $k_\eta(x, x') = e^{-(x-x')^\top \text{diag}(\eta)(x-x')}$, for $x, x' \in \mathbb{R}^d$ be the Gaussian kernel with precision η . The function k_η is a translation invariant kernel, since $k_\eta(x, x') = v(x - x')$ with $v(z) = e^{-\|D^{1/2}z\|^2}$, $D = \text{diag}(\eta)$ and $\mathcal{F}[v](\omega) = c_\eta e^{-\pi^2 \|D^{-1/2}\omega\|^2}$, $c_\eta = \pi^{d/2} \det(D)^{-1/2}$, for $\omega \in \mathbb{R}^d$ is in $L^1(\mathbb{R}^d)$ and satisfies $\mathcal{F}[v](\omega) \geq 0$ for all $\omega \in \mathbb{R}^d$. The associated reproducing kernel Hilbert space \mathcal{H}_η is defined according to (A.1), with norm

$$\|f\|_{\mathcal{H}_\eta}^2 = \frac{1}{c_\eta} \int_{\mathbb{R}^d} |\mathcal{F}[f](\omega)|^2 e^{\pi^2 \|D^{-1/2}\omega\|^2} d\omega, \quad \forall f \in L^2(\mathbb{R}^d). \quad (\text{A.3})$$

The inner product and the feature map ϕ_η are defined as in the discussion above.

A.1 Sobolev spaces

Let $\beta \in \mathbb{N}$, $p \in [1, \infty]$ and let $\Omega \subseteq \mathbb{R}^d$ be an open set. The set $L^p(\Omega)$ denotes the set of p -integrable functions on Ω for $p \in [1, \infty)$ and that of the essentially bounded on Ω when $p = \infty$. The set $W_p^\beta(\Omega)$ denotes the Sobolev space, i.e., the set of measurable functions with their distributional derivatives up to β -th order belonging to $L^p(\Omega)$,

$$W_p^\beta(\Omega) = \{f \in L^p(\Omega) \mid \|f\|_{W_p^\beta(\Omega)} < \infty\}, \quad \|f\|_{W_p^\beta(\Omega)}^p = \sum_{|\alpha| \leq \beta} \|D^\alpha f\|_{L^p(\Omega)}^p, \quad (\text{A.4})$$

where D^α denotes the distributional derivative. In the case of $p = \infty$,

$$\|f\|_{W_\infty^\beta(\Omega)} = \max_{|\alpha| \leq \beta} \|D^\alpha f\|_{L^\infty(\Omega)} \quad (\text{A.5})$$

We now recall some basic results about Sobolev spaces that are useful for the proofs in this paper. First we start by recalling the restriction properties of Sobolev spaces. Let $\Omega \subseteq \Omega' \subseteq \mathbb{R}^d$ be two open sets. Let $\beta \in \mathbb{N}$ and $p \in [1, \infty]$. By definition of the Sobolev norm above we have

$$\|g|_{\Omega}\|_{W_p^s(\Omega)} \leq \|g\|_{W_p^s(\Omega')},$$

and so $g|_{\Omega} \in W_p^s(\Omega)$ for any $g \in W_p^s(\Omega')$. Now we recall the extension properties of Sobolev spaces.

Proposition A.2 (Stein total extension theorem, 5.24 in [1] or [23]). *Let Ω be a bounded open subset of \mathbb{R}^d with locally Lipschitz boundary [1]. For any measurable function $h : \Omega \rightarrow \mathbb{R}$, there exists a function $\tilde{h} : \mathbb{R}^d \rightarrow \mathbb{R}$, such that $\tilde{h}|_{\Omega} = h$ almost everywhere on Ω and for any $\beta \geq 0$ and $p \in [1, \infty]$, the condition $\|h\|_{W_p^\beta(\Omega)} < \infty$ implies $\|\tilde{h}\|_{W_p^\beta(\mathbb{R}^d)} \leq C_{\Omega, \beta, p} \|h\|_{W_p^\beta(\Omega)}$.*

Corollary A.3. *Let $\mathcal{X} \subset \mathbb{R}^d$ be a non-empty open set with Lipschitz boundary. Let $\beta \in \mathbb{N}$, $p \in [1, \infty]$. Then for any function $f \in W_p^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$ there exists an extension \tilde{f} on \mathbb{R}^d , i.e. a function $\tilde{f} \in W_p^\beta(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$ such that*

$$f = \tilde{f}|_{\mathcal{X}} \text{ a.e. on } \mathcal{X}, \quad \|\tilde{f}\|_{L^\infty(\mathbb{R}^d)} \leq C \|f\|_{L^\infty(\mathcal{X})}, \quad \|\tilde{f}\|_{W_p^\beta(\mathbb{R}^d)} \leq C' \|f\|_{L^\infty(\mathcal{X})}. \quad (\text{A.6})$$

The constant C depends only on \mathcal{X} , d , and the constant C' only on \mathcal{X} , β , d , p

Proposition A.4. *Let \mathcal{X} be an open bounded set with Lipschitz boundary. Let f be a function that is m times differentiable on the closure of \mathcal{X} . Then there exists a function $\tilde{f} \in W_p^m(\mathcal{X}) \cap L^\infty(\mathcal{X})$ for any $p \in [1, \infty]$, such that $\tilde{f} = f$ on \mathcal{X} .*

Proof. A function f that is m -times differentiable on the closure of \mathcal{X} belongs also to $W_\infty^m(\mathcal{X})$ since each derivative up to order m is continuous and the set \mathcal{X} is bounded. Then f satisfies also $f \in W_p^m(\mathcal{X}) \cap L^\infty(\mathcal{X})$ since $W_\infty^m(\mathcal{X}) \subset L^\infty(\mathcal{X})$ and $W_\infty^m(\mathcal{X}) \subset W_p^\beta(\mathcal{X})$, by construction, for bounded \mathcal{X} and any $p \in [1, \infty]$. \square

The following proposition provides a useful characterization of the space $W_2^\beta(\mathbb{R}^d)$

Proposition A.5 (Characterization of the Sobolev space $W_2^k(\mathbb{R}^d)$, [39]). *Let $k \in \mathbb{N}$. The norm of the Sobolev space $\|\cdot\|_{W_2^k(\mathbb{R}^d)}$ is equivalent to the following norm*

$$\|f\|_{W_2^k(\mathbb{R}^d)}^2 = \int_{\mathbb{R}^d} |\mathcal{F}[f](\omega)|^2 (1 + \|\omega\|^2)^k d\omega, \quad \forall f \in L^2(\mathbb{R}^d) \quad (\text{A.7})$$

and satisfies

$$\frac{1}{(2\pi)^{2k}} \|f\|_{W_2^k(\mathbb{R}^d)}^2 \leq \|f\|_{W_2^k(\mathbb{R}^d)}^2 \leq 2^{2k} \|f\|_{W_2^k(\mathbb{R}^d)}^2, \quad \forall f \in L^2(\mathbb{R}^d) \quad (\text{A.8})$$

Moreover, when $k > d/2$, then $W_2^k(\mathbb{R}^d)$ is a reproducing kernel Hilbert space.

Proof. Consider first the seminorm $|g|_{W_2^t(\mathbb{R}^d)}^2 = \sum_{|\alpha| \leq t} \|D^\alpha g\|_{L^2(\mathbb{R}^d)}^2$. We have that $\|g\|_{W_2^k(\mathbb{R}^d)}^2 = \sum_{t=0}^k |g|_{W_2^t(\mathbb{R}^d)}^2$. Now let $0 \leq t \leq k$. By using the properties of the Fourier transform (in particular, the Plancherel theorem and the transform of a distributional derivative Prop. A.1) we have that

$$|g|_{W_2^t(\mathbb{R}^d)}^2 = \sum_{|\alpha|=t} \|D^\alpha g\|_{L^2(\mathbb{R}^d)}^2 = \sum_{|\alpha|=t} \|(2\pi i)^\alpha \omega^\alpha \mathcal{F}[g](\omega)\|_{L^2(\mathbb{R}^d)}^2 = \int \sum_{|\alpha|=t} (2\pi\omega)^{2\alpha} |\mathcal{F}[g](\omega)|^2 d\omega.$$

Now note that, by the multinomial theorem, $\|2\pi\omega\|^{2t} = (2\pi\omega_1^2 + \dots + 2\pi\omega_d^2)^t = \sum_{|\alpha|=t} \binom{t}{\alpha} (2\pi\omega)^\alpha$. Since $1 \leq \binom{t}{\alpha} \leq 2^t$ for any $\alpha \in \mathbb{N}_0^d$, $|\alpha| = t$, then $2^{-t} \|2\pi\omega\|^{2t} \leq \sum_{|\alpha|=t} (2\pi\omega)^{2\alpha} \leq \|2\pi\omega\|^{2t}$, so

$$|g|_{W_2^t(\mathbb{R}^d)}^2 \leq (2\pi)^{2t} \int \|\omega\|^{2t} |\mathcal{F}[g](\omega)|^2 d\omega \leq 2^t |g|_{W_2^t(\mathbb{R}^d)}^2. \quad (\text{A.9})$$

Since, $(1 + \|\omega\|^2)^k = \sum_{t=0}^k \binom{k}{t} \|\omega\|^{2t}$ and so $\sum_{t=0}^k \|\omega\|^{2t} \leq (1 + \|\omega\|^2)^k \leq 2^k \sum_{t=0}^k \|\omega\|^{2t}$, then

$$\|g\|_{W_2^k(\mathbb{R}^d)}^2 = \sum_{t=0}^k |g|_{W_2^t(\mathbb{R}^d)}^2 \leq (2\pi)^{2k} \sum_{t=0}^k \int \|\omega\|^{2t} |\mathcal{F}[g](\omega)|^2 d\omega \leq (2\pi)^{2k} \int (1 + \|\omega\|^2)^k |\mathcal{F}[g](\omega)|^2,$$

moreover

$$\int (1+\|\omega\|^2)^k |\mathcal{F}[g](\omega)|^2 \leq 2^k \sum_{t=0}^k (2\pi)^{2t} \int \|\omega\|^{2t} |\mathcal{F}[g](\omega)|^2 \leq 2^{2k} \sum_{t=0}^k |g|_{W_2^k(\mathbb{R}^d)}^2 = 2^{2k} \|g\|_{W_2^k(\mathbb{R}^d)}.$$

To conclude, we recall that when $k > d/2$ the space $W_2^k(\mathbb{R}^d)$ endowed with the equivalent norm $\|\cdot\|_{W_2^k(\mathbb{R}^d)}$ is a reproducing kernel Hilbert space [39]. \square

B Useful linear operators in RKHS

Consider the space $\mathcal{G}_\eta = \mathcal{H}_\eta \otimes \mathcal{H}_\eta = \{v \otimes v' \mid v, v' \in \mathcal{H}_\eta\}$ endowed with the inner product $\langle u \otimes u', v \otimes v' \rangle_{\mathcal{G}_\eta} = \langle u, v \rangle_{\mathcal{H}_\eta} \langle u', v' \rangle_{\mathcal{H}_\eta}$ for any $u, u', v, v' \in \mathcal{H}_\eta$. Denote by vec the unitary map that maps the Hilbert-Schmidt operators on \mathcal{H}_η in vectors in \mathcal{H}_η^\otimes . In particular, for any $u, v \in \mathcal{H}_\eta$, we have $\text{vec}(uv^\top) = v \otimes u$, moreover for any $M, M' : \mathcal{H}_\eta \rightarrow \mathcal{H}_\eta$ with finite Hilbert-Schmidt norm

$$\langle \text{vec}(M), \text{vec}(M') \rangle_{\mathcal{G}_\eta} = \text{Tr}(M^* M'), \quad \langle \text{vec}(M), v \otimes u \rangle_{\mathcal{G}_\eta} = v^\top M u. \quad (\text{B.1})$$

Now denote by ψ_η the feature map $\psi_\eta(x) = \phi_\eta(x) \otimes \phi_\eta(x)$ for any $x \in \mathbb{R}^d$. We define the operator $Q \in \mathbb{S}_+(\mathcal{G}_\eta)$ and the vectors $\hat{v}, v \in \mathcal{G}_\eta$ as follows

$$Q = \int_{\mathcal{X}} \psi_\eta(x) \psi_\eta(x)^\top dx, \quad \hat{v} = \frac{1}{n} \sum_{i=1}^n \psi_\eta(x_i), \quad v = \int_{\mathcal{X}} \psi_\eta(x) p(x) dx. \quad (\text{B.2})$$

Define the operator $S : \mathcal{G}_\eta \rightarrow L^2(\mathcal{X})$ as

$$Sf = \langle \psi_\eta(\cdot), f \rangle_{\mathcal{G}_\eta} \in L^2(\mathcal{X}), \quad \forall f \in \mathcal{G}_\eta \quad (\text{B.3})$$

$$S^* \alpha = \int \alpha(x) \psi_\eta(x) dx \in \mathcal{G}_\eta, \quad \forall \alpha \in L^2(\mathcal{X}). \quad (\text{B.4})$$

Note, in particular, that Q and v are characterized by

$$Q = S^* S, \quad v = S^* p. \quad (\text{B.5})$$

Given $\tilde{x}_1, \dots, \tilde{x}_m \in \mathbb{R}^d$ define the operator $\tilde{Z} : \mathcal{H}_\eta \rightarrow \mathbb{R}^m$ as $\tilde{Z} = (\phi_\eta(\tilde{x}_1)^\top, \dots, \phi_\eta(\tilde{x}_m)^\top)^\top$, in particular we have

$$\begin{aligned} \tilde{Z}u &= (\phi_\eta(\tilde{x}_1)^\top u, \dots, \phi_\eta(\tilde{x}_m)^\top u), \forall u \in \mathcal{H}_\eta \\ \tilde{Z}^* \alpha &= \sum_{i=1}^m \phi_\eta(\tilde{x}_i) \alpha_i, \forall \alpha \in \mathbb{R}^m, \end{aligned} \quad (\text{B.6})$$

In particular, note that for any $A \in \mathbb{R}^{m \times m}$

$$\tilde{Z}^* A \tilde{Z} = \sum_{i,j=1}^m A_{i,j} \phi_\eta(\tilde{x}_i) \phi_\eta(\tilde{x}_j)^\top, \quad \tilde{Z} \tilde{Z}^* = K_{\tilde{X}, \tilde{X}, \eta}. \quad (\text{B.7})$$

Given $\tilde{Z} : \mathcal{H}_\eta \rightarrow \mathbb{R}^m$, define the associated projection operator $\tilde{P} : \mathcal{H}_\eta \rightarrow \mathcal{H}_\eta$ on the range of the adjoint \tilde{Z}^* . In particular, note that

$$\tilde{P} = \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z} \quad (\text{B.8})$$

indeed, since $\tilde{Z} \tilde{Z}^* = K_{\tilde{X}, \tilde{X}, \eta}$ and it is invertible for any $\eta \in \mathbb{R}_{++}^d$, then

$$\tilde{P}^2 = \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z} \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z}^* = \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} K_{\tilde{X}, \tilde{X}, \eta} K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z} = \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z} = \tilde{P}. \quad (\text{B.9})$$

and

$$\tilde{P} \tilde{Z}^* = \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z} \tilde{Z}^* = \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} K_{\tilde{X}, \tilde{X}, \eta} = \tilde{Z}^*. \quad (\text{B.10})$$

and analogously $\tilde{Z} \tilde{P} = \tilde{Z}$. This implies also that $\tilde{P} \phi_\eta(\tilde{x}_i) = \phi_\eta(\tilde{x}_i)$ for any $i = 1, \dots, m$.

C Compression of a PSD model

Let $\mathcal{X} \subset \mathbb{R}^d$ be an open set with Lipschitz boundary, contained in the hypercube $[-R, R]^d$ with $R > 0$. Given $\tilde{x}_1, \dots, \tilde{x}_m \in \mathcal{X}$ be m points in $[-R, R]^d$. Define the base point matrix $\tilde{X} \in \mathbb{R}^{m \times d}$ to be the matrix whose j -th row is the point \tilde{x}_j . The following result holds. We introduce the so called *fill distance* [39]

$$h = \max_{x \in [-R, R]^d} \min_{z \in \tilde{X}} \|x - z\|, \quad (\text{C.1})$$

In the next lemma we specialize Theorem 4.5 of [22], to obtain explicit constants in terms of R and of our η . In particular, we identify the scale parameter $\sigma = \min(R, 1/\sqrt{\max_i \eta_i})$. This is interesting since it shows the effect of the precision η of the kernel (if it was a Gaussian probability, its variance would scale exactly as $1/\sqrt{\eta}$).

Lemma C.1 (Norm of functions with scattered zeros). *Let $T = (-R, R)^d$ and $\eta \in \mathbb{R}_{++}^d$. Let $u \in \mathcal{H}_\eta$ satisfying $u(\tilde{x}_1) = \dots = u(\tilde{x}_m) = 0$. There exists three constants c, C', C'' depending only on d (and in particular, independent from $R, \eta, u, \tilde{x}_1, \dots, \tilde{x}_m$), such that, when $h \leq \sigma/C'$, then,*

$$\|u\|_{L^\infty(T)} \leq Cq_\eta e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} \|u\|_{\mathcal{H}_\eta}, \quad (\text{C.2})$$

with $q_\eta = \det(\frac{1}{\eta_+} \text{diag}(\eta))^{-1/4}$ and $\sigma = \min(R, \frac{1}{\sqrt{\eta_+}})$ and $\eta_+ = \max_{i=1, \dots, d} \eta_i$.

Proof. By Theorem 4.3 of [22] there exists two constants B_d, B'_d depending only on d (and independent from R, \tilde{x}_j), such that for any $k \in \mathbb{N}, k > d/2 + 1$ and $u \in W_2^k(T)$ satisfying $u(\tilde{x}_1) = \dots = u(\tilde{x}_m) = 0$, the following holds,

$$\|u\|_{L^\infty(T)} \leq \frac{B_d^k k^{k-d/2}}{k!} h^{k-d/2} |u|_{W_2^k(T)}, \quad (\text{C.3})$$

when $kh \leq R/B'_d$. Here the seminorm $|u|_{W_2^k(T)}$, by using the multinomial notation recalled in Appendix A, corresponds to $|u|_{W_2^k(T)}^2 = \sum_{|\alpha|=k} \|D^\alpha u\|_{L^2(T)}^2$. Our goal is to apply the result above to $f \in \mathcal{H}_\eta$. First we recall that $\mathcal{H}_\eta \subset W_2^k(\mathbb{R}^d)$ for any k [22]. Then, since $T \subset \mathbb{R}^d$, we have $|f|_{W_2^k(T)}^2 = \sum_{|\alpha|=k} \|D^\alpha f\|_{L^2(T)}^2 \leq \sum_{|\alpha|=k} \|D^\alpha f\|_{L^2(\mathbb{R}^d)}^2 = |f|_{W_2^k(\mathbb{R}^d)}^2$. Then, using (A.9) we have $|f|_{W_2^k(\mathbb{R}^d)}^2 \leq (2\pi)^{2k} \int \|\omega\|^{2k} |\mathcal{F}[f](\omega)|^2 d\omega$. Now, denote by D the matrix $D = \text{diag}(\eta)$ and $c_\eta = \pi^{d/2} \det(D)^{-1/2}$. By the characterization of the norm $\|\cdot\|_{\mathcal{H}_\eta}$ in terms of the Fourier transform reported in Example 1, we have

$$\|f\|_{W_2^k(\mathbb{R}^d)}^2 \leq (2\pi)^{2k} \int \|\omega\|^{2k} |\mathcal{F}[f](\omega)|^2 d\omega \quad (\text{C.4})$$

$$= \int \|\omega\|^{2k} c_\eta (2\pi)^{2k} e^{-\pi^2 \|D^{-1/2}\omega\|^2} \frac{1}{c_\eta} e^{\pi^2 \|D^{-1/2}\omega\|^2} |\mathcal{F}[f](\omega)|^2 d\omega \quad (\text{C.5})$$

$$\leq c_\eta (2\pi)^{2k} \sup_{t \in \mathbb{R}^d} \|t\|^{2k} e^{-\pi^2 \|D^{-1/2}t\|^2} \int \frac{1}{c_\eta} e^{\pi^2 \|D^{-1/2}\omega\|^2} |\mathcal{F}[f](\omega)|^2 d\omega \quad (\text{C.6})$$

$$= \|f\|_{\mathcal{H}_\eta}^2 c_\eta (2\pi)^{2k} \sup_{t \in \mathbb{R}^d} \|t\|^{2k} e^{-\pi^2 \|D^{-1/2}t\|^2}. \quad (\text{C.7})$$

Now, since $\sup_{z \in \mathbb{R}^d} \|z\|^{2k} e^{-\|z\|^2} = \sup_{r \geq 0} r^{2k} e^{-r^2} = k! e^{-k} \leq k!$ and $\|D\| = \max_i \eta_i = \eta_+$,

$$\sup_{t \in \mathbb{R}^d} \|t\|^{2k} e^{-\pi^2 \|D^{-1/2}t\|^2} = \sup_{z \in \mathbb{R}^d} \left\| \frac{1}{\pi} D^{\frac{1}{2}} z \right\|^{2k} e^{-\|z\|^2} \leq \frac{\|D\|^k}{\pi^{2k}} \sup_{z \in \mathbb{R}^d} \|z\|^{2k} e^{-\|z\|^2} = \frac{\eta_+^k k!}{\pi^{2k}}. \quad (\text{C.8})$$

Then,

$$\|f\|_{W_2^k(\mathbb{R}^d)} \leq \|f\|_{\mathcal{H}_\eta} c_\eta^{1/2} (4\eta_+)^{k/2} \sqrt{k!}. \quad (\text{C.9})$$

By plugging the bound above in (C.3), we obtain that when $k, h_{\tilde{X}, R}$ satisfy $kh_{\tilde{X}, R} \leq RB'_d$ then,

$$\|f\|_{L^\infty(T)} \leq C \frac{(C_2 kh)^{k-d/2}}{\sqrt{k!}} \|f\|_{\mathcal{H}_\eta}, \quad (\text{C.10})$$

where $C_2 = 2\sqrt{\eta_+}B_d$, and $C = (4B_d^2\pi)^{d/4} \det(\eta/\eta_+)^{-1/4}$. Let now $C_3 = 1/\max(B'_d, 2B_d)$. Assume that $h \leq \frac{C_3}{2+d} \min(R, 1/\sqrt{\eta_+})$ and set $k = \lfloor s \rfloor$ and $s = \frac{C_3}{h} \min(R, 1/\sqrt{\eta_+})$. Note first, that with this choice of h and k we satisfy $hk \leq R/B'_d$ so we can apply Theorem 4.3 of [22]. Moreover, by construction $d/2 + 1 \leq \frac{s}{2} \leq s - 1 \leq \lfloor s \rfloor \leq s$. Then $C_2 \lfloor s \rfloor h \leq 1$ and $\lfloor s \rfloor - d/2 \geq 0$, so $(C_2 h \lfloor s \rfloor)^{\lfloor s \rfloor - d/2} \leq 1$. Moreover $\frac{1}{\sqrt{k!}} \leq e^{-\frac{k}{2} \log \frac{k}{2}}$ so we have

$$\|f\|_{L^\infty(T)} \leq C(C_2 h \lfloor s \rfloor)^{\lfloor s \rfloor - d/2} e^{-\frac{1}{2} \lfloor s \rfloor \log \frac{1}{2} \lfloor s \rfloor} \|f\|_{\mathcal{H}_\eta} \quad (\text{C.11})$$

$$\leq C e^{-\frac{s-1}{2} \log \frac{s-1}{2}} \|f\|_{\mathcal{H}_\eta} \leq C e^{-\frac{s}{4} \log \frac{s}{4}} \|f\|_{\mathcal{H}_\eta}. \quad (\text{C.12})$$

The final result is obtained by writing $s/4 = c\sigma/h$ with $\sigma = \min(R, 1/\sqrt{\eta_+})$ and $c = C_3/4$ and by writing the assumption on h as $h \leq \sigma/C'$ with $C' = (2+d)/C_3$. \square

Lemma C.2 (Lemma 3, page 28 [20]). *Let $\mathcal{X} \subset \mathbb{R}^d$ with non-zero volume. Let \mathcal{H} be a reproducing kernel Hilbert space on \mathcal{X} , associated to a continuous uniformly bounded feature map $\phi : \mathcal{X} \rightarrow \mathcal{H}$. Let $A : \mathcal{H} \rightarrow \mathcal{H}$ be a bounded linear operator. Then,*

$$\sup_{x \in \mathcal{X}} \|A\phi(x)\|_{\mathcal{H}} \leq \sup_{\|f\| \leq 1} \|A^*f\|_{C(\mathcal{X})}. \quad (\text{C.13})$$

*In particular, if $\mathcal{X} \subset \mathbb{R}^d$ is a non-empty open set, then $\sup_{x \in \mathcal{X}} \|A\phi(x)\|_{\mathcal{H}} \leq \|A^*f\|_{L^\infty(\mathcal{X})}$.*

Proof. We recall the variational characterization of the norm $\|\cdot\|_{\mathcal{H}}$ in terms of the inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ as $\|v\|_{\mathcal{H}} = \sup_{\|f\| \leq 1} \langle f, v \rangle_{\mathcal{H}}$. We have the following

$$\sup_{x \in \mathcal{X}} \|A\phi(x)\|_{\mathcal{H}} = \sup_{x \in \mathcal{X}, \|f\| \leq 1} \langle f, A\phi(x) \rangle_{\mathcal{H}} \leq \sup_{x \in \mathcal{X}, \|f\| \leq 1} |\langle A^*f, \phi(x) \rangle_{\mathcal{H}}| \quad (\text{C.14})$$

$$= \sup_{\|f\| \leq 1} \sup_{x \in \mathcal{X}} |(A^*f)(x)| = \sup_{\|f\| \leq 1} \|A^*f\|_{C(\mathcal{X})}. \quad (\text{C.15})$$

Finally, note that when $\mathcal{X} \subset \mathbb{R}^d$ is a non-empty open set $\|A^*f\|_{C(\mathcal{X})} = \|A^*f\|_{L^\infty(\mathcal{X})}$, since $A^*f \in \mathcal{H}$ and all the functions in \mathcal{H} are continuous and bounded due to the continuity of ϕ . \square

Theorem C.3 (Approximation properties of the projection). *Let $R > 0, \eta \in \mathbb{R}_{++}^d, m \in \mathbb{N}$. Let $\mathcal{X} \subseteq T = (-R, R)^d$ be a non-empty open set and let $\tilde{x}_1, \dots, \tilde{x}_m$ be a set of distinct points. Let $h > 0$ be the fill distance associated to the points w.r.t T (defined in (C.1)). Let $\tilde{P} : \mathcal{H}_\eta \rightarrow \mathcal{H}_\eta$ be the associated projection operator (see definition in (B.8)). There exists three constants c, C, C' , such that, when $h \leq \sigma/C'$,*

$$\sup_{x \in \mathcal{X}} \|(I - \tilde{P})\phi_\eta(x)\|_{\mathcal{H}_\eta} \leq Cq_\eta e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}}. \quad (\text{C.16})$$

Here $q_\eta = \det(\frac{1}{\eta_+} \text{diag}(\eta))^{-1/4}$ and $\sigma = \min(R, \frac{1}{\sqrt{\eta_+}})$, $\eta_+ = \max_i \eta_i$. The constants c, C', C'' depend only on d and, in particular, are independent from $R, \eta, \tilde{x}_1, \dots, \tilde{x}_m$.

Proof. We first recall some basic properties of the projection operator $\tilde{P} : \mathcal{H}_\eta \rightarrow \mathcal{H}_\eta$ on the span of $\phi_\eta(\tilde{x}_1), \dots, \phi_\eta(\tilde{x}_m)$, defined in (B.8). By construction $\tilde{P}\phi_\eta(\tilde{x}_i) = \phi_\eta(\tilde{x}_i)$ is of rank m for any $i = 1, \dots, m$. Now note that for any $f \in \mathcal{H}_\eta$, the function $(\tilde{P}f)(\tilde{x}_i) = f(\tilde{x}_i)$, indeed, by the reproducing property of \mathcal{H}_η

$$(\tilde{P}u)(\tilde{x}_i) = \left\langle \tilde{P}u, \phi_\eta(\tilde{x}_i) \right\rangle_{\mathcal{H}_\eta} = \left\langle u, \tilde{P}\phi_\eta(\tilde{x}_i) \right\rangle_{\mathcal{H}_\eta} = \langle u, \phi_\eta(\tilde{x}_i) \rangle_{\mathcal{H}_\eta} = u(\tilde{x}_i). \quad (\text{C.17})$$

Then $(f - \tilde{P}f)(\tilde{x}_i) = 0$ for any $i = 1, \dots, m$. By Lemma C.1, we know that there exist three constants c, C, C' depending only on d such that when $h \leq \sigma/C'$ we have that the following holds $\|u\|_{L^\infty(T)} \leq Cq_\eta e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}}$, for any $u \in \mathcal{H}_\eta$ such that $u(\tilde{x}_1) = \dots = u(\tilde{x}_m) = 0$. Since, for any $f \in \mathcal{H}_\eta$, we have that $f - \tilde{P}f$ belongs to \mathcal{H}_η and satisfies such property, we can apply Lemma C.1 with $u = (I - \tilde{P})f$, obtaining, under the same assumption on h ,

$$\|(I - \tilde{P})f\|_{L^\infty(T)} \leq Cq_\eta e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} \|f\|_{\mathcal{H}_\eta}, \quad \forall f \in \mathcal{H}_\eta, \quad (\text{C.18})$$

where we used the fact that $\|(I - \tilde{P})f\|_{\mathcal{H}_\eta} \leq \|I - \tilde{P}\| \|f\|_{\mathcal{H}_\eta}$ and $\|I - \tilde{P}\| \leq 1$, since P is a projection operator and so also $I - P$ satisfies this property. The final result is obtained by applying [Lemma C.2](#) with $A = I - \tilde{P}$, from which we have

$$\sup_{x \in \mathcal{X}} \|(I - \tilde{P})\phi(x)\|_{\mathcal{H}} \leq \sup_{\|f\| \leq 1} \|(I - \tilde{P})f\|_{L^\infty(T)} \leq \sup_{\|f\| \leq 1} Cq_\eta e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} \|f\|_{\mathcal{H}_\eta} \quad (\text{C.19})$$

$$= Cq_\eta e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}}. \quad (\text{C.20})$$

□

Theorem C.4 (Compression of a PSD model). *Let $\eta \in \mathbb{R}_{++}^d$ and let $M \in \mathbb{S}_+(\mathcal{H}_\eta)$. Let \mathcal{X} be an open bounded subset with Lipschitz boundary of the cube $[-R, R]^d$, $R > 0$. Let $\tilde{x}_1, \dots, \tilde{x}_m \in \mathcal{X}$ and \tilde{X}_m be the base point matrix whose j -rows are the points \tilde{x}_j with $j = 1, \dots, m$. Consider the model $p = f(\cdot; M, \phi_\eta)$ and the compressed model $\tilde{p} = f(\cdot; A_m, \tilde{X}_m, \eta)$ with*

$$A_m = K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z} M \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1}, \quad (\text{C.21})$$

where $\tilde{Z} : \mathcal{H}_\eta \rightarrow \mathbb{R}^m$ is defined in [\(B.6\)](#) in terms of \tilde{X}_m . Let h be the fill distance (defined in [\(C.1\)](#)) associated to the points $\tilde{x}_1, \dots, \tilde{x}_m$. There exist three constants c, C, C' depending only on d such that, when $h \leq \sigma/C'$, with $\sigma = \min(R, 1/\sqrt{\eta_+})$, $\eta_+ = \max_{i=1, \dots, d} \eta_i$, $q_\eta = \det(\frac{1}{\eta_+} \text{diag}(\eta))^{-1/4}$, then

$$|p(x) - \tilde{p}(x)| \leq 2Cq_\eta \sqrt{\|M\|p(x)} e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} + C^2 q_\eta^2 \|M\| e^{-\frac{2c\sigma}{h} \log \frac{c\sigma}{h}}, \quad \forall x \in \mathcal{X}. \quad (\text{C.22})$$

Proof. Let $\varepsilon \in (0, 1/e]$ and $\eta = \tau \mathbf{1}_d \in \mathbb{R}^d$ with $\tau > 0$. Consider the projection operator $\tilde{P} : \mathcal{H}_\eta \rightarrow \mathcal{H}_\eta$ associated to the points $\tilde{x}_1, \dots, \tilde{x}_m$, defined in [\(B.8\)](#). Note that the adjoint \tilde{Z}^* has range equal to $\text{span}\{\phi_\eta(\tilde{x}_1), \dots, \phi_\eta(\tilde{x}_m)\}$ and that by construction $\tilde{Z} M_\varepsilon \tilde{Z}^* \in \mathbb{S}_+^m$ and so $A_m \in \mathbb{S}_+^m$.

Step 1. Error induced by a projection. By the reproducing property $k_\eta(x, x') = \phi_\eta(x)^\top \phi_\eta(x')$ (see [Example 1](#) and [\(A.2\)](#)) and the fact that $\tilde{P} = \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z}$, (see [\(B.8\)](#)), then, for any $x \in \mathbb{R}^d$

$$f(x; A_m, \tilde{X}_m, \eta) = \sum_{i,j=1}^m (A_m)_{i,j} k_\eta(x, \tilde{x}_i) k_\eta(x, \tilde{x}_j) \quad (\text{C.23})$$

$$= \phi_\eta(x)^\top \left(\sum_{i,j=1}^m (A_m)_{i,j} \phi_\eta(\tilde{x}_i) \phi_\eta(\tilde{x}_j)^\top \right) \phi_\eta(x) \quad (\text{C.24})$$

$$= \phi_\eta(x)^\top \tilde{Z}^* A_m \tilde{Z} \phi_\eta(x) \quad (\text{C.25})$$

$$= \phi_\eta(x)^\top \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z} M_\varepsilon \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z} \phi_\eta(x) \quad (\text{C.26})$$

$$= \phi_\eta(x)^\top \tilde{P} M \tilde{P} \phi_\eta(x) \quad (\text{C.27})$$

$$= f(x; \tilde{P} M \tilde{P}, \phi_\eta). \quad (\text{C.28})$$

This implies that, for all $x \in \mathbb{R}^d$ the following holds

$$f(x; A_m, \tilde{X}_m, \eta) - f(x; M, \phi_\eta) = f(x; \tilde{P} M \tilde{P}, \phi_\eta) - f(x; M, \phi_\eta) \quad (\text{C.29})$$

$$= \phi_\eta(x)^\top (\tilde{P} M \tilde{P} - M_\varepsilon) \phi_\eta(x). \quad (\text{C.30})$$

Step 2. Bounding $|\phi_\eta(x)^\top (\tilde{P} M_\varepsilon \tilde{P} - M) \phi_\eta(x)|$. Now, consider that

$$\tilde{P} M \tilde{P} - M_\varepsilon = M(I - \tilde{P}) + (I - \tilde{P})M + (I - \tilde{P})M(I - \tilde{P}). \quad (\text{C.31})$$

Since $|a^\top A B A a| \leq \|A a\|_{\mathcal{H}}^2 \|B\|$ and $a^\top A B a \leq \|A a\|_{\mathcal{H}} \|B^{1/2}\| \|B^{1/2} a\|_{\mathcal{H}}$, for any a in a Hilbert space \mathcal{H} and for A, B bounded symmetric linear operators with $B \in \mathbb{S}_+(\mathcal{H})$, by bounding the terms of the equation above, we have for any $x \in \mathbb{R}^d$,

$$|\phi_\eta(x)^\top (\tilde{P} M \tilde{P} - M) \phi_\eta(x)| \leq 2\|(I - \tilde{P})\phi_\eta(x)\|_{\mathcal{H}_\eta} \|M\|^{1/2} \|M^{1/2} \phi_\eta(x)\|_{\mathcal{H}_\eta} \quad (\text{C.32})$$

$$+ \|(I - \tilde{P})\phi_\eta(x)\|_{\mathcal{H}_\eta}^2 \|M\| \quad (\text{C.33})$$

$$= 2c_M^{1/2} f(x; M, \phi_\eta)^{1/2} C e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} + c_M \tilde{u}(x)^2, \quad (\text{C.34})$$

where $c_M = \|M\|$ and we denoted by $\tilde{u}(x)$ the quantity $\tilde{u}(x) = \|(I - \tilde{P})\phi_\eta(x)\|_{\mathcal{H}_\eta}$ and we noted that $\|M^{1/2}\phi_\eta(x)\|_{\mathcal{H}_\eta}^2 = \phi_\eta(x)^\top M \phi_\eta(x) = f(x; M, \phi_\eta)$.

Step 3. Bounding \tilde{u} . Now, by [Thm. C.3](#) we have that when the fill distance h (defined in [\(C.1\)](#)) satisfies $h \leq \sigma/C'$ with $\sigma = \min(R, 1/\sqrt{\tau})$, then

$$\|\tilde{u}\|_{L^\infty(\mathcal{X})} = \sup_{x \in \mathcal{X}} \|(I - \tilde{P})\phi_\eta(x)\|_{\mathcal{H}_\eta} \leq C e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}}. \quad (\text{C.35})$$

with c, C, C' depending only on d . \square

C.1 Proof of [Thm. 8](#)

[Thm. 8](#) is a corollary of the next theorem, considering that $1/\sigma \leq (1 + \eta_+)^{1/2}$ and moreover $\det(\frac{1}{\eta_+} \text{diag}(\eta)) = \prod_{i=1}^d \eta_i/\eta_+ \leq 1$, $\text{Tr}(AK_{X,X,\eta}) \leq \|A\| \text{Tr}(K_{X,X,\eta})$ since both $A, K_{X,X,\eta} \in \mathbb{S}_+^n$, and by construction $\text{Tr}(K_{X,X,\eta}) = n$, then $c_{A,\eta} \leq \|A\|n$.

Theorem C.5. *Let $A \in \mathbb{S}_+^n$, $X \in \mathbb{R}^{n \times d}$, $\eta \in \mathbb{R}_{++}^d$. Let $\tilde{x}_1, \dots, \tilde{x}_m$ be sampled independently and uniformly at random from $[-1, 1]^d$. Let $\delta \in (0, 1]$, $\eta_+ = \max(1, \max_{i=1, \dots, d} \eta_i)$. When m satisfies $m \geq Q' \eta_+^{d/2} (\log \frac{Q\|A\|n}{\varepsilon})^d \log(\frac{Q''(1+\eta_+)}{\delta} \log \frac{Q\|A\|n}{\varepsilon})$, then the following holds with probability at least $1 - \delta$,*

$$|p(x) - \tilde{p}(x)| \leq \varepsilon^2 + \varepsilon \sqrt{p(x)}, \quad \forall x \in [-1, 1]^d, \quad (\text{C.36})$$

Here the three constants Q, Q', Q'' depend only on d .

Proof. First let us rewrite $f(\cdot; A, X, \eta)$ in the equivalent form $f(\cdot; M, \phi_\eta)$ with $M \in \mathbb{S}_+(\mathcal{H}_\eta)$ defined as $M = \sum_{ij=1}^n A_{ij} \phi_\eta(x_i) \phi_\eta(x_j)$. In particular, by the cyclicity of the trace

$$\text{Tr}(M) = \sum_{ij=1}^n A_{i,j} \phi_\eta(x_i)^\top \phi_\eta(x_j) = \sum_{ij=1}^n A_{i,j} k_\eta(x_i, x_j) = \text{Tr}(AK_{X,X,\eta}). \quad (\text{C.37})$$

The proof of this theorem is an application of the approximation result in [Thm. C.4](#) to the model $f(\cdot; M, \phi_\eta)$ where we use as compression points, the points $\tilde{x}_1, \dots, \tilde{x}_m$ sampled independently and uniformly at random from $[-1, 1]^d$.

The result of the theorem depends on the fill distance h , defined in [\(C.1\)](#), and associated to the points $\tilde{x}_1, \dots, \tilde{x}_m$. Let c, C, C' be the constants depending only on d from [Thm. C.4](#). To apply [Thm. C.4](#) we have to guarantee that $h \leq \sigma/C'$ with $\sigma = \min(1, 1/\sqrt{\eta_+})$, in particular, choosing h such that $h \leq \min(c, 1/C')\sigma/(e \log(2Cc_{A,\eta}/\varepsilon))$ guarantees that $h \leq \sigma/C'$ and by applying the theorem we have, for all $x \in [-1, 1]^d$

$$|p(x) - \tilde{p}(x)| \leq 2Cq_\eta \sqrt{\|M\|p(x)} e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} + Cq_\eta^2 \|M\| e^{-\frac{2c\sigma}{h} \log \frac{c\sigma}{h}}, \quad (\text{C.38})$$

$$\leq 2Cc_{A,\eta} \sqrt{p(x)} e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} + C^2 c_{A,\eta}^2 e^{-\frac{2c\sigma}{h} \log \frac{c\sigma}{h}}, \quad (\text{C.39})$$

with $q_\eta = \det(\frac{1}{\eta_+} \text{diag}(\eta))^{-1/4}$ where in the last step we used the fact that $\|M\| \leq \text{Tr}(M) = \text{Tr}(AK_{X,X,\eta})$ and so $\|M\|q_\eta^2 \leq \text{Tr}(AK_{X,X,\eta})q_\eta^2 = c_{A,\eta}^2$. Note now, that by the choice we made for h , we have $\log(c\sigma/h) \geq 1$ and so that $e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} \leq \varepsilon/(2Cc_{A,\eta})$. This implies

$$|p(x) - \tilde{p}(x)| \leq \varepsilon \sqrt{p(x)} + \varepsilon^2. \quad (\text{C.40})$$

The final result is obtained by controlling the number of points m such that h satisfy the required bound in high probability. By, e.g. Lemma 12, page 19 of [\[37\]](#) and the fact that $[-1, 1]^d$ is a convex set, we have that there exists two constants C_1, C_2 depending only on d such that $h \leq C_1 m^{-1/d} (\log(C_2 m/\delta))^{1/d}$, with probability at least $1 - \delta$. In particular m satisfying

$$m \geq \left(\frac{eC_1}{\min(c, 1/C')} \frac{1}{\sigma} \log \frac{2Cc_{A,\varepsilon}}{\varepsilon} \right)^d \log \frac{C_2 m}{\delta} \quad (\text{C.41})$$

guarantees that $C_1 m^{-1/d} (\log \frac{C_2 m}{\delta})^{1/d} \leq \frac{\sigma \min(c, 1/C')}{\epsilon} \log \frac{2C c_{A,\eta}}{\epsilon}$. Note that, given $A \geq e, B \geq e$, the inequality $m \geq B \log(Am)$ is satisfied by $m \geq 2B \log 2AB$, indeed $\log A \geq \log \log A$ and $\log 2B \geq \log \log 2B$ and so, when $m = m_0 = 2B \log 2AB$ we have

$$B \log(Am_0) = B \log(2AB \log(2AB)) \quad (\text{C.42})$$

$$= B \log A + B \log 2B + B \log \log A + B \log \log 2B \quad (\text{C.43})$$

$$\leq 2B \log A + 2B \log 2B = 2B \log 2AB = m_0, \quad (\text{C.44})$$

and moreover $m - B \log(Am)$ is increasing for $m \geq B$. Then, to satisfy (C.41) we choose $m \geq 2B \log(2AB)$ with $B = \left(\frac{eC_1}{\min(c, 1/C')} \frac{1}{\sigma} \log \frac{2C c_{A,\eta}}{\epsilon} \right)^d$ and $A = \frac{C_2}{\delta}$, in particular

$$m = Q \left(\frac{1}{\sigma} \log \frac{Q' c_{A,\eta}}{\epsilon} \right)^d \log \left(\frac{Q''}{\delta \sigma} \log \frac{Q' c_{A,\eta}}{\epsilon} \right). \quad (\text{C.45})$$

with $Q = 2d \left(\frac{eC_1}{\min(c, 1/C')} \right)^d$, with $Q' = 2C, Q'' = 2C_2 Q^d$. \square

D Approximation of a probability via a PSD model

In this section we prove [Prop. 5](#) and [Thm. 6](#).

D.1 Proof of [Prop. 5](#)

Lemma D.1. *Let $\beta > 0$ and \tilde{p} be a continuous strictly positive function such that $\tilde{p} \in W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$, with $\beta > 0$, be a function such that $p > 0$. Let \mathcal{X} be a compact set with Lipschitz boundary. Then there exist a function in $\tilde{f} \in W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$ such that $p(x) = f(x)^2$ for all $x \in \mathcal{X}$.*

Proof. First, let $\tilde{p} \in W_2^\beta(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$ be the extension of p on all \mathbb{R}^d (see [Cor. A.3](#)), i.e. $\tilde{p} = p$ on \mathcal{X} and $\|\tilde{p}\|_{W_2^\beta(\mathbb{R}^d)} \leq C \|p\|_{W_2^\beta(\mathcal{X})}$ for a fixed constant C . Let $c = \inf_{x \in \mathcal{X}} p(x)$ and $C = \sup_{x \in \mathcal{X}} p(x)$, we have that $c > 0$ since \mathcal{X} is compact and p is continuous. Note that $g(z) = \sqrt{z}$ is C^∞ on the open interval $(0, +\infty)$. Let $u \in C^\infty(\mathbb{R})$ be the bump function that is identically 0 on $J = (\infty, c/2] \cup [2C, \infty)$ and identically 1 on the interval $I = [c, C]$. Then the function $h(z) = u(z)g(z)$ is identically 0 on J and $h(z) = \sqrt{z}$ on I and $h \in C^\infty(\mathbb{R})$ since $h = 0$ on J and both $u, g \in C^\infty([c/2, 2C])$. Now, denote by f the function $f(x) = h(p(x))$ for all $x \in \mathbb{R}^d$. Since $p(x) \in I$ for any $x \in \mathcal{X}$, we have $f(x) = h(p(x)) = \sqrt{p(x)}$ for any $x \in \mathcal{X}$. Moreover, by Theorem 1, page 8, in [\[28\]](#), $f \in W_2^\beta(\mathbb{R}^d)$, since $h \in C^\infty(\mathbb{R}) \subset C^{\beta+1}(\mathbb{R})$ and $p \in W_2^\beta(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$. The fact that $f \in L^\infty(\mathbb{R}^d)$ derive from the fact that it is the composition of a compactly supported smooth function h and an $L^\infty(\mathbb{R}^d)$ function p . The proof is concluded by taking \tilde{f} to be the restriction of f to \mathcal{X} . \square

Lemma D.2 ([\[24\]](#) Corollary 2, page 23). *Let \mathcal{X} be an open bounded subset of \mathbb{R}^d with Lipschitz boundary. Let p be a probability density that is $\beta + 2$ -times differentiable on the closure of \mathcal{X} , with $\beta > 0$. The zeros of p are isolated points with strictly positive Hessian and their number is finite. Moreover there are no zeros of p on the boundary. Then there exist $q \in \mathbb{N}$ and q functions $f_1 \dots f_q \in W^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$, such that*

$$p(x) = \sum_{i=1}^q f_i(x)^2, \quad \forall x \in \mathcal{X}. \quad (\text{D.1})$$

Proof. Let \tilde{p} be the $\beta + 2$ -times differentiable extension to \mathbb{R}^d of p (via the Whitney extension theorem [\[12\]](#)), i.e. $\tilde{p} = p$ on the closure of \mathcal{X} . We apply [\[24\]](#) Corollary 2, page 23 on \tilde{p} , obtaining q functions $f_1, \dots, f_q \in C^\beta(\mathbb{R}^d)$ such that $\tilde{p}(x) = \sum_i f_i(x)^2$ for all $x \in \mathcal{X}$. The result is obtained by applying [Prop. A.4](#) on the restrictions f_1, \dots, f_q on \mathcal{X} . \square

Now we are ready to prove [Prop. 5](#). We restate here fore convenience.

Proposition 5 (Generality of [Assumption 1](#)). *The assumption above is satisfied by*

- (a) any probability density $p \in W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$ and strictly positive on $[-1, 1]^d$,
- (b) any exponential model $p(x) = e^{-v(x)}$ with $v \in W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$,
- (c) any mixture model of Gaussians or, more generally, of exponential models from (b),
- (d) any p that is $\beta + 2$ -times differentiable on $[-1, 1]^d$, with a finite set of zeros, all in $(-1, 1)^d$, and with positive definite Hessian in each zero. E.g. $p(x) \propto x^2 e^{-x^2}$.

Moreover if p is β -times differentiable over $[-1, 1]^d$, then it belongs to $W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$.

Proof. Let $\mathcal{X} = (-1, 1)^d$. The case (a) is proven in [Lemma D.1](#). Note that (b) and (c) are a subclass of (a). For the case (b), let $\tilde{v} \in W_2^\beta(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$ be the extension of v (see [Cor. A.3](#)), by Theorem 1, page 8, in [\[28\]](#) the function $e^{\tilde{v}} - 1 \in W_2^\beta(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$ since $\exp(\cdot) - 1$ is analytic and 0 in 0. Let $q = (e^{\tilde{v}} - 1)|_{\mathcal{X}}$, $q \in W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$ and so also $g = q + 1$ belongs to $W_2^\beta(\mathcal{X}) \cap L^\infty(\mathcal{X})$, since \mathcal{X} is a bounded set. Finally note that $g = e^v$ on \mathcal{X} and $\min_{x \in \mathcal{X}} g(x) = \min_{x \in \mathcal{X}} e^{-v(x)} \geq e^{-\|v\|_{L^\infty(\mathcal{X})}} > 0$, so it satisfies the point (a). The point (c) is a consequence of (b) indeed if $p = \sum_{i=1}^t \alpha_i e^{-v_i}$ and each v_i satisfies (b), then $e^{-v_i} = \sum_{j=1}^{q_i} f_{i,j}^2$ with $f_{i,j} \in W_2^\beta(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$, so $p = \sum_{i=1}^t \sum_{j=1}^{q_i} g_{i,j}^2$, with $g_{i,j} = \sqrt{\alpha_i} f_{i,j} \in W_2^\beta(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$. Finally, (d) is proven in [Lemma D.2](#). \square

D.2 Additional results required to prove [Thm. 6](#)

We now focus on proving the result in [Thm. 6](#). To this end, we first prove some preliminary result that will be useful in the following.

Let $S \subseteq \mathbb{R}^d$. We recall the definition of the function $\mathbf{1}_S$, that is $\mathbf{1}_S(\omega) = 1$ for any $\omega \in S$ and $\mathbf{1}_S(\omega) = 0$ for any $\omega \notin S$. Define moreover,

$$g(x) = \frac{1}{V_d} \|x\|^{-d} J_{d/2}(2\pi\|x\|) J_{d/2}(4\pi\|x\|) \quad (\text{D.2})$$

where $J_{d/2}$ is the Bessel function of the first kind of order $d/2$ and $V_d = \int_{\|x\| \leq 1} dx = \frac{\pi^{d/2}}{\Gamma(d/2+1)}$.

Lemma D.3. *The function g defined above satisfies $g \in L^1(\mathbb{R}^d) \cap L^2(\mathbb{R}^d)$ and $\int g(x) dx = 1$. Moreover, for any $\omega \in \mathbb{R}^d$, we have*

$$\mathbf{1}_{\{\|\omega\| < 1\}}(\omega) \leq \mathcal{F}[g](\omega) \leq \mathbf{1}_{\{\|\omega\| \leq 3\}}(\omega). \quad (\text{D.3})$$

Proof. In this proof we will use the notation about the convolution and the Fourier transform in [Prop. A.1](#). Define $b(x) = \|x\|^{-d/2} J_{d/2}(2\pi\|x\|)$ where $J_{d/2}$ is the Bessel function of the first kind of order $d/2$. Note that $b \in L^2(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$, since there exists a constant $c > 0$ $|J_{d/2}(z)| \leq c \min(z^{d/2}, z^{-1/2})$ for any $z \geq 0$ [\[33\]](#). Moreover note that the Fourier transform of b is $\mathcal{F}[b](\omega) = \mathbf{1}_{\{\|\omega\| < 1\}}$ (see [\[33\]](#), [Thm. 4.15](#), page 171). Define now $g(x) = \frac{1}{V_d} b(x)b(2x) = \frac{1}{V_d} \|x\|^{-d} J_{d/2}(2\pi\|x\|) J_{d/2}(4\pi\|x\|)$. Note that $g \in L^1(\mathbb{R}^d)$ since

$$\|g\|_{L^1(\mathbb{R}^d)} = \|b(\cdot)b(2\cdot)\|_{L^1(\mathbb{R}^d)} \leq \|b(\cdot)\|_{L^2(\mathbb{R}^d)} \|b(2\cdot)\|_{L^2(\mathbb{R}^d)} < \infty, \quad (\text{D.4})$$

and analogously

$$\|g\|_{L^2(\mathbb{R}^d)} = \|b(\cdot)b(2\cdot)\|_{L^2(\mathbb{R}^d)} \leq \|b(\cdot)\|_{L^\infty(\mathbb{R}^d)} \|b(2\cdot)\|_{L^2(\mathbb{R}^d)} < \infty. \quad (\text{D.5})$$

By the properties of the Fourier transform, we have $\mathcal{F}[g] = \frac{1}{V_d} \mathcal{F}[b] \star \mathcal{F}[b(2\cdot)] = \frac{1}{V_d} \int \mathbf{1}_{\{\|\omega\| \leq 1\}} \mathbf{1}_{\{\|\omega - z\| \leq 2\}} dz$. Note that for any $\omega \in \mathbb{R}^d$, since $\mathbf{1}_{\{\|\omega - z\| \leq 2\}} \leq 1$,

$$\mathcal{F}[g](\omega) = \frac{1}{V_d} \int \mathbf{1}_{\{\|z\| \leq 1\}} \mathbf{1}_{\{\|\omega - z\| \leq 2\}} dz \leq \frac{1}{V_d} \int \mathbf{1}_{\{\|z\| \leq 1\}} dz \leq 1. \quad (\text{D.6})$$

Now, note that when $\|\omega\|, \|z\| \leq 1$, then $\|\omega - z\| \leq \|\omega\| + \|z\| \leq 2$. So we have $\mathbf{1}_{\{\|\omega\| \leq 1\}} \mathbf{1}_{\{\|z\| \leq 1\}} \mathbf{1}_{\{\|\omega - z\| \leq 2\}} = \mathbf{1}_{\{\|z\| \leq 1\}}$. Then

$$\mathbf{1}_{\{\|\omega\| \leq 1\}} \mathcal{F}[g](\omega) = \frac{1}{V_d} \int \mathbf{1}_{\{\|\omega\| \leq 1\}} \mathbf{1}_{\{\|z\| \leq 1\}} \mathbf{1}_{\{\|\omega - z\| \leq 2\}} dz \quad (\text{D.7})$$

$$= \frac{1}{V_d} \int \mathbf{1}_{\{\|z\| \leq 1\}} dz = \mathbf{1}_{\{\|\omega\| \leq 1\}}. \quad (\text{D.8})$$

Moreover note that for all $\|\omega\| > 3, \|z\| \leq 1$ we have $\|\omega - z\| \geq \|\omega\| - \|z\| > 2$, then $\mathbf{1}_{\{\|\omega\|>3\}} \mathbf{1}_{\{\|z\|\leq 1\}} \mathbf{1}_{\{\|\omega-z\|\leq 2\}} = 0$. So for any $\|\omega\| > 3$

$$\mathbf{1}_{\{\|\omega\|>3\}} \mathcal{F}[g](\omega) = \frac{1}{V_d} \int \mathbf{1}_{\{\|\omega\|>3\}} \mathbf{1}_{\{\|z\|\leq 1\}} \mathbf{1}_{\{\|\omega-z\|\leq 2\}} dz = 0. \quad (\text{D.9})$$

To conclude $\int g(x) dx = \int g(x) e^{-2\pi i \omega^\top 0} dx = \mathcal{F}[g](0) = 1$. \square

Theorem D.4. Let $\beta > 0, q \in \mathbb{N}$. Let $f_1, \dots, f_q \in W_2^\beta(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$ and denote by p the function $p = \sum_{i=1}^q f_i^2$. Let $\varepsilon \in (0, 1]$ and let $\eta \in \mathbb{R}_{++}^d$. Denote by $\eta_0 = \min_{j=1, \dots, d} \eta_j$. Let ϕ_η be the feature map of the Gaussian kernel with bandwidth η and let \mathcal{H}_η be the associated RKHS. Then there exists $M_\varepsilon \in \mathbb{S}_+(\mathcal{H}_\eta)$ with $\text{rank}(M_\varepsilon) \leq q$, such that

$$\|f(\cdot; M_\varepsilon, \phi_\eta) - p(\cdot)\|_{L^r(\mathbb{R}^d)} \leq \varepsilon, \quad \text{Tr}(M_\varepsilon) \leq C|\eta|^{1/2} (1 + \varepsilon^2 \exp(\frac{C'}{\eta_0} \varepsilon^{-\frac{2}{\beta}})), \quad (\text{D.10})$$

for all $r \in [1, 2]$, where $|\eta| = \det(\text{diag}(\eta))$ and C, C' depend only on $\beta, d, \|f_i\|_{W_2^\beta(\mathbb{R}^d)}, \|f_i\|_{L^\infty(\mathbb{R}^d)}$.

Proof. Let $t > 0$ (to be set later) and let g be defined according to (D.2). Define $g_t(x) = t^{-d} g(x/t)$. Given the properties of g in Lemma D.3, we have that $\int g_t(x) dx = 1, g_t \in L^1(\mathbb{R}^d) \cap L^2(\mathbb{R}^d)$, that $\mathcal{F}[g_t](\omega) = \mathcal{F}[g](t\omega)$ and so that $|\mathcal{F}[g_t](\omega)| = |\mathcal{F}[g](t\omega)| \leq \mathbf{1}_{\{t\|\omega\|\leq 3\}}(\omega)$. Moreover we have that $|1 - \mathcal{F}[g_t](\omega)| = |1 - \mathcal{F}[g](t\omega)| \leq \mathbf{1}_{\{t\|\omega\|\geq 1\}}(\omega)$.

Now, note that $\int (1 + \|\omega\|^2)^\beta |\mathcal{F}[f](\omega)|^2 d\omega \leq 2^{2\beta} \|f\|_{W_2^\beta(\mathbb{R}^d)}^2$, as discussed in Prop. A.5.

Step 1. Bounding $\|f - f \star g_t\|_{L^2(\mathbb{R}^d)}$.

Since, we have seen that $|1 - \mathcal{F}[g_t](\omega)| \leq \mathbf{1}_{\{t\|\omega\|\geq 1\}}$, then for any $f \in W_2^\beta(\mathbb{R}^d)$ we have

$$\|f - f \star g_t\|_{L^2(\mathbb{R}^d)}^2 = \|\mathcal{F}[f] - \mathcal{F}[f \star g_t]\|_{L^2(\mathbb{R}^d)}^2 = \|\mathcal{F}[f](1 - \mathcal{F}[g_t])\|_{L^2(\mathbb{R}^d)}^2 \quad (\text{D.11})$$

$$= \int |1 - \mathcal{F}[g](t\omega)|^2 |\mathcal{F}[f](\omega)|^2 d\omega \leq \int_{t\|\omega\|\geq 1} |\mathcal{F}[f](\omega)|^2 d\omega \quad (\text{D.12})$$

$$= \int_{t\|\omega\|\geq 1} (1 + \|\omega\|^2)^{-\beta} (1 + \|\omega\|^2)^\beta |\mathcal{F}[f](\omega)|^2 d\omega \quad (\text{D.13})$$

$$\leq \sup_{t\|\omega\|\geq 1} (1 + \|\omega\|^2)^{-\beta} \int (1 + \|\omega\|^2)^\beta |\mathcal{F}[f](\omega)|^2 d\omega \quad (\text{D.14})$$

$$= 2^{2\beta} \|f\|_{W_2^\beta(\mathbb{R}^d)}^2 \frac{(2t)^{2\beta}}{(1 + t^2)^\beta} \leq \|f\|_{W_2^\beta(\mathbb{R}^d)}^2 (2t)^{2\beta}. \quad (\text{D.15})$$

Step 2. Bounding $\|f \star g_t\|_{\mathcal{H}_\eta}$.

However, the function $f \star g_t$ belongs to \mathcal{H}_η for any $\eta \in \mathbb{R}_{++}^d$. Indeed, as discussed in Example 1, we have that $\|u\|_{\mathcal{H}_\eta}$ is characterized as

$$\|u\|_{\mathcal{H}_\eta}^2 = c_\eta \int |\mathcal{F}[u](\omega)|^2 e^{\pi^2 \omega^\top \text{diag}(\eta)^{-1} \omega} d\omega,$$

and $u \in \mathcal{H}_\eta$ iff $\|u\|_{\mathcal{H}_\eta} < \infty$, with $c_\eta = \pi^{-d/2} \det(\text{diag}(\eta))^{1/2}$. Now, let $\eta_0 = \min_{i=1 \dots d} \eta_i$, since we have seen that $|\mathcal{F}[g_t](\omega)| \leq \mathbf{1}_{\{t\|\omega\|\leq 3\}}(\omega)$, then we have that

$$\|f \star g_t\|_{\mathcal{H}_\eta}^2 = c_\eta \int |\mathcal{F}[f](\omega) \mathcal{F}[g](t\omega)|^2 e^{\omega^\top \text{diag}(\eta)^{-1} \omega} d\omega \quad (\text{D.16})$$

$$\leq c_\eta \int |\mathcal{F}[f](\omega) \mathcal{F}[g](t\omega)|^2 e^{\frac{\pi^2}{\eta_0} \|\omega\|^2} d\omega = c_\eta \int_{t\|\omega\|\leq 3} |\mathcal{F}[f](\omega)|^2 e^{\frac{4}{\eta_0} \|\omega\|^2} d\omega \quad (\text{D.17})$$

$$= c_\eta \int_{t\|\omega\|\leq 3} |\mathcal{F}[f](\omega)|^2 (1 + \|\omega\|^2)^\beta \frac{e^{\frac{4}{\eta_0} \|\omega\|^2}}{(1 + \|\omega\|^2)^\beta} d\omega \quad (\text{D.18})$$

$$\leq c_\eta \sup_{t\|\omega\|\leq 3} \frac{e^{\frac{4}{\eta_0} \|\omega\|^2}}{(1 + \|\omega\|^2)^\beta} \int |\mathcal{F}[f](\omega)|^2 (1 + \|\omega\|^2)^\beta d\omega \quad (\text{D.19})$$

$$\leq \|f\|_{W_2^\beta(\mathbb{R}^d)}^2 c_\eta 2^{2\beta} \sup_{r \leq 3/t} \frac{e^{r^2 \pi^2 / \eta_0}}{(1 + r^2)^\beta} \quad (\text{D.20})$$

Step 3. Bounding $\text{Tr}(\mathbf{M}_\varepsilon)$.

Note that the function $\frac{1}{(1+r^2)^\beta} \exp(\frac{r^2\pi^2}{\eta_0})$ has only one critical point in r that is a minimum, then $\sup_{r \leq 3/t} \frac{1}{(1+r^2)^\beta} \exp(\frac{r^2\pi^2}{\eta_0}) \leq \max[1, \frac{t^{2\beta}}{(t^2+9)^\beta} \exp(\frac{9\pi^2}{\eta_0 t^2})] \leq 1 + (t/3)^{2\beta} \exp(\frac{89}{\eta_0 t^2})$. Now let consider the functions $f_{i,t} = f_i \star g_t$ for $i \in \{1, \dots, q\}$ and note that, by the results above $\|f_{i,t} - f_i\|_{L^2(\mathbb{R}^d)} \leq \|f_i\|_{W_2^\beta(\mathbb{R}^d)} (2t)^\beta$ and $\|f_{i,t}\|_{\mathcal{H}_\eta}^2 \leq \|f_i\|_{W_2^\beta(\mathbb{R}^d)}^2 c_\eta 2^{2\beta} (1 + (t/3)^{2\beta} \exp(\frac{89}{\eta_0 t^2}))$. Since $f_{i,t}$ belong to the reproducing kernel Hilbert space \mathcal{H}_η , define the operator \mathbf{M}_ε as

$$\mathbf{M}_\varepsilon = \sum_{i=1}^q f_{i,t} f_{i,t}^\top.$$

First note that $\mathbf{M}_\varepsilon \in \mathbb{S}_+(\mathcal{H}_\eta)$, moreover $\text{rank}(\mathbf{M}_\varepsilon) = q$ and

$$\text{Tr}(\mathbf{M}_\varepsilon) = \sum_{i=1}^q \|f_{i,t}\|_{\mathcal{H}_\eta}^2 \leq c_\eta 2^{2\beta} (1 + (t/3)^{2\beta} e^{\frac{89}{\eta_0 t^2}}) \sum_{i=1}^q \|f_i\|_{W_2^\beta(\mathbb{R}^d)}^2. \quad (\text{D.21})$$

Step 4. Bounding $\|p - f(x; \mathbf{M}_\varepsilon, \phi_\eta)\|_{L^1(\mathbb{R}^d)}$.

Note that

$$f(x; \mathbf{M}_\varepsilon, \phi_\eta) = \phi_\eta(x)^\top \mathbf{M}_\varepsilon \phi_\eta(x) = \sum_{i=1}^q (f_{i,t}^\top \phi(x))^2 = \sum_{i=1}^q f_{i,t}(x)^2.$$

Then, since $a^2 - b^2 = (a - b)(a + b)$ for any $a, b \in \mathbb{R}$, by applying the Hölder inequality

$$\|p - f(x; \mathbf{M}_\varepsilon, \phi_\eta)\|_{L^1(\mathbb{R}^d)} = \left\| \sum_{i=1}^q f_i^2 - f_{i,t}^2 \right\|_{L^1(\mathbb{R}^d)} = \left\| \sum_{i=1}^q (f_i - f_{i,t})(f_i + f_{i,t}) \right\|_{L^1(\mathbb{R}^d)} \quad (\text{D.22})$$

$$\leq \sum_{i=1}^q \|f_i - f_{i,t}\|_{L^2(\mathbb{R}^d)} (\|f_i\|_{L^2(\mathbb{R}^d)} + \|f_{i,t}\|_{L^2(\mathbb{R}^d)}), \quad (\text{D.23})$$

finally, by the Young convolution inequality,

$$\|f_{i,t}\|_{L^2(\mathbb{R}^d)} = \|f_i \star g_t\|_{L^2(\mathbb{R}^d)} \leq \|f_i\|_{L^2(\mathbb{R}^d)} \|g_t\|_{L^1(\mathbb{R}^d)}.$$

By the change of variable $x = zt$, $dx = t^d dz$, we have

$$\|g_t\|_{L^1(\mathbb{R}^d)} = \int |g_t(tx)| dx = \int t^{-d} |g(x/t)| dx = \int |g(z)| dz = \|g\|_{L^1(\mathbb{R}^d)}. \quad (\text{D.24})$$

then we obtain

$$\|p - f(x; \mathbf{M}_\varepsilon, \phi_\eta)\|_{L^1(\mathbb{R}^d)} \leq (2t)^\beta (1 + \|g\|_{L^1(\mathbb{R}^d)}) \sum_{i=1}^q \|f_i\|_{W_2^\beta(\mathbb{R}^d)} \|f_i\|_{L^2(\mathbb{R}^d)}. \quad (\text{D.25})$$

Step 5. Bounding $\|p - f(x; \mathbf{M}_\varepsilon, \phi_\eta)\|_{L^2(\mathbb{R}^d)}$.

With the same reasoning above, we have

$$\|p - f(x; \mathbf{M}_\varepsilon, \phi_\eta)\|_{L^2(\mathbb{R}^d)} = \left\| \sum_{i=1}^q f_i^2 - f_{i,t}^2 \right\|_{L^2(\mathbb{R}^d)} = \left\| \sum_{i=1}^q (f_i - f_{i,t})(f_i + f_{i,t}) \right\|_{L^2(\mathbb{R}^d)} \quad (\text{D.26})$$

$$\leq \sum_{i=1}^q \|f_i - f_{i,t}\|_{L^2(\mathbb{R}^d)} (\|f_i\|_{L^\infty(\mathbb{R}^d)} + \|f_{i,t}\|_{L^\infty(\mathbb{R}^d)}) \quad (\text{D.27})$$

finally, by the Young convolution inequality,

$$\|f_{i,t}\|_{L^\infty(\mathbb{R}^d)} = \|f_i \star g_t\|_{L^\infty(\mathbb{R}^d)} \leq \|f_i\|_{L^\infty(\mathbb{R}^d)} \|g_t\|_{L^1(\mathbb{R}^d)}.$$

Then,

$$\|p - f(x; \mathbf{M}_\varepsilon, \phi_\eta)\|_{L^2(\mathbb{R}^d)} \leq (2t)^\beta (1 + \|g\|_{L^1(\mathbb{R}^d)}) \sum_{i=1}^q \|f_i\|_{W_2^\beta(\mathbb{R}^d)} \|f_i\|_{L^\infty(\mathbb{R}^d)} \quad (\text{D.28})$$

Step 6. Setting t appropriately.

Finally, noting that by construction $\|f\|_{L^2(\mathbb{R}^d)} \leq \|f\|_{W_2^\beta(\mathbb{R}^d)}$ and setting

$$t = \left(\frac{\varepsilon}{C_1}\right)^{\frac{1}{\beta}}, \quad C_1 = 2^\beta(1 + \|g\|_{L^1(\mathbb{R}^d)}) \sum_{i=1}^q \|f_i\|_{W_2^\beta(\mathbb{R}^d)} \max(\|f_i\|_{L^\infty(\mathbb{R}^d)}, \|f_i\|_{W_2^\beta(\mathbb{R}^d)}) \quad (\text{D.29})$$

then, $\|p - f(x; \mathbf{M}_\varepsilon, \phi_\eta)\|_{L^j(\mathbb{R}^d)} \leq \varepsilon$, with $j = 1, 2$. By Littlewood's interpolation inequality, $\|\cdot\|_{L^r(\mathbb{R}^d)} \leq \|\cdot\|_{L^1(\mathbb{R}^d)}^{(2/r)-1} \|\cdot\|_{L^2(\mathbb{R}^d)}^{2-(2/r)}$ when $r \in [1, 2]$ (see, e.g, Thm. 8.5 pag 316 of [6]), we have

$$\|p - f(x; \mathbf{M}_\varepsilon, \phi_\eta)\|_{L^r(\mathbb{R}^d)} \leq \varepsilon, \quad \forall r \in [1, 2]. \quad (\text{D.30})$$

By setting $C_2 = 2^{2\beta} \sum_{i=1}^q \|f_i\|_{W_2^\beta(\mathbb{R}^d)}^2$, we have

$$\text{Tr}(\mathbf{M}_\varepsilon) \leq c_\eta C_2 (1 + e^{\frac{89}{\eta_0 t^2}} (t/3)^{2\beta}) \leq c_\eta C_2 (1 + \frac{3^{-2\beta}}{C_1^2} \varepsilon^2 e^{\frac{89}{\eta_0} (\frac{C_1}{\varepsilon})^{2/\beta}}) \leq C |\eta|^{1/2} (1 + \varepsilon^2 e^{\frac{C'}{\eta_0 \varepsilon^{2/\beta}}}),$$

where $|\eta| = \det(\text{diag}(\eta))$ and $C' = 89C_1^{2/\beta}$ and $C = \pi^{-d/2} C_2 \max(1, 3^{-2\beta}/C_1^2)$. \square

D.3 Proof of Thm. 6

We can now prove [Thm. 6](#). We will prove a more general result [Thm. D.5](#), from which [Thm. 6](#) follows when $R = 1$ and $\mathcal{X} = (-1, 1)^d$ applied to $\tilde{f}_1, \dots, \tilde{f}_q \in W_2^\beta(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$ that are the extension to \mathbb{R}^d of the functions f_1, \dots, f_q characterizing p via [Assumption 1](#). The details of the extension are in [Cor. A.3](#)

Theorem D.5. *Let $R > 0$ and let $\mathcal{X} \subseteq T = (R, R)^d$ be a non-empty open set with Lipschitz boundary. Let $f_1, \dots, f_q \in W_2^\beta(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$ and let $p = \sum_{i=1}^q \tilde{f}_i^2$. Then, for any $\varepsilon \in (0, 1/e]$, there exists $m \in \mathbb{N}$, $\eta \in \mathbb{R}_{++}^d$, a base point matrix $\tilde{X} \in \mathbb{R}^{m \times d}$ and a matrix $A \in \mathbb{S}_+^m$ such that $\|f(\cdot; A, \tilde{X}, \eta) - p\|_{L^2(\mathcal{X})} \leq 2\varepsilon$, with*

$$m^{1/d} \leq C + C' \log \frac{1+R}{\varepsilon} + C'' R \varepsilon^{-\frac{1}{\beta}} (\log \frac{1+R}{\varepsilon})^{\frac{1}{2}} \quad (\text{D.31})$$

where C, C', C'' depend only on $\mathcal{X}, \beta, d, \|f_j\|_{W_2^\beta(\mathbb{R}^d)}, \|f_j\|_{L^\infty(\mathbb{R}^d)}$ for $j = 1, \dots, q$. This implies that there exists a model of dimension m such that $\|f(\cdot; A, \tilde{X}, \eta) - p\|_{L^2(\mathcal{X})} \leq \varepsilon$,

$$m = O\left(R^d \varepsilon^{-d/\beta} (\log \frac{1+R}{\varepsilon})^{d/2}\right). \quad (\text{D.32})$$

Proof. Let $\varepsilon \in (0, 1/e]$ and $\eta = \tau \mathbf{1}_d \in \mathbb{R}^d$ with $\tau > 0$ and $m \in \mathbb{N}$. Let $\mathbf{M}_\varepsilon \in \mathbb{S}_+(\mathcal{H}_\eta)$ be the operator constructed in [Thm. D.4](#). We consider the compression of the model $p_\varepsilon = f(\cdot; \mathbf{M}_\varepsilon, \phi_\eta)$ as in [Thm. C.4](#). In particular, let $\tilde{x}_1, \dots, \tilde{x}_m$ be a covering of T with ℓ_2 . We consider the following model $\tilde{p}_m = f(\cdot; A_m, \tilde{X}_m, \eta)$ where $\tilde{X}_m \in \mathbb{R}^{m \times d}$ is the base point matrix whose j -th row is the point \tilde{x}_j for $j = 1, \dots, m$, and where $A_m \in \mathbb{S}_+^m$ is defined as

$$A_m = K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z} \mathbf{M}_\varepsilon \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1}, \quad (\text{D.33})$$

where $\tilde{Z} : \mathcal{H}_\eta \rightarrow \mathbb{R}^m$ is defined in [\(B.6\)](#) and its adjoint \tilde{Z}^* has range equal to $\text{span}\{\phi_\eta(\tilde{x}_1), \dots, \phi_\eta(\tilde{x}_m)\}$. Note that, by construction $\tilde{Z} \mathbf{M}_\varepsilon \tilde{Z}^* \in \mathbb{S}_+^m$ and so $A_m \in \mathbb{S}_+^m$.

Step 1. Approximation error decomposition. We will split the approximation error as follows,

$$\|\tilde{p}_m - p\|_{L^2(\mathcal{X})} \leq \|\tilde{p}_m - p_\varepsilon\|_{L^2(\mathcal{X})} + \|p_\varepsilon - p\|_{L^2(\mathcal{X})}. \quad (\text{D.34})$$

Note that for the second term, by [Thm. D.4](#), we have

$$\|p_\varepsilon - p\|_{L^r(\mathcal{X})} \leq \|p_\varepsilon - p\|_{L^r(\mathbb{R}^d)} \leq \varepsilon, \quad \forall r \in [1, 2] \quad (\text{D.35})$$

Step 2. Error induced by a projection. Since an h -covering of a set has fill distance h , by definition of fill distance [\(C.1\)](#), then we will choose the m base points $\tilde{x}_1, \dots, \tilde{x}_m$ to be an h -covering of the hypercube T . Since the ℓ_2 -ball of diameter 1 contains a cube of side $1/\sqrt{d}$, it is possible to cover a

cube of side $2R$ with $m \leq (1 + 2R\sqrt{d}/h)^d$ balls of diameter $2h$ (and so of radius h), see, e.g., Thm. 5.3, page 76 of [8]. Now, by Thm. C.4 applied to M_ε , we have that when the fill distance h (defined in (C.1)) satisfies $h \leq \sigma/C'$ with $\sigma = \min(R, 1/\sqrt{\tau})$, then

$$|p_\varepsilon(x) - \tilde{p}_m(x)| \leq 2C\sqrt{\|M_\varepsilon\|p_\varepsilon(x)} e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} + C^2\|M_\varepsilon\| e^{-\frac{2c\sigma}{h} \log \frac{c\sigma}{h}}, \quad \forall x \in \mathcal{X} \quad (\text{D.36})$$

with c, C, C' depending only on d and $c_M =$. Now denoting by $\alpha = 2C\sqrt{\|M_\varepsilon\|} e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}}$ and $\beta = C^2\|M_\varepsilon\| e^{-\frac{2c\sigma}{h} \log \frac{c\sigma}{h}}$, we have

$$\|p_\varepsilon - \tilde{p}_m\|_{L^2(\mathcal{X})} \leq \|\alpha\sqrt{p_\varepsilon} + \beta\|_{L^2(\mathcal{X})} \leq \alpha\|\sqrt{p_\varepsilon}\|_{L^2(\mathcal{X})} + \beta\|1\|_{L^2(\mathcal{X})} \quad (\text{D.37})$$

$$\leq (2R)^{d/2}\beta + \alpha\|p_\varepsilon^{1/2}\|_{L^2(\mathcal{X})}^2. \quad (\text{D.38})$$

where we used the fact that $\|1\|_{L^2(\mathcal{X})}^2 \leq \|1\|_{L^2(T)}^2 = \int_T dx = (2R)^d$.

Step 3. Final bound. First, note that by Thm. D.4

$$\text{Tr}(M_\varepsilon) \leq C_1\tau^{d/2}(1 + \varepsilon^2 \exp(\frac{C_2}{\tau}\varepsilon^{-\frac{2}{\beta}})) \quad (\text{D.39})$$

where C_1, C_2 are independent on ε, τ and depend only on $\beta, d, \|f_i\|_{W_2^\beta(\mathbb{R}^d)}, \|f_i\|_{L^\infty(\mathbb{R}^d)}$. By setting $\tau = \frac{C_2\varepsilon^{-2/\beta}}{2\log \frac{1+R}{\varepsilon}}$ we have

$$\|M_\varepsilon\| \leq \text{Tr}(M_\varepsilon) \leq C_1\tau^{d/2}(1 + \varepsilon^2 \exp(\frac{C_2}{\tau}\varepsilon^{-\frac{2}{\beta}})) \leq (1+R)^2 C_3 \varepsilon^{-d/\beta}. \quad (\text{D.40})$$

with $C_3 = 2^{-d/2}C_1C_2^{d/2}$. Then, note that $\|p_\varepsilon^{1/2}\|_{L^2(\mathcal{X})} = \|p_\varepsilon\|_{L^1(\mathcal{X})}^{1/2}$, so, using (D.35), we have

$$\|p_\varepsilon\|_{L^1(\mathcal{X})} \leq \|p_\varepsilon - p\|_{L^1(\mathcal{X})} + \|p\|_{L^1(\mathcal{X})} \leq 1 + \varepsilon \leq 2. \quad (\text{D.41})$$

By choosing $h = c\sigma/s$ with $s = \max(C', (1 + \frac{d}{2\beta}) \log \frac{1}{\varepsilon} + (1 + \frac{d}{4}) \log(1+R) + \log(C\sqrt{C_3}) + e)$, since $s \geq e$, then $\log s \geq 1$, so

$$C e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} = C e^{-s \log s} \leq C e^{-s} \leq \frac{1}{15\sqrt{C_3}} (1+R)^{-d/4} \varepsilon^{1+\frac{d}{2\beta}}. \quad (\text{D.42})$$

Gathering the results from the previous steps, we have

$$\|\tilde{p}_m - p\|_{L^2(\mathcal{X})} \leq \varepsilon + 4\|M_\varepsilon\|^{1/2} C e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} + \|M_\varepsilon\| R^{d/2} C^2 e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} \quad (\text{D.43})$$

$$\leq \varepsilon + \frac{4}{15} (1+R)^{-(d/4)} \varepsilon + \frac{1}{225} \left(\frac{R}{R+1}\right)^{d/2} \varepsilon^2 \quad (\text{D.44})$$

$$\leq 2\varepsilon. \quad (\text{D.45})$$

To conclude we recall the fact that $\tilde{x}_1, \dots, \tilde{x}_m$ is a h -covering of T , guarantees that the number of centers m in the covering satisfies

$$m \leq (1 + \frac{2R\sqrt{d}}{h})^d. \quad (\text{D.46})$$

Then, since $h \geq c\sigma/(C_4 \log \frac{C_5 \log(1+R)}{\varepsilon})$ with $C_4 = 1 + d/\min(2\beta, 4)$ and $C_5 = (C\sqrt{C_3}e)^{1/C_4}$, and since $\sigma = \min(R, 1/\sqrt{\tau})$, then $R/\sigma = \max(1, R\sqrt{\tau}) \leq 1 + \sqrt{C_2/2} \varepsilon^{-1/\beta} (\log \frac{1+R}{\varepsilon})^{-1/2}$, so we have

$$m^{\frac{1}{d}} \leq 1 + 2R\sqrt{d}/h \leq 1 + C_4(1 + R\sqrt{d}(\frac{C_2}{2})^{1/2} \varepsilon^{-\frac{1}{\beta}} (\log \frac{1+R}{\varepsilon})^{-\frac{1}{2}}) \log \frac{C_5(1+R)}{\varepsilon} \quad (\text{D.47})$$

$$= (1 + C_4 \log C_5) + C_4 \log \frac{1+R}{\varepsilon} + RC_4\sqrt{d}(\frac{C_2}{2})^{\frac{1}{2}} \varepsilon^{-\frac{1}{\beta}} (\log \frac{1+R}{\varepsilon})^{-\frac{1}{2}} \log \frac{C_5(1+R)}{\varepsilon} \quad (\text{D.48})$$

$$\leq C_6 + C_4 \log \frac{1+R}{\varepsilon} + RC_7 \varepsilon^{-\frac{1}{\beta}} (\log \frac{1+R}{\varepsilon})^{\frac{1}{2}} \quad (\text{D.49})$$

with $C_6 = 1 + C_4 \log C_5$, $C_7 = C_4\sqrt{d}(\frac{C_2}{2})^{1/2} \log(eC_5)$, since $\log(eC_5) \geq 1$, then

$$\log \frac{C_5(1+R)}{\varepsilon} = \log eC_5 + \log \frac{(1+R)}{e\varepsilon} \leq (\log eC_5)(1 + \frac{\log(1+R)}{\log e\varepsilon}) \quad (\text{D.50})$$

$$\leq (\log eC_5)(1 + \log \frac{(1+R)}{e\varepsilon}) \leq (\log eC_5)(1 + \log \frac{(1+R)}{\varepsilon}). \quad (\text{D.51})$$

The constants C, C', C'' in the statement of the theorem correspond to $C = C_6, C' = C_4, C'' = C_7$. \square

E Learning a PSD model from examples

In this section we provide a proof for [Thm. 7](#), which characterizes the learning capabilities of PSD models. We first provide intermediate results that will be useful for the proof.

Let \mathcal{X} be a compact space and let $p : \mathcal{X} \rightarrow \mathbb{R}$ be a probability density which we assume to belong to $p \in L^2(\mathcal{X})$. Let x_1, \dots, x_n sampled i.i.d. according to p . We will study an estimator for p in terms of the squared L^2 norm $\|\cdot\|_{L^2(\mathcal{X})}$. Let $\eta = \eta_0 \mathbf{1}_d$ with $\eta_0 > 0$ and $\tilde{X} \in \mathbb{R}^{m \times d}$ the base point matrix whose rows are some points $\tilde{x}_1, \dots, \tilde{x}_m$. We will consider the following estimator \hat{p} for p

$$\hat{p}(x) = f(x; \hat{A}, \tilde{X}, \eta), \quad \hat{A} = \min_{A \in \mathbb{S}_+^m} \hat{L}_\lambda(A), \quad (\text{E.1})$$

and, denoting by \tilde{R} the Cholesky decomposition of $K_{\tilde{X}, \tilde{X}, \eta}$, i.e. the upper triangular matrix such that $K_{\tilde{X}, \tilde{X}, \eta} = \tilde{R}^\top \tilde{R}$, we define

$$\hat{L}_\lambda(A) = \int_{\mathcal{X}} f(x; A, \tilde{X}, \eta)^2 dx - \frac{2}{n} \sum_{i=1}^n \mathbf{1}_{\mathcal{X}}(x_i) f(x_i; A, \tilde{X}, \eta) + \lambda \|\tilde{R} A \tilde{R}^\top\|_F^2. \quad (\text{E.2})$$

Denote by $L_\lambda(A)$ the following functional

$$L_\lambda(A) = \int_{\mathcal{X}} f(x; A, \tilde{X}, \eta)^2 dx - \int_{\mathcal{X}} f(x; A, \tilde{X}, \eta) p(x) dx + \lambda \|\tilde{R} A \tilde{R}^\top\|_F^2. \quad (\text{E.3})$$

and by $\bar{A}_{\eta, \lambda} \in \mathbb{S}_+^m$ the matrix $\bar{A}_{\eta, \lambda} = \min_{A \in \mathbb{S}_+^m} L_\lambda(A)$. We have the following result

E.1 Operatorial characterization of $\hat{L}_\lambda, L_\lambda$

We can now rewrite the loss functions as follows

Lemma E.1 (Characterization of $\hat{L}_\lambda, L_\lambda$ in terms of \hat{v}, v). *For any $\lambda \geq 0$ the following holds*

$$\hat{L}_\lambda(A) = \|\text{Svec}(\tilde{Z}^* A \tilde{Z})\|_{L^2(\mathcal{X})}^2 + \lambda \|\text{vec}(\tilde{Z}^* A \tilde{Z})\|_{\mathcal{G}_\eta}^2 - 2 \left\langle \hat{v}, \text{vec}(\tilde{Z}^* A \tilde{Z}) \right\rangle_{\mathcal{G}_\eta} \quad (\text{E.4})$$

$$L_\lambda(A) = \|\text{Svec}(\tilde{Z}^* A \tilde{Z})\|_{L^2(\mathcal{X})}^2 + \lambda \|\text{vec}(\tilde{Z}^* A \tilde{Z})\|_{\mathcal{G}_\eta}^2 - 2 \left\langle v, \text{vec}(\tilde{Z}^* A \tilde{Z}) \right\rangle_{\mathcal{G}_\eta}. \quad (\text{E.5})$$

Proof. With the notation [Appendix A](#) and by using the operators defined in [Appendix B](#) for any $M \in \mathbb{S}_+(\mathcal{H}_\eta)$ we have

$$f(x; M, \phi_\eta) = \langle \psi_\eta(x), \text{vec}(M) \rangle_{\mathcal{G}_\eta}, \quad \forall x \in \mathbb{R}^d \quad (\text{E.6})$$

and in particular for any $A \in \mathbb{S}_+^m$, we have

$$f(x; A, \tilde{X}, \eta) = f(x; \tilde{Z}^* A \tilde{Z}, \phi_\eta) = \left\langle \psi_\eta(x), \text{vec}(\tilde{Z}^* A \tilde{Z}) \right\rangle_{\mathcal{G}_\eta}, \quad \forall x \in \mathbb{R}^d \quad (\text{E.7})$$

Now note that, by ciclicity of the trace, for any matrix $A, B \in \mathbb{R}^{m \times m}$ we have

$$\|B^{1/2} A B^{1/2}\|_F^2 = \text{Tr}(B^{1/2} A B^{1/2} B^{1/2} A B^{1/2}) = \text{Tr}(A B A B). \quad (\text{E.8})$$

This implies that $\|K_{\tilde{X}, \tilde{X}, \eta}^{1/2} A K_{\tilde{X}, \tilde{X}, \eta}^{1/2}\|_F^2 = \text{Tr}(A K_{\tilde{X}, \tilde{X}, \eta} A K_{\tilde{X}, \tilde{X}, \eta})$. Moreover, by ciclicity of the trace, definition of Frobenious norm and since $\tilde{Z} \tilde{Z}^* = K_{\tilde{X}, \tilde{X}, \eta}$ we have

$$\text{Tr}(A K_{\tilde{X}, \tilde{X}, \eta} A K_{\tilde{X}, \tilde{X}, \eta}) = \text{Tr}(A \tilde{Z} \tilde{Z}^* A \tilde{Z} \tilde{Z}^*) = \text{Tr}(\tilde{Z}^* A \tilde{Z} \tilde{Z}^* A \tilde{Z}) \quad (\text{E.9})$$

$$= \left\langle \text{vec}(\tilde{Z}^* A \tilde{Z}), \text{vec}(\tilde{Z}^* A \tilde{Z}) \right\rangle_{\mathcal{G}_\eta} = \|\text{vec}(\tilde{Z}^* A \tilde{Z})\|_{\mathcal{G}_\eta}^2. \quad (\text{E.10})$$

By linearity of the integral and the inner product and since ϕ_η and so ψ_η are uniformly bounded,

$$\int_{\mathcal{X}} f(x; M, \phi_\eta)^2 dx = \int_{\mathcal{X}} \langle \text{vec}(M), (\psi_\eta(x) \psi_\eta(x)^\top) \text{vec}(M) \rangle_{\mathcal{G}_\eta} dx \quad (\text{E.11})$$

$$= \left\langle \text{vec}(M), \left(\int_{\mathcal{X}} \psi_\eta(x) \psi_\eta(x)^\top dx \right) \text{vec}(M) \right\rangle_{\mathcal{G}_\eta} \quad (\text{E.12})$$

$$= \langle \text{vec}(M), Q \text{vec}(M) \rangle_{\mathcal{G}_\eta} = \langle \text{vec}(M), S^* \text{Svec}(M) \rangle_{\mathcal{G}_\eta} \quad (\text{E.13})$$

$$= \langle \text{Svec}(M), \text{Svec}(M) \rangle_{\mathcal{G}_\eta} = \|\text{Svec}(M)\|_{L^2(\mathcal{X})}^2 \quad (\text{E.14})$$

Then, we have

$$\hat{L}_\lambda(A) = \|S\text{vec}(\tilde{Z}^* A \tilde{Z})\|_{L^2(\mathcal{X})}^2 - 2 \langle \hat{v}, \text{vec}(\tilde{Z}^* A \tilde{Z}) \rangle_{\mathcal{G}_\eta} + \lambda \|\text{vec}(\tilde{Z}^* A \tilde{Z})\|_{\mathcal{G}_\eta}^2. \quad (\text{E.15})$$

The identical reasoning holds for $L_\lambda(A)$, with respect to v . \square

Theorem E.2 (Error decomposition). *Let \hat{A} be a minimizer of \hat{L}_λ over a set $S \subseteq \mathbb{R}^{m \times m}$ (non-necessarily convex). Denote by $\mu(A)$ the vector $\mu(A) = \text{vec}(\tilde{Z}^* A \tilde{Z}) \in \mathcal{G}_\eta$ for any $A \in S$. Then for any $A \in S$ the following holds*

$$\begin{aligned} (\|S\mu(\hat{A}) - p\|_{L^2(\mathcal{X})}^2 + \lambda \|\mu(\hat{A})\|_{\mathcal{G}_\eta}^2)^{1/2} &\leq (\|S\mu(A) - p\|_{L^2(\mathcal{X})}^2 + \lambda \|\mu(A)\|_{\mathcal{G}_\eta}^2)^{1/2} \\ &\quad + 3\sqrt{2} \|(Q + \lambda I)^{-1/2}(\hat{v} - v)\|_{\mathcal{G}_\eta}. \end{aligned} \quad (\text{E.16})$$

Proof. We start noting that since \hat{A} is the minimizer over S of \hat{L}_λ , then $\hat{L}_\lambda(\hat{A}) \leq \hat{L}_\lambda(A)$ for any $A \in S$. In particular, since $\bar{A} \in S$ this means that $\hat{L}_\lambda(\hat{A}) \leq \hat{L}_\lambda(\bar{A})$, then $\hat{L}_\lambda(\hat{A}) - \hat{L}_\lambda(\bar{A}) \leq 0$, this implies that

$$L_\lambda(\hat{A}) - L_\lambda(A) = L_\lambda(\hat{A}) - \hat{L}_\lambda(\hat{A}) + \hat{L}_\lambda(\hat{A}) - \hat{L}_\lambda(A) + \hat{L}_\lambda(A) - L_\lambda(A) \quad (\text{E.17})$$

$$\leq L_\lambda(\hat{A}) - \hat{L}_\lambda(\hat{A}) + \hat{L}_\lambda(A) - L_\lambda(A). \quad (\text{E.18})$$

Denote $\hat{\mu} = \text{vec}(\tilde{Z}^* \hat{A} \tilde{Z})$ and $\bar{\mu} = \text{vec}(\tilde{Z}^* \bar{A} \tilde{Z})$. Now note that, by the characterization of L_λ , \hat{L}_λ in [Lemma E.1](#), we have

$$L_\lambda(\hat{A}) - \hat{L}_\lambda(\hat{A}) + \hat{L}_\lambda(A) - L_\lambda(A) = 2 \langle \hat{v} - v, \hat{\mu} - \bar{\mu} \rangle_{\mathcal{G}_\eta}. \quad (\text{E.19})$$

Step 1. Decomposing the error.

Note that, since $v = S^*p$, then $\langle v, w \rangle_{\mathcal{G}_\eta} = \langle S^*p, w \rangle_{\mathcal{G}_\eta} = \langle p, Sw \rangle_{L^2(\mathcal{X})}$ for any $w \in \mathcal{G}_\eta$. Then, for any $A \in \mathbb{S}_+^m$, denoting by $\mu = \text{vec}(\tilde{Z}^* A \tilde{Z})$, and substituting $\langle v, \mu \rangle_{\mathcal{G}_\eta}$ with $\langle p, S\mu \rangle_{L^2(\mathcal{X})}$ in the definition of $L_\lambda(A)$, we have

$$L_\lambda(A) + \|p\|^2 = \|S\mu\|_{L^2(\mathcal{X})}^2 - 2 \langle p, S\mu \rangle_{L^2(\mathcal{X})} + \|p\|_{L^2(\mathcal{X})}^2 + \lambda \|\mu\|_{\mathcal{G}_\eta}^2 \quad (\text{E.20})$$

$$= \|S\mu - p\|_{L^2(\mathcal{X})}^2 + \lambda \|\mu\|_{\mathcal{G}_\eta}^2. \quad (\text{E.21})$$

From [\(E.17\)](#) and [\(E.19\)](#) and the equation above, we obtain

$$\begin{aligned} \|S\hat{\mu} - p\|_{L^2(\mathcal{X})}^2 + \lambda \|\hat{\mu}\|_{\mathcal{G}_\eta}^2 &= L_\lambda(\hat{A}) + \|p\|_{L^2(\mathcal{X})}^2 \\ &\leq L_\lambda(\bar{A}) + \|p\|_{L^2(\mathcal{X})}^2 + 2 \langle \hat{v} - v, \hat{\mu} - \bar{\mu} \rangle_{\mathcal{G}_\eta} \\ &= \|S\bar{\mu} - p\|_{L^2(\mathcal{X})}^2 + \lambda \|\bar{\mu}\|_{\mathcal{G}_\eta}^2 + 2 \langle \hat{v} - v, \hat{\mu} - \bar{\mu} \rangle_{\mathcal{G}_\eta} \end{aligned} \quad (\text{E.22})$$

Now, note that

$$\|S\hat{\mu} - p\|_{L^2(\mathcal{X})}^2 = \|S(\hat{\mu} - \bar{\mu}) + (S\bar{\mu} - p)\|_{L^2(\mathcal{X})}^2 \quad (\text{E.23})$$

$$= \|S(\hat{\mu} - \bar{\mu})\|_{L^2(\mathcal{X})}^2 + 2 \langle S(\hat{\mu} - \bar{\mu}), (S\bar{\mu} - p) \rangle_{L^2(\mathcal{X})} + \|S\bar{\mu} - p\|_{L^2(\mathcal{X})}^2 \quad (\text{E.24})$$

$$\|\hat{\mu}\|_{\mathcal{G}_\eta}^2 = \|(\hat{\mu} - \bar{\mu}) + \bar{\mu}\|_{\mathcal{G}_\eta}^2 = \|\hat{\mu} - \bar{\mu}\|_{\mathcal{G}_\eta}^2 + 2 \langle \bar{\mu}, \hat{\mu} - \bar{\mu} \rangle_{\mathcal{G}_\eta} + \|\bar{\mu}\|_{\mathcal{G}_\eta}^2. \quad (\text{E.25})$$

Expanding $\|S\hat{\mu} - p\|_{L^2(\mathcal{X})}^2$ and $\|\hat{\mu}\|_{\mathcal{G}_\eta}^2$ in [\(E.22\)](#) and reorganizing the terms, we obtain

$$\begin{aligned} \|S(\hat{\mu} - \bar{\mu})\|_{L^2(\mathcal{X})}^2 + \lambda \|\hat{\mu} - \bar{\mu}\|_{\mathcal{G}_\eta}^2 &\leq -2 \langle S(\hat{\mu} - \bar{\mu}), (S\bar{\mu} - p) \rangle_{L^2(\mathcal{X})} \\ &\quad - 2\lambda \langle \bar{\mu}, \hat{\mu} - \bar{\mu} \rangle_{\mathcal{G}_\eta} \\ &\quad + 2 \langle \hat{v} - v, \hat{\mu} - \bar{\mu} \rangle_{\mathcal{G}_\eta}. \end{aligned} \quad (\text{E.26})$$

Step 2. Bounding the three terms of the decomposition.

The proof is concluded by bounding the three terms of the right hand side of the equation above, indeed

$$-2 \langle S(\hat{\mu} - \bar{\mu}), (S\bar{\mu} - p) \rangle_{L^2(\mathcal{X})} \leq 2 \|S(\hat{\mu} - \bar{\mu})\|_{L^2(\mathcal{X})} \|S\bar{\mu} - p\|_{L^2(\mathcal{X})} \quad (\text{E.27})$$

$$-2\lambda \langle \bar{\mu}, \hat{\mu} - \bar{\mu} \rangle_{\mathcal{G}_\eta} \leq 2\lambda \|\bar{\mu}\|_{\mathcal{G}_\eta} \|\hat{\mu} - \bar{\mu}\|_{\mathcal{G}_\eta}. \quad (\text{E.28})$$

For the third term we multiply and divide by $Q + \lambda$ (it is invertible since $Q \in \mathbb{S}_+(\mathcal{G}_\eta)$ and $\lambda > 0$), so

$$\begin{aligned} 2 \langle \hat{v} - v, \hat{\mu} - \bar{\mu} \rangle_{\mathcal{G}_\eta} &= 2 \left\langle (Q + \lambda)^{-1/2}(\hat{v} - v), (Q + \lambda)^{1/2}(\hat{\mu} - \bar{\mu}) \right\rangle_{\mathcal{G}_\eta} \\ &\leq 2 \|(Q + \lambda)^{-1/2}(\hat{v} - v)\|_{\mathcal{G}_\eta} \|(Q + \lambda)^{1/2}(\hat{\mu} - \bar{\mu})\|_{\mathcal{G}_\eta}. \end{aligned} \quad (\text{E.29})$$

Note that for any $w \in \mathcal{G}_\eta$, since Q is characterized as $Q = S^*S$, we have

$$\begin{aligned} \|(Q + \lambda)^{1/2}w\|_{\mathcal{G}_\eta}^2 &= \langle w, (Q + \lambda)w \rangle_{\mathcal{G}_\eta} = \langle w, Qw \rangle_{\mathcal{G}_\eta} + \lambda \langle w, w \rangle_{\mathcal{G}_\eta} \\ &= \langle Sw, Sw \rangle_{\mathcal{G}_\eta} + \lambda \langle w, w \rangle_{\mathcal{G}_\eta} = \|Sw\|_{\mathcal{G}_\eta}^2 + \lambda \|w\|_{\mathcal{G}_\eta}^2. \end{aligned} \quad (\text{E.30})$$

By applying the equation above to $w = \hat{\mu} - \bar{\mu}$, since $\sqrt{a^2 + b^2} \leq a + b, \forall a, b \geq 0$, we have

$$\|(Q + \lambda)^{1/2}(\hat{\mu} - \bar{\mu})\|_{\mathcal{G}_\eta} \leq \|S(\hat{\mu} - \bar{\mu})\|_{L^2(\mathcal{X})} + \sqrt{\lambda} \|(\hat{\mu} - \bar{\mu})\|_{\mathcal{G}_\eta}. \quad (\text{E.31})$$

Combining (E.26) with the bounds in (E.27) and (E.28) for the first two terms of its right hand side and with the bounds in (E.29) and (E.31) for the third term, we have

$$\|S(\hat{\mu} - m)\|_{L^2(\mathcal{X})}^2 + \lambda \|\hat{\mu} - \bar{\mu}\|_{\mathcal{G}_\eta}^2 \leq 2\alpha \|S(\hat{\mu} - m)\|_{L^2(\mathcal{X})} + 2\beta \sqrt{\lambda} \|\hat{\mu} - \bar{\mu}\|_{\mathcal{G}_\eta}, \quad (\text{E.32})$$

with $\alpha = \|S\bar{\mu} - p\|_{L^2(\mathcal{X})} + \|(Q + \lambda)^{-1/2}(\hat{v} - v)\|_{\mathcal{G}_\eta}$, $\beta = \sqrt{\lambda} \|\bar{\mu}\|_{\mathcal{G}_\eta} + \|(Q + \lambda)^{-1/2}(\hat{v} - v)\|_{\mathcal{G}_\eta}$.

Step 3. Solving the inequality associated to the bound of the three terms.

Now denoting by $x = \|S(\hat{\mu} - \bar{\mu})\|_{L^2(\mathcal{X})}$ and $y = \sqrt{\lambda} \|\hat{\mu} - \bar{\mu}\|_{\mathcal{G}_\eta}$, the inequality above becomes

$$x^2 + y^2 \leq 2\alpha x + 2\beta y. \quad (\text{E.33})$$

By completing the squares it is equivalent to $(x - \alpha)^2 + (y - \beta)^2 \leq \alpha^2 + \beta^2$, from which we derive that $(x - \alpha)^2 \leq (x - \alpha)^2 + (y - \beta)^2 \leq \alpha^2 + \beta^2$. This implies $x \leq \alpha + \sqrt{\alpha^2 + \beta^2} \leq 2\alpha + \beta$. With the same reasoning we derive $y \leq \alpha + 2\beta$, that corresponds to

$$\|S(\hat{\mu} - \bar{\mu})\|_{L^2(\mathcal{X})} \leq 2\|S\bar{\mu} - p\|_{L^2(\mathcal{X})} + 3\|(Q + \lambda)^{-1/2}(\hat{v} - v)\|_{\mathcal{G}_\eta} + \sqrt{\lambda} \|\bar{\mu}\|_{\mathcal{G}_\eta} \quad (\text{E.34})$$

$$\|\hat{\mu} - \bar{\mu}\|_{\mathcal{G}_\eta} \leq \frac{1}{\sqrt{\lambda}} \|S\bar{\mu} - p\|_{L^2(\mathcal{X})} + \frac{3}{\sqrt{\lambda}} \|(Q + \lambda)^{-1/2}(\hat{v} - v)\|_{\mathcal{G}_\eta} + 2\|\bar{\mu}\|_{\mathcal{G}_\eta} \quad (\text{E.35})$$

Step 4. The final bound.

The final result is obtained by bounding the term $\langle \hat{v} - v, \hat{\mu} - \bar{\mu} \rangle_{\mathcal{G}_\eta}$ in (E.22). In particular, we will bound it by using (E.29), and by bounding the resulting term $\|(Q + \lambda)^{1/2}(\hat{\mu} - \bar{\mu})\|_{\mathcal{G}_\eta}$ with (E.31) and the resulting terms $\|S(\hat{\mu} - \bar{\mu})\|_{L^2(\mathcal{X})}$, $\|\hat{\mu} - \bar{\mu}\|_{\mathcal{G}_\eta}$ via (E.34). This leads to

$$\begin{aligned} \|S\hat{\mu} - p\|_{L^2(\mathcal{X})}^2 + \lambda \|\hat{\mu}\|_{\mathcal{G}_\eta}^2 &\leq \|S\bar{\mu} - p\|_{L^2(\mathcal{X})}^2 + \lambda \|\bar{\mu}\|_{\mathcal{G}_\eta}^2 \\ &\quad + 6\tau (\|S\bar{\mu} - p\|_{L^2(\mathcal{X})} + \sqrt{\lambda} \|\bar{\mu}\|_{\mathcal{G}_\eta} + 2\tau). \end{aligned} \quad (\text{E.36})$$

with $\tau = \|(Q + \lambda)^{-1/2}(\hat{v} - v)\|_{\mathcal{G}_\eta}$. We can optimize the writing of the theorem by noting that since $2ab \leq a^2 + b^2$ for any $a, b \in \mathbb{R}$, we have $a + b \leq \sqrt{2(a^2 + b^2)}$ and so $\|S\bar{\mu} - p\|_{L^2(\mathcal{X})} + \sqrt{\lambda} \|\bar{\mu}\|_{\mathcal{G}_\eta} \leq \sqrt{2(\|S\bar{\mu} - p\|_{L^2(\mathcal{X})}^2 + \lambda \|\bar{\mu}\|_{\mathcal{G}_\eta}^2)}$ and the bound in (E.36) becomes

$$\|S\hat{\mu} - p\|_{L^2(\mathcal{X})}^2 + \lambda \|\hat{\mu}\|_{\mathcal{G}_\eta}^2 \leq z^2 + 6\tau(\sqrt{2}z + 2\tau) \leq (z + 3\sqrt{2}\tau)^2 \quad (\text{E.37})$$

where $z^2 = \|S\bar{\mu} - p\|_{L^2(\mathcal{X})}^2 + \lambda \|\bar{\mu}\|_{\mathcal{G}_\eta}^2$. The final result is obtained by noting that, since \bar{A} is the minimizer of L_λ , then for any $A \in \mathbb{S}_+^m$ the following holds

$$z^2 = \|p\|^2 + L_\lambda(\bar{A}) = \|p\|^2 + \min_{A' \in \mathbb{S}_+^m} L_\lambda(A') \quad (\text{E.38})$$

$$\leq \|p\|^2 + L_\lambda(A) = \|S\mu(A) - p\|_{L^2(\mathcal{X})}^2 + \lambda \|\mu(A)\|_{\mathcal{G}_\eta}^2. \quad (\text{E.39})$$

□

Lemma E.3. Let $\mathcal{X} \subseteq \mathbb{R}^d$ and let $\eta \in \mathbb{R}_{++}^d$. Let $\tilde{x}_1, \dots, \tilde{x}_m$ be a set of points in \mathbb{R}^d and let $\tilde{X} \in \mathbb{R}^{m \times d}$ be the associated base point matrix, i.e., the j -th row of \tilde{X} corresponds to \tilde{x}_j . With the notation and the definitions of [Appendix B](#) denote by $\mu(A) = \text{vec}(\tilde{Z}^* A \tilde{Z})$. Then, for any $A \in \mathbb{S}_+^m$ and any $p \in L^2(\mathcal{X})$

$$\|S\mu(A) - p\|_{L^2(\mathcal{X})}^2 = \|f(\cdot; A, \tilde{X}, \eta) - p\|_{L^2(\mathcal{X})}^2, \quad (\text{E.40})$$

and moreover

$$\|\mu(A)\|_{\mathcal{G}_\eta}^2 = \|\tilde{Z}^* A \tilde{Z}\|_F^2. \quad (\text{E.41})$$

Proof. Denote by $\mu(A) = \text{vec}(\tilde{Z}^* A \tilde{Z})$. We recall that the operator used in the rest of the section are defined in [Appendix B](#). First we recall from [\(E.20\)](#) that, by definition of $S : \mathcal{G}_\eta \rightarrow L^2(\mathcal{X})$, we have $S\mu(A) = \langle \psi_\eta(\cdot), \mu(A) \rangle_{\mathcal{G}_\eta} \in L^2(\mathcal{X})$. In particular, by definition of vec , for any $x \in \mathcal{X}$

$$(S\mu(A))(x) = \langle \mu(A), \psi_\eta(x) \rangle_{\mathcal{G}_\eta} = \left\langle \text{vec}(\tilde{Z}^* A \tilde{Z}), \phi_\eta(x) \otimes \phi_\eta(x) \right\rangle_{\mathcal{G}_\eta} = \phi_\eta(x)^\top \tilde{Z}^* A \tilde{Z} \phi_\eta(x). \quad (\text{E.42})$$

Now, since $\tilde{Z}^* A \tilde{Z} = \sum_{i,j=1}^m A_{i,j} \phi_\eta(x_i) \phi_\eta(x_j)^\top$, and, by the representer property of the kernel k_η we have $k_\eta(x, x') = \phi_\eta(x)^\top \phi_\eta(x')$ (see [\(A.2\)](#) and [Example 1](#)), then

$$\phi_\eta(x)^\top \tilde{Z}^* A \tilde{Z} \phi_\eta(x) = \sum_{i,j=1}^m A_{i,j} (\phi_\eta(x)^\top \phi_\eta(x_i)) (\phi_\eta(x_j)^\top \phi_\eta(x)) \quad (\text{E.43})$$

$$= \sum_{i,j=1}^m A_{i,j} k_\eta(x, x_i) k_\eta(x, x_j) = f(x; A, \tilde{X}, \eta). \quad (\text{E.44})$$

Then $S\mu(A) = f(\cdot; A, \tilde{X}, \eta) \in L^2(\mathcal{X})$. So

$$\|S\mu(A) - p\|_{L^2(\mathcal{X})}^2 = \|f(\cdot; A, \tilde{X}, \eta) - p\|_{L^2(\mathcal{X})}^2. \quad (\text{E.45})$$

To conclude note that, by the properties of vec (see [Appendix B](#)), we have

$$\|\mu(A)\|_{\mathcal{G}_\eta}^2 = \left\langle \text{vec}(\tilde{Z}^* A \tilde{Z}), \text{vec}(\tilde{Z}^* A \tilde{Z}) \right\rangle_{\mathcal{G}_\eta} = \|\tilde{Z}^* A \tilde{Z}\|_F^2 \quad (\text{E.46})$$

□

Lemma E.4. Let $s \in \mathbb{N}$ and $\delta \in (0, 1]$, $\lambda > 0$ let $\mathcal{X} \subset \mathbb{R}^d$ be an open set with Lipschitz boundary and let $p \in L^1(\mathcal{X}) \cap L^\infty(\mathcal{X})$. Then the following holds with probability $1 - \delta$

$$\|(Q + \lambda I)^{-1/2}(\hat{v} - v)\|_{\mathcal{G}_\eta} \leq \frac{C\tau^{d/4} \log \frac{2}{\delta}}{n(\lambda\tau^d)^{d/4s}} + \sqrt{\frac{2\text{Tr}(Q_\lambda^{-1}Q) \log \frac{2}{\delta}}{n}}, \quad (\text{E.47})$$

where C is a constant that depends only on \mathcal{X} , s , d .

Proof. We are going to use here a Bernstein inequality for random vectors in separable Hilbert spaces (see, e.g., Thm 3.3.4 of [\[41\]](#)). Define the random variable $\zeta = Q_\lambda^{-1/2} \psi_\eta(x) \mathbf{1}_\mathcal{X}(x)$ with x distributed according to ρ and $Q_\lambda = Q + \lambda I$. To apply such inequality, we need to control the second moment and the norm of ζ . First note that ζ

Step 1. Bounding the variance of ζ . We have

$$\mathbb{E}\|\zeta\|^2 = \text{Tr} \left(\int \zeta \zeta^\top p(x) dx \right) = \text{Tr} \left(\int_{\mathcal{X}} Q_\lambda^{-1/2} (\psi_\eta(x) \psi_\eta(x)^\top) Q_\lambda^{-1/2} p(x) dx \right) \quad (\text{E.48})$$

$$\leq \|p\|_{L^\infty(\mathcal{X})} \text{Tr} \left(Q_\lambda^{-1/2} \left(\int_{\mathcal{X}} \psi_\eta(x) \psi_\eta(x)^\top dx \right) Q_\lambda^{-1/2} \right) \quad (\text{E.49})$$

$$= \|p\|_{L^\infty(\mathcal{X})} \text{Tr}(Q_\lambda^{-1/2} Q Q_\lambda^{-1/2}) = \text{Tr}(Q Q_\lambda^{-1}). \quad (\text{E.50})$$

Bounding the norm of ζ . Applying [Lemma C.2](#) the the operator $Q_\lambda^{-1/2}$ we have that

$$\text{ess-sup}\|\zeta\|_{\mathcal{G}_\eta} \leq \sup_{x \in \mathcal{X}} \|Q_\lambda^{-1/2} \phi_\eta(x)\| \leq \sup_{\|f\|_{\mathcal{G}_\eta} \leq 1} \|Q_\lambda^{-1/2} f\|_{L^\infty(\mathcal{X})}. \quad (\text{E.51})$$

Now, according to the definitions in [Appendix B](#), note that the reproducing kernel Hilbert space \mathcal{G}_η is associated to the kernel $h(x, x') = \psi_\eta(x)^\top \psi_\eta(x')$, that corresponds to

$$h(x, x') = \psi_\eta(x)^\top \psi_\eta(x') = \langle \phi_\eta(x) \otimes \phi_\eta(x), \phi_\eta(x') \otimes \phi_\eta(x') \rangle_{\mathcal{H}_\eta \otimes \mathcal{H}_\eta} \quad (\text{E.52})$$

$$= (\phi_\eta(x)^\top \phi_\eta(x'))^2 = k_{2\eta}(x, x')^2 = k_{2\eta}(x, x'). \quad (\text{E.53})$$

Then \mathcal{G}_η is still a RKHS associated to a Gaussian kernel, in particular $k_{2\eta}$, so $\mathcal{G}_\eta \subset W_2^s(\mathbb{R}^d)$ for any $s \geq 0$. In particular, by [\(C.9\)](#) and the fact that $\mathcal{X} \subset \mathbb{R}^d$, we have

$$\|f\|_{W_2^s(\mathcal{X})} \leq \|g\|_{W_2^s(\mathbb{R}^d)} \leq C \|f\|_{\mathcal{G}_\eta} \tau^{(s-d)/2}, \quad \forall f \in \mathcal{G}_\eta, \quad (\text{E.54})$$

where C depends only on s, d . Now by the interpolation inequality for Sobolev spaces (see, e.g. Thm 5.9, page 139 of [1]), we have that for any $g \in W_2^s(\mathcal{X})$ the following holds

$$\|g\|_{L^\infty(\mathcal{X})} \leq C' \|g\|_{W_2^s(\mathcal{X})}^{d/(2s)} \|g\|_{L^2(\mathcal{X})}^{1-d/(2s)}. \quad (\text{E.55})$$

Applying the inequality above to the function $g = Q_\lambda^{-1/2} f$, with $f \in \mathcal{G}_\eta$ and $\|f\|_{\mathcal{G}_\eta} \leq 1$, we have

$$\|Q_\lambda^{-1/2} f\|_{L^\infty(\mathcal{X})} \leq C' \|Q_\lambda^{-1/2} f\|_{W_2^s(\mathcal{X})}^{d/(2s)} \|Q_\lambda^{-1/2} f\|_{L^2(\mathcal{X})}^{1-d/(2s)} \quad (\text{E.56})$$

$$\leq C' C^{d/2s} \tau^{\frac{(s-d)d}{4s}} \|Q_\lambda^{-1/2} f\|_{\mathcal{G}_\eta}^{d/2s} \|Q_\lambda^{-1/2} f\|_{L^2(\mathcal{X})}^{1-d/(2s)}, \quad (\text{E.57})$$

$$\leq C' C^{d/2s} \tau^{\frac{(s-d)d}{4s}} \lambda^{-d/4s} \|Q_\lambda^{-1/2} f\|_{L^2(\mathcal{X})}^{1-d/(2s)}. \quad (\text{E.58})$$

Finally note that, since by the reproducing property $g(x) = \langle \psi_\eta(x), g \rangle_{\mathcal{G}_\eta}$ for any $g \in \mathcal{G}_\eta$, we have that for any $f \in \mathcal{G}_\eta$ such that $\|f\|_{\mathcal{G}_\eta} \leq 1$, we have

$$\|Q_\lambda^{1/2} f\|_{L^2(\mathcal{X})}^2 = \int_{\mathcal{X}} (Q_\lambda^{1/2} f)(x)^2 dx \quad (\text{E.59})$$

$$= \int_{\mathcal{X}} \langle Q_\lambda^{1/2} f, (\psi_\eta(x) \psi_\eta(x)^\top Q_\lambda^{1/2} f) \rangle_{\mathcal{G}_\eta} \quad (\text{E.60})$$

$$= \int \left\langle f, Q_\lambda^{1/2} \left(\int_{\mathcal{X}} \psi_\eta(x) \psi_\eta(x)^\top dx \right) Q_\lambda^{1/2} f \right\rangle_{\mathcal{G}_\eta} \quad (\text{E.61})$$

$$\leq \|f\|^2 \|Q_\lambda^{-1/2} Q Q_\lambda^{-1/2}\| \leq 1. \quad (\text{E.62})$$

Then, to recap

$$\text{ess-sup}\|\zeta\|_{\mathcal{G}_\eta} \leq \sup_{\|f\|_{\mathcal{G}_\eta} \leq 1} \|Q_\lambda^{-1/2} f\|_{L^\infty(\mathcal{X})} \leq \sup_{\|f\|_{\mathcal{G}_\eta} \leq 1} \|Q_\lambda^{-1/2} f\|_{W_2^s(\mathcal{X})}^{d/(2s)} \|Q_\lambda^{-1/2} f\|_{L^2(\mathcal{X})}^{1-d/(2s)} \quad (\text{E.63})$$

$$\leq C' C^{d/2s} \tau^{\frac{(s-d)d}{4s}} \lambda^{-d/4s}. \quad (\text{E.64})$$

Step 3. Bernstein inequality for random vectors in separable Hilbert spaces. The points x_1, \dots, x_n are independently and identically distributed according to p . Define the random variables $\zeta_i = Q_\lambda^{-1/2} \phi_\eta(x_i) \mathbf{1}_{\mathcal{X}}(x_i)$ for $i = 1, \dots, n$. Note that ζ_i are independent and identically distributed with the same distribution as ζ . Now note that

$$\mathbb{E} \zeta_i = \mathbb{E} \zeta = Q_\lambda^{-1/2} \int_{\mathcal{X}} p(x) \psi_\eta(x) dx = Q_\lambda^{-1/2} v, \quad (\text{E.65})$$

moreover $\frac{1}{n} \sum_{i=1}^n \zeta_i = Q_\lambda^{-1/2} \hat{v}$. Then, given the bounds on the variance and on the norm for ζ , by applying a Bernstein inequality for random vectors in separable Hilbert spaces, as, e.g., Thm 3.3.4 of [41] (we will use the notation of Prop. 2 of [25]), the following holds with probability $1 - \delta$

$$\|Q_\lambda^{-1/2}(\hat{v} - v)\|_{\mathcal{G}_\eta} = \left\| \frac{1}{n} \sum_{i=1}^n \zeta_i - \mu \right\|_{\mathcal{G}_\eta} \leq \frac{M \log \frac{2}{\delta}}{n} + \sqrt{\frac{S \log \frac{2}{\delta}}{n}}, \quad (\text{E.66})$$

with $M = C' C^{d/2s} \tau^{\frac{(s-d)d}{4s}} \lambda^{-d/4s}$ and $S = 2\text{Tr}(Q_\lambda^{-1} Q)$. \square

Lemma E.5. Let $\eta = \tau \mathbf{1}_d$ with $\tau \geq 1$ and $\lambda \leq 1/2$. Let $\mathcal{X} \subset [-1, 1]^d$. Then

$$\mathrm{Tr}((Q + \lambda I)^{-1}Q) \leq C\tau^{d/2} \left(\log \frac{1}{\lambda}\right)^d, \quad (\text{E.67})$$

Proof. Let $\sigma_j(Q)$ with $j \in \mathbb{N}$ be the sequence of singular values of Q in non-increasing order. First note that for any $k \in \mathbb{N}$,

$$\sigma_k(Q) \leq \sum_{j \in \mathbb{N}} \sigma_j(Q) \leq \mathrm{Tr}(Q) = \mathrm{Tr} \left(\int_{\mathcal{X}} \psi_\eta(x) \psi_\eta(x)^\top dx \right) = \int_{\mathcal{X}} \|\psi_\eta(x)\|_{\mathcal{G}_\eta}^2 dx \leq 2^d, \quad (\text{E.68})$$

since $\|\psi_\eta(x)\|_{\mathcal{G}_\eta}^2 = 1$ and $\mathcal{X} \subseteq [-1, 1]^d$. Moreover, for $k \in \mathbb{N}$, let $\bar{x}_{1,k}, \dots, \bar{x}_{k,k}$ be a minimal covering of $[-1, 1]^d$. Let \tilde{P}_k be the projection operator whose range is $\mathrm{span}\{\bar{x}_{1,k}, \dots, \bar{x}_{k,k}\}$. Note that \tilde{P}_k has rank k . By the Eckart-Young theorem, we have that $\sigma_{k+1}(Q) = \inf_{\mathrm{rank}(A)=k} \|Q - A\|_F^2$, then

$$\sigma_{k+1}(Q) = \inf_{\mathrm{rank}(A)=k} \|Q - A\|_F^2 \leq \|Q - \tilde{P}_k Q\|_F^2 \quad (\text{E.69})$$

$$\leq \left\| \int_{\mathcal{X}} (I - \tilde{P}_k) \psi_\eta(x) \psi_\eta(x) dx \right\|_F^2 \leq \int_{\mathcal{X}} \|(I - \tilde{P}_k) \psi_\eta(x)\|_F^2 dx \quad (\text{E.70})$$

$$= \int_{\mathcal{X}} \|(I - \tilde{P}_k) \psi_\eta(x)\|_{\mathcal{G}_\eta} \|\psi_\eta(x)\|_{\mathcal{G}_\eta} dx \leq 2^d \sup_{x \in \mathcal{X}} \|(I - \tilde{P}_k) \psi_\eta(x)\|_{\mathcal{G}_\eta}. \quad (\text{E.71})$$

Since $\bar{x}_{1,k}, \dots, \bar{x}_{k,k}$ is a minimal covering of $[-1, 1]^d$ and since the ℓ_2 -ball of diameter 1 contains a cube of side $1/\sqrt{d}$, it is possible to cover a cube of side 2 with k balls of diameter $2h$ (and so of radius h) with $h \leq 2\sqrt{d}/(k^{1/d} - 1)$, see, e.g., Thm. 5.3, page 76 of [8]. Since \mathcal{G}_η is a reproducing kernel Hilbert space with Gaussian kernel $k_{2\eta}$ as discussed in (E.52), by applying Thm. C.3, we have that when $h \leq 1/(C'\sqrt{\tau})$ and

$$\sup_{x \in \mathcal{X}} \|(I - \tilde{P}_k) \psi_\eta(x)\|_{\mathcal{G}_\eta} \leq C e^{-\frac{c}{\sqrt{\tau}h} \log \frac{c}{\sqrt{\tau}h}} \leq C e^{-\frac{c(k^{1/d}-1)}{2\sqrt{d}\sqrt{\tau}} \log \frac{c(k^{1/d}-1)}{2\sqrt{d}\sqrt{\tau}}}, \quad (\text{E.72})$$

with c, C, C' depending only on d . Take $k_\tau \geq (1 + 2 \max(1, e/c, C')\sqrt{d\tau})^d$. When $k \geq k_\tau$, we have $h \leq 1/(C'\sqrt{\tau})$, $c/(h\sqrt{\tau}) \leq e$, and $k \geq 2^d$, so $(k^{1/d} - 1) \geq k^{1/d}/2$, then

$$\sup_{x \in \mathcal{X}} \|(I - \tilde{P}_k) \psi_\eta(x)\|_{\mathcal{G}_\eta} \leq C e^{-\frac{c_1 k^{1/d}}{\sqrt{\tau}}}, \quad \forall k \geq k_\tau. \quad (\text{E.73})$$

with $c_1 = c/(4\sqrt{d})$. Let $g(x) = C e^{-\frac{c_1}{\sqrt{\tau}} x^{1/d}}$ for $x \geq 0$. Since g is non-increasing then $g(n+1) \geq \int_n^{n+1} g(x) dx$. Since $\int_k^\infty e^{-\frac{c_1}{\sqrt{\tau}} x^{1/d}} \leq d(\frac{c_1}{\sqrt{\tau}})^{-d} \Gamma(d, k^{1/d} \frac{c_1}{\sqrt{\tau}})$ where $\Gamma(a, x)$ is the *incomplete Gamma function* and is bounded by $\Gamma(a, x) \leq 2x^a e^{-x}$ for any $x, a \geq 1$ and $x \geq 2a$ (see Lemma P, page 31 of [2]). The condition to apply the bound on the incomplete Gamma function in our case corresponds to require k to satisfy $k \geq (2d/c_1)^d \tau^{d/2}$. Then, for any $k \in \mathbb{N}$, we have

$$\sum_{t > k+2} \sigma_t(Q) \leq 2^d \begin{cases} t & t \leq k'_\tau \\ d k e^{-\frac{c_1}{\sqrt{\tau}} k^{1/d}} & t > k'_\tau, \end{cases} \quad (\text{E.74})$$

where $k'_\tau = \max(k_\tau, 2 + (2d/c_1)^d \tau^{d/2})$. The bound on $\mathrm{Tr}((Q + \lambda I)^{-1}Q)$ is obtained by considering the characterization of the trace of an operator in terms of its singular values and the fact that $\frac{z}{z+\lambda} \leq \min(1, \frac{z}{\lambda})$. For any $k \geq k'_\tau$, we have

$$\mathrm{Tr}((Q + \lambda)^{-1}Q) = \sum_{t \in \mathbb{N}} \frac{\sigma_t(Q)}{\sigma_t(Q) + \lambda} = \sum_{t \leq (k+1)} \frac{\sigma_t(Q)}{\sigma_t(Q) + \lambda} + \sum_{t \geq k+2} \frac{\sigma_t(Q)}{\sigma_t(Q) + \lambda} \quad (\text{E.75})$$

$$\leq 1 + k + \frac{1}{\lambda} \sum_{t \geq k+2} \sigma_t(Q) \leq 1 + k(1 + \frac{1}{\lambda}) e^{-\frac{c_1}{\sqrt{\tau}} k^{1/d}}. \quad (\text{E.76})$$

In particular, choosing $\bar{k}_\tau = (\sqrt{\tau} C_2 \log(1 + \frac{1}{\lambda}))^d$ with $C_2 = 2d/c_1 + \max(1, e/c, C')\sqrt{d}$, then $\bar{k} \geq k'_\tau$ and so

$$\mathrm{Tr}((Q + \lambda)^{-1}Q) \leq 1 + (\sqrt{\tau} C_2 \log(1 + \frac{1}{\lambda}))^d \leq (\sqrt{\tau} 4 C_2 \log(\frac{1}{\lambda}))^d. \quad (\text{E.77})$$

□

E.2 Proof of Thm. 7

We are finally ready to prove [Thm. 7](#). We restate here the theorem for convenience.

Theorem 7. *Let $n, m \in \mathbb{N}$, $\lambda > 0$, $\eta \in \mathbb{R}_{++}^d$ and p be a density satisfying [Assumption 1](#). With the definitions above, let $\hat{p}_{n,m}$ be the model $\hat{p}_{n,m}(x) = f(x; \hat{A}, \tilde{X}, \eta)$, with \hat{A} the minimizer of [\(12\)](#). Let $\eta = n^{\frac{2}{2\beta+d}} \mathbf{1}_d$ and $\lambda = n^{-\frac{2\beta+2d}{2\beta+d}}$. When $m \geq C'n^{\frac{d}{2\beta+d}} (\log n)^d \log(C''n(\log n))$, the following holds with probability at least $1 - \delta$,*

$$\|p - \hat{p}_{n,m}\|_{L^2(\mathcal{X})} \leq Cn^{-\frac{\beta}{2\beta+d}} (\log n)^{d/2}, \quad (13)$$

where constant C depends only on β, d and p and the constants C', C'' depend only on β, d .

Proof. Let $\varepsilon \in (0, 1]$, $h > 0$ and $\eta = \tau \mathbf{1}_d$ with $\tau \in [1, \infty)$, to be fixed later. Denote by \hat{p} the model associated to matrix $\hat{A} \in \mathbb{S}_+^m$ that minimizes \hat{L}_λ over the set \mathbb{S}_+^m , i.e., $\hat{p} = f(\cdot; \hat{A}, \tilde{X}, \eta)$. The goal is then to bound $\|\hat{p} - p\|_{L^2(\mathbb{R}^d)}$. First we introduce the probabilities $p_\varepsilon, \tilde{p}_\varepsilon$, useful to perform the analysis. Let $M_\varepsilon \in \mathbb{S}_+(\mathcal{H}_\eta)$ be the operator such that $p_\varepsilon = f(\cdot; M_\varepsilon, \phi_\eta)$ approximates p with error ε as defined in [Thm. D.4](#), on the functions $\tilde{f}_1, \dots, \tilde{f}_q \in W_2^\beta(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$ that are the extension to \mathbb{R}^d of the functions f_1, \dots, f_q characterizing p via [Assumption 1](#). The details of the extension are in [Cor. A.3](#). Now, consider the model the compressed model $\tilde{p}_\varepsilon = f(\cdot; \tilde{A}_\varepsilon, \tilde{X}_m, \eta)$ with

$$\tilde{A}_\varepsilon = K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z} M_\varepsilon \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1}, \quad (E.78)$$

where $\tilde{Z} : \mathcal{H}_\eta \rightarrow \mathbb{R}^m$ is defined in [\(B.6\)](#) in terms of \tilde{X}_m .

Step 1. Decomposition of the error.

By applying [Thm. E.2](#) with $A = \tilde{A}_\varepsilon$ and [Lemma E.3](#) to simplify the notation, we derive

$$\|\hat{p} - p\|_{L^2(\mathcal{X})} \leq \|\tilde{p}_\varepsilon - p\|_{L^2(\mathcal{X})} + \sqrt{\lambda} \|\tilde{M}_\varepsilon\|_F + 5 \|Q_\lambda^{-\frac{1}{2}}(\hat{v} - v)\|_{\mathcal{G}_\eta}, \quad (E.79)$$

where $Q_\lambda = Q + \lambda I$ and $\tilde{M}_\varepsilon = \tilde{Z}^* \tilde{A}_\varepsilon \tilde{Z}_\varepsilon$. Note that

$$\tilde{M}_\varepsilon = \tilde{Z}^* \tilde{A}_\varepsilon \tilde{Z}_\varepsilon = \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z} M_\varepsilon \tilde{Z}^* K_{\tilde{X}, \tilde{X}, \eta}^{-1} \tilde{Z}_\varepsilon = \tilde{P} M_\varepsilon \tilde{P}, \quad (E.80)$$

where $\tilde{P} : \mathcal{H}_\eta \rightarrow \mathcal{H}_\eta$ is defined in [Appendix B](#) and is the projection operator on the range of \tilde{Z}^* , so

$$\|\tilde{M}_\varepsilon\|_F = \|\tilde{P} M_\varepsilon \tilde{P}\|_F \leq \|\tilde{P}\|^2 \|M_\varepsilon\|_F \leq \|M_\varepsilon\|_F. \quad (E.81)$$

Bounding $\|\tilde{p}_\varepsilon - p\|_{L^2(\mathcal{X})}$ with $\|\tilde{p}_\varepsilon - p\|_{L^2(\mathcal{X})} \leq \|\tilde{p}_\varepsilon - p_\varepsilon\|_{L^2(\mathcal{X})} + \|p_\varepsilon - p\|_{L^2(\mathcal{X})}$, we obtain

$$\|\hat{p} - p\|_{L^2(\mathcal{X})} \leq \|\tilde{p}_\varepsilon - p_\varepsilon\|_{L^2(\mathcal{X})} + \|p_\varepsilon - p\|_{L^2(\mathcal{X})} + \sqrt{\lambda} \|M_\varepsilon\|_F + 5 \|Q_\lambda^{-\frac{1}{2}}(\hat{v} - v)\|_{\mathcal{G}_\eta}, \quad (E.82)$$

Step 2. Bounding the terms of the decomposition.

Let h be the fill distance (defined in [\(C.1\)](#)) associated to the points $\tilde{x}_1, \dots, \tilde{x}_m$. By [Thm. C.4](#), there exist three constants c, C, C' depending only on d such that, when $h \leq \sigma/C'$, with $\sigma = \min(1, 1/\sqrt{\tau})$, then for any $x \in \mathcal{X}$

$$|\tilde{p}_\varepsilon(x) - p_\varepsilon(x)| \leq 2C \sqrt{\|M_\varepsilon\|_{p_\varepsilon(x)}} e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} + C^2 \|M_\varepsilon\| e^{-\frac{2c\sigma}{h} \log \frac{c\sigma}{h}}. \quad (E.83)$$

Since $q_\eta = 1$, $p_\varepsilon(x) = \phi_\eta(x)^\top M_\varepsilon \phi_\eta(x) \leq \|M_\varepsilon\| \|\phi_\eta(x)\|_{\mathcal{H}_\eta}^2 = \|M_\varepsilon\|$ for any $x \in \mathbb{R}^d$, since $\|\phi_\eta(x)\|_{\mathcal{H}_\eta}^2 = \phi_\eta(x)^\top \phi_\eta(x) = k_\eta(x, x) = 1$, then

$$\|\tilde{p}_\varepsilon - p_\varepsilon\|_{L^2(\mathcal{X})} \leq \text{vol}(\mathcal{X}) (2C + C^2) \|M_\varepsilon\| e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} \quad (E.84)$$

where $\text{vol}(\mathcal{X})$ is the volume of \mathcal{X} and is $\text{vol}(\mathcal{X}) = 2^d$. By [Thm. D.4](#) we know also that $\|p_\varepsilon - p\|_{L^2(\mathcal{X})} \leq \|p_\varepsilon - p\|_{L^2(\mathbb{R}^d)} \leq \varepsilon$. Moreover, we also know that there exists two constants C_1, C_2 depending on \mathcal{X}, β, d and the norms of $\tilde{f}_1, \dots, \tilde{f}_q$ (and so on the norms of f_1, \dots, f_q via [Cor. A.3](#)) such that

$$\|M_\varepsilon\| \leq \|M_\varepsilon\|_F \leq \text{Tr}(M_\varepsilon) \leq C_1 \tau^{d/2} (1 + \varepsilon^2 \exp(\frac{C_2}{\tau} \varepsilon^{-\frac{2}{\beta}})) \leq 2C_1 \tau^{d/2} \exp(\frac{C_2}{\tau} \varepsilon^{-\frac{2}{\beta}}) \quad (E.85)$$

By bounding $\|Q_\lambda^{-\frac{1}{2}}(\hat{v} - v)\|_{\mathcal{G}_\eta}$ via [Lemma E.4](#), with $s = d$ and [Lemma E.5](#), then [\(E.82\)](#) is bounded by

$$\|\hat{p} - p\|_{L^2(\mathcal{X})} \leq \varepsilon + C_3 \tau^{d/2} e^{\frac{C_2}{\tau} \varepsilon^{-\frac{2}{\beta}}} \left(\sqrt{\lambda} + e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} \right) \quad (\text{E.86})$$

$$+ \frac{C_4 \log \frac{2}{\delta}}{n \lambda^{1/4}} + \left(\frac{C_5 \tau^{d/2} (\log \frac{1}{\lambda})^d \log \frac{2}{\delta}}{n} \right)^{1/2} \quad (\text{E.87})$$

$C_3 = 2C_1 \text{vol}(\mathcal{X})(2C + C^2)$, $C_4 = 5C'_4$ where C'_4 is from [Lemma E.4](#) and depends only on d , while $C_5 = 50C'_5$, where C'_5 is from [Lemma E.5](#) and depends only on d . Setting $\varepsilon = n^{-\frac{\beta}{2\beta+d}}$, $\tau = \varepsilon^{-2/\beta}$ and $\lambda = \varepsilon^{2+2d/\beta} = n^{-(2\beta+2d)/(2\beta+d)}$, since $1/(n\lambda^{1/4}) = \varepsilon$ and $\varepsilon^{d/\beta}/n = \varepsilon^2$, then

$$\|\hat{p} - p\|_{L^2(\mathcal{X})} \leq (1 + C_3 e^{C_2}) \varepsilon + C_3 e^{C_2} \varepsilon^{d/\beta} e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} \quad (\text{E.88})$$

$$+ C_4 \log \frac{2}{\delta} \varepsilon + C_5^{1/2} \left(\frac{2\beta+2d}{2\beta+d} \right)^{d/2} \varepsilon (\log n)^{d/2} (\log \frac{2}{\delta})^{1/2}. \quad (\text{E.89})$$

Step 3. Controlling the number m in terms of h .

The final result is obtained by controlling the number of points m such that $h \leq \frac{1}{C' \sqrt{\tau}}$ (in order to be able to apply [Thm. C.4](#)) and such that $h \leq \frac{c}{\sqrt{\tau}} \max(e, (1 + \frac{d}{\beta}) \log \frac{1}{\varepsilon})^{-1}$, so $\log \frac{c\sigma}{h} \geq 1$ and

$$\varepsilon^{d/\beta} e^{-\frac{c\sigma}{h} \log \frac{c\sigma}{h}} \leq \varepsilon^{d/\beta} e^{-\frac{c\sigma}{h}} \leq \varepsilon. \quad (\text{E.90})$$

By, e.g. [Lemma 12](#), page 19 of [\[37\]](#) and the fact that $[-1, 1]^d$ is a convex set, we have that there exists two constants C_6, C_7 depending only on d such that $h \leq C_6 m^{-1/d} (\log(C_7 m/\delta))^{1/d}$, with probability at least $1 - \delta$. In particular, we want to find m that satisfy

$$C_6 C_8^d \tau^{d/2} (\log 1/\varepsilon)^d m \geq \log \frac{C_7 m}{\delta}, \quad (\text{E.91})$$

with $C_8 = \max(\frac{1}{c}, C', e, 1 + d/\beta)$. With the same reasoning as in [\(C.42\)](#), we see that any m satisfying $m \geq 2B \log(2AB)$ with $A = C_6 C_8^d \tau^{d/2} (\log 1/\varepsilon)^d$ and $B = C_7/\delta$ suffices to guarantee the inequality above. In particular, since $\varepsilon \leq n$ and $\varepsilon^{d/\beta} = n^{\frac{d}{2\beta+d}} \leq n^d$ we choose

$$m = C_9 n^{\frac{d}{2\beta+d}} (\log n)^d d \log(C_{10}^{1/d} n (\log n)). \quad (\text{E.92})$$

with $C_9 = C_6 C_8^d$, $C_{10} = 2C_9 C_7$. With this choice, we have

$$\|\hat{p} - p\|_{L^2(\mathcal{X})} \leq (1 + 2C_3 e^{C_2}) \varepsilon + C_4 \varepsilon \log \frac{2}{\delta} + C_5^{\frac{1}{2}} \left(\frac{2\beta+2d}{2\beta+d} \right)^{\frac{d}{2}} \varepsilon (\log n)^{\frac{d}{2}} (\log \frac{2}{\delta})^{\frac{1}{2}} \quad (\text{E.93})$$

$$\leq (1 + 2C_3 e^{C_2} + C_4 + C_5^{\frac{1}{2}} \left(\frac{2\beta+2d}{2\beta+d} \right)^{\frac{d}{2}}) \varepsilon (\log n)^{\frac{d}{2}} (\log \frac{2}{\delta}) \quad (\text{E.94})$$

$$= C_{11} n^{-\frac{\beta}{2\beta+d}} (\log n)^{\frac{d}{2}} (\log \frac{2}{\delta}). \quad (\text{E.95})$$

with $C_{11} = 1 + 2C_3 e^{C_2} + C_4 + C_5^{\frac{1}{2}} \left(\frac{2\beta+2d}{2\beta+d} \right)^{\frac{d}{2}}$. \square

F Operations

We report here the derivation of the operations discussed in [Sec. 2.1](#) for PSD models. For simplicity in the following, given a vector $x \in \mathbb{R}^d$ and a positive vector $\eta \in \mathbb{R}_{++}^n$, with some abuse of notation, in the following we will denote $\eta \|x\|^2$ for $x^\top \text{diag}(\eta)x$ when clear from the context.

F.1 Properties of the Gaussian Function

We recall the following classical properties of Gaussian functions, which are key to derive the results in the following. For any two points $x_1, x_2 \in \mathbb{R}^d$ and $\eta_1, \eta_2 \in \mathbb{R}_{++}^d$, with the notation introduced above, let $c_\eta = \pi^{d/2} \det(\text{diag}(\eta))^{-1/2}$ and $x_3 = \frac{\eta_1 x_1 + \eta_2 x_2}{\eta_1 + \eta_2}$ and $\eta_3 = \frac{\eta_1 \eta_2}{\eta_1 + \eta_2}$. We have

$$k_{\eta_1}(x, x_1) k_{\eta_2}(x, x_2) = k_{\eta_1 + \eta_2}(x, x_3) k_{\eta_3}(x_1, x_2) \quad \text{and} \quad \int k_\eta(x, x_1) dx = c_\eta. \quad (\text{F.1})$$

Additionally, we recall that the joint Gaussian kernel corresponds to the product kernels in the two variables, namely $k_{(\eta_1, \eta_2)}((x, y), (x', y')) = k_{\eta_1}(x, x')k_{\eta_2}(y, y')$.

We begin by recalling how the equalities in (F.1) can be derived. First, we recall that for any positive definite matrix $A \in \mathbb{S}_{++}^n$ (namely an invertible positive semi-definite matrix), the integral of the Gaussian function $e^{-\langle x, Ax \rangle}$ is

$$\int_{\mathbb{R}^d} e^{-\eta \|x\|^2} dx = \pi^{d/2} \det A^{-1/2}, \quad (\text{F.2})$$

which yields the required equality in (F.1) for $A = \text{diag}(\eta)$ and $\eta \in \mathbb{R}_{++}^d$.

For the second property in (F.1), let $x_1, x_2 \in \mathbb{R}^d$ and $\eta_1, \eta_2 \in \mathbb{R}_{++}^d$. For any $x \in \mathbb{R}^d$ we have

$$k_{\eta_1}(x, x_1)k_{\eta_2}(x, x_2) = e^{-\eta_1 \|x-x_1\|^2 - \eta_2 \|x-x_2\|^2}. \quad (\text{F.3})$$

By expanding the argument in the exponent we have

$$\eta_1 \|x - x_1\|^2 + \eta_2 \|x - x_2\|^2 \quad (\text{F.4})$$

$$= (\eta_1 + \eta_2) \|x\|^2 - 2 \langle x, \eta_1 x_1 + \eta_2 x_2 \rangle + \eta_1 \|x_1\|^2 + \eta_2 \|x_2\|^2 \quad (\text{F.5})$$

$$= (\eta_1 + \eta_2) \left(\|x\|^2 - 2 \left\langle x, \frac{\eta_1 x_1 + \eta_2 x_2}{\eta_1 + \eta_2} \right\rangle \pm \left\| \frac{\eta_1 x_1 + \eta_2 x_2}{\eta_1 + \eta_2} \right\|^2 \right) \quad (\text{F.6})$$

$$+ \eta_1 \|x_1\|^2 + \eta_2 \|x_2\|^2 \quad (\text{F.7})$$

$$= (\eta_1 + \eta_2) \|x - x_3\|^2 + \eta_1 \|x_1\|^2 + \eta_2 \|x_2\|^2 - (\eta_1 + \eta_2) \|x_3\|^2, \quad (\text{F.8})$$

where $x_3 = \frac{\eta_1 x_1 + \eta_2 x_2}{\eta_1 + \eta_2}$. Now,

$$\eta_1 \|x_1\|^2 + \eta_2 \|x_2\|^2 - (\eta_1 + \eta_2) \|x_3\|^2 \quad (\text{F.9})$$

$$= \eta_1 \|x_1\|^2 + \eta_2 \|x_2\|^2 - \frac{\|\eta_1 x_1 + \eta_2 x_2\|^2}{\eta_1 + \eta_2} \quad (\text{F.10})$$

$$= \eta_1 \|x_1\|^2 + \eta_2 \|x_2\|^2 - \frac{\eta_1^2 \|x_1\|^2 + 2\eta_1 \eta_2 \langle x_1, x_2 \rangle + \eta_2^2 \|x_2\|^2}{\eta_1 + \eta_2} \quad (\text{F.11})$$

$$= \frac{\eta_1 \eta_2 \|x_1\|^2 - 2\eta_1 \eta_2 \langle x_1, x_2 \rangle + \eta_1 \eta_2 \|x_2\|^2}{\eta_1 + \eta_2} \quad (\text{F.12})$$

$$= \frac{\eta_1 \eta_2}{\eta_1 + \eta_2} \|x_1 - x_2\|. \quad (\text{F.13})$$

We conclude that, for $\eta_3 = \frac{\eta_1 \eta_2}{\eta_1 + \eta_2}$, we have

$$k_{\eta_1}(x, x_1)k_{\eta_2}(x, x_2) = e^{-\eta_1 \|x-x_1\|^2 - \eta_2 \|x-x_2\|^2} \quad (\text{F.14})$$

$$= e^{-(\eta_1 + \eta_2) \|x-x_3\|^2 - \frac{\eta_1 \eta_2}{\eta_1 + \eta_2} \|x_1 - x_2\|^2} \quad (\text{F.15})$$

$$= k_{(\eta_1 + \eta_2)}(x, x_3)k_{\eta_3}(x_1, x_2), \quad (\text{F.16})$$

as required.

Finally, we recall that for any vectors $x, x' \in \mathbb{R}^{d_1}$ and $y, y' \in \mathbb{R}^{d_2}$ and positive weights $\eta_1 \in \mathbb{R}_{++}^{d_1}, \eta_2 \in \mathbb{R}_{++}^{d_2}$, we will make use of the fact that

$$k_{(\eta_1, \eta_2)}((x, y), (x', y')) = e^{-(x-x', y-y')^\top \text{diag}(\eta_1, \eta_2)(x-x', y-y')} \quad (\text{F.17})$$

$$= e^{-(x-x')^\top \text{diag}(\eta_1)x - y^\top \text{diag}(\eta_2)y'} \quad (\text{F.18})$$

$$= k_{\eta_1}(x, x')k_{\eta_2}(y, y') \quad (\text{F.19})$$

F.2 Evaluation

We recall that, given a PSD model $f(x; A, X, \eta)$ evaluating it in a point x_0 writes as

$$f(x = x_0; A, X, \eta) = K_{X, x_0, \eta}^\top A K_{X, x_0, \eta}. \quad (\text{F.20})$$

Given a PSD model of the form $f(x, y; A, [X, Y], (\eta_1, \eta_2))$ we denote partial evaluation in a vector $y_0 \in \mathcal{Y}$ as

$$f(x, y = y_0; A, [X, Y], (\eta_1, \eta_2)) = \sum_{i,j=1}^n A_{i,j} k_{\eta_1}(x_i, x) k_{\eta_1}(x_j, x) k_{\eta_2}(y_i, y_0) k_{\eta_2}(y_j, y_0) \quad (\text{F.21})$$

$$= \sum_{i,j=1}^n \left[k_{\eta_2}(y_i, y_0) A_{i,j} k_{\eta_2}(y_j, y_0) \right] k_{\eta_1}(x_i, x) k_{\eta_1}(x_j, x) \quad (\text{F.22})$$

$$= f(x; A \circ (K_{Y, y_0, \eta_2} K_{Y, y_0, \eta_2}^\top), X, \eta_1) \quad (\text{F.23})$$

F.3 Integration and Marginalization

We begin by showing the result characterizing the marginalization of a PSD model with respect to a number of random variables.

Proposition 1 (Sum Rule – Marginalization). *Let $X \in \mathbb{R}^{n \times d}$, $Y \in \mathbb{R}^{n \times d'}$, $A \in \mathbb{S}_+(\mathbb{R}^n)$ and $\eta \in \mathbb{R}_{++}^d$, $\eta' \in \mathbb{R}_{++}^{d'}$. Then, the following integral is a PSD model*

$$\int f(x, y; A, [X, Y], (\eta, \eta')) dx = f(y; B, Y, \eta'), \quad \text{with} \quad B = c_{2\eta} A \circ K_{X, X, \frac{\eta}{2}}, \quad (5)$$

Proof. The result is obtained as follows

$$\int f(x, y; A, [X, Y], (\eta, \eta')) dy = \sum_{i,j=1}^n A_{i,j} k_{\eta'}(y_i, y) k_{\eta'}(y_j, y) \int k_{\eta}(x_i, y) k_{\eta}(x_j, y) dx \quad (\text{F.24})$$

$$= \sum_{i,j=1}^n A_{i,j} k_{\frac{\eta}{2}}(x_i, x_j) k_{\eta'}(y_i, y) k_{\eta'}(y_j, y) \int k_{2\eta}\left(\frac{x_i + x_j}{2}, x\right) dx \quad (\text{F.25})$$

$$= c_{2\eta} \sum_{i,j=1}^n \left[A_{i,j} k_{\frac{\eta}{2}}(x_i, x_j) \right] k_{\eta'}(y_i, y) k_{\eta'}(y_j, y) \quad (\text{F.26})$$

$$= f(y; c_{2\eta}(A \circ K_{X, X, \frac{\eta}{2}}), Y, \eta') \quad (\text{F.27})$$

□

Integration. Analogously, we can write in matrix form the full integral of a PSD model (with respect to all its variables): let $f(x; A, X, \eta)$, we have

$$\int f(x; A, X, \eta) dx = \sum_{i,j=1}^n A_{i,j} \int k_{\eta}(x_i, x) k_{\eta}(x_j, x) dx \quad (\text{F.28})$$

$$= \sum_{i,j=1}^n A_{i,j} k_{\frac{\eta}{2}}(x_i, x_j) \int k_{2\eta}\left(\frac{x_i + x_j}{2}, x\right) dx \quad (\text{F.29})$$

$$= c_{2\eta} \sum_{i,j=1}^n A_{i,j} k_{\frac{\eta}{2}}(x_i, x_j) \quad (\text{F.30})$$

$$= c_{2\eta} \text{Tr}(A K_{X, X, \frac{\eta}{2}}), \quad (\text{F.31})$$

which yields (4).

Integration on the Hypercube. In Remark 3 we commented upon restricting integration and marginalization on the hypercube $H = \prod_{t=1}^d [a_t, b_t]$. Both operations can be performed by slightly changing the integrals above. In particular, let $G \in \mathbb{R}^{n \times n}$ be the matrix with i, j -th entry equal to

$$G_{ij} = c_{2\eta} \prod_{t=1}^d \text{erf} \left(\sqrt{2\eta_t} \left(b_t - \frac{x_{i,t} + x_{j,t}}{2} \right) \right) - \text{erf} \left(\sqrt{2\eta_t} \left(a_t - \frac{x_{i,t} + x_{j,t}}{2} \right) \right) \quad (\text{F.32})$$

Then, integration becomes

$$\int_H f(X; A, X, \eta; =) \text{Tr}((A \circ K_{X, X, \frac{\eta}{2}})G) \quad (\text{F.33})$$

and marginalization

$$\int f(x, y; A, [X, Y], (\eta, \eta')) dy = f(y; c_{2\eta}A \circ K_{X, X, \frac{\eta}{2}} \circ G.) \quad (\text{F.34})$$

This result is a corollary of [Prop. 9](#) that we prove below.

Proposition 9. *Let $p(x) = f(x; A, X, \eta)$ with $X \in \mathbb{R}^{n \times d}$, $A \in \mathbb{S}_+^n$, $\eta \in \mathbb{R}_{++}^d$. Let $g: \mathbb{R}^d \rightarrow \mathbb{R}$ and define $c_{g, \eta}(z) = \int g(x) e^{-\eta \|x-z\|^2} dx$ for any $z \in \mathbb{R}^d$. Then*

$$\mathbb{E}_{x \sim p} g(x) = \text{Tr}((A \circ K_{X, X, \eta/2})G) \quad \text{with} \quad G_{ij} = c_{g, 2\eta}\left(\frac{x_i + x_j}{2}\right). \quad (\text{18})$$

Proof. Following the same proof for [Prop. 1](#), we have

$$\int g(x) f(x; A, [X, Y], (\eta)) dy = \sum_{i,j=1}^n A_{ij} \int g(x) k_{\eta}(x_i, y) k_{\eta}(x_j, y) dx \quad (\text{F.35})$$

$$= \sum_{i,j=1}^n A_{ij} k_{\frac{\eta}{2}}(x_i, x_j) \int g(x) k_{2\eta}\left(\frac{x_i + x_j}{2}, x\right) dx \quad (\text{F.36})$$

$$= \sum_{i,j=1}^n \left[A_{ij} k_{\frac{\eta}{2}}(x_i, x_j) \right] c_{g, 2\eta}\left(\frac{x_i + x_j}{2}\right) \quad (\text{F.37})$$

$$= \text{Tr}((A \circ K_{X, X, \eta/2})G). \quad (\text{F.38})$$

□

F.4 Multiplication

Proposition 2 (Multiplication). *Let $X \in \mathbb{R}^{n \times d_1}$, $Y \in \mathbb{R}^{n \times d_2}$, $Y' \in \mathbb{R}^{m \times d_2}$, $Z \in \mathbb{R}^{m \times d_3}$, $A \in \mathbb{S}_+^n$, $B \in \mathbb{S}_+^m$ and $\eta_1 \in \mathbb{R}_{++}^{d_1}$, $\eta_2, \eta'_2 \in \mathbb{R}_{++}^{d_2}$, $\eta_3 \in \mathbb{R}_{++}^{d_3}$. Then*

$$f(x, y; A, [X, Y], (\eta_1, \eta_2)) f(y, z; B, [Y', Z], (\eta'_2, \eta_3)) = f(x, y, z; C, W, \eta), \quad (6)$$

is a PSD model, where $C = (A \otimes B) \circ (\text{vec}(K_{Y, Y', \tilde{\eta}_2}) \text{vec}(K_{Y, Y', \tilde{\eta}_2})^\top)$, with $\tilde{\eta}_2 = \frac{\eta_2 \eta'_2}{\eta_2 + \eta'_2}$, base matrix $W = [X \otimes \mathbf{1}_m, \frac{\eta_2}{\eta_2 + \eta'_2} Y \otimes \mathbf{1}_m + \frac{\eta'_2}{\eta_2 + \eta'_2} \mathbf{1}_n \otimes Y', \mathbf{1}_n \otimes Z]$ and $\eta = (\eta_1, \eta_2 + \eta'_2, \eta_3)$.

Proof. We begin by explicitly writing the product between the two PSD models

$$f(x, y; A, [X, Y], (\eta_1, \eta_2)) f(y, z; B, [Y', Z], (\eta'_2, \eta_3)) \quad (\text{F.39})$$

$$= \left(\sum_{i,j=1}^n A_{ij} k_{\eta_1}(x_i, x) k_{\eta_1}(x_j, x) k_{\eta_2}(y_i, x) k_{\eta_2}(y_j, y) \right) \quad (\text{F.40})$$

$$\left(\sum_{\ell, h=1}^n B_{\ell h} k_{\eta'_2}(y'_\ell, y) k_{\eta'_2}(y'_h, y) k_{\eta_3}(z_\ell, z) k_{\eta_3}(z_h, z) \right) \quad (\text{F.41})$$

$$= \sum_{i,j,\ell,h=1}^{n,m} [A_{ij} B_{\ell h} k_{\tilde{\eta}_2}(y_i, y'_\ell) k_{\tilde{\eta}_2}(y_j, y'_h)] k_{\eta_2 + \eta'_2}\left(\frac{\eta_2 y_i + \eta'_2 y'_\ell}{\eta_2 + \eta'_2}, y\right) k_{\eta_2 + \eta'_2}\left(\frac{\eta_2 y_j + \eta'_2 y'_h}{\eta_2 + \eta'_2}, y\right) \quad (\text{F.42})$$

$$k_{\eta_1}(x_i, x) k_{\eta_1}(x_j, x) k_{\eta_3}(z_\ell, z) k_{\eta_3}(z_h, z), \quad (\text{F.43})$$

where we have coupled together the pairs (i, ℓ) and (j, h) using the product rule between Gaussian functions. Let

$$C = (A \otimes B) \circ (\text{vec}(K_{Y, Y', \tilde{\eta}_2}) \text{vec}(K_{Y, Y', \tilde{\eta}_2})^\top), \quad (\text{F.44})$$

and denote by $\mathbf{i} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ now the indexing function such that $\mathbf{i}(i, \ell) = (i-1)n + \ell$. It follows that the term

$$A_{ij} B_{\ell h} k_{\tilde{\eta}_2}(y_i, y'_\ell) k_{\tilde{\eta}_2}(y_j, y'_h) = C_{\mathbf{i}(i, \ell) \mathbf{i}(j, h)}, \quad (\text{F.45})$$

corresponds to the $(\mathbf{i}(i, \ell) \mathbf{i}(j, h))$ -th entry of the matrix C . Analogously, let:

- $\tilde{y}_{\mathbf{i}(i, \ell)} = \frac{\eta_2 y_i + \eta'_2 y'_\ell}{\eta_2 + \eta'_2}$ is the $\mathbf{i}(i, \ell)$ -th row of the matrix $\tilde{Y} = \frac{\eta_2}{\eta_2 + \eta'_2} Y \otimes \mathbf{1}_m + \frac{\eta'_2}{\eta_2 + \eta'_2} \mathbf{1}_n \otimes Y'$, which is the $nm \times d_2$ matrix whose rows correspond to all possible pairs from Y and Y' respectively.
- $\tilde{x}_{\mathbf{i}(i, \ell)} = x_i$ is the $\mathbf{i}(i, \ell)$ -th row of the matrix $\tilde{X} = X \otimes \mathbf{1}_m$, namely the $nm \times d_1$ matrix containing m copies of each row of X .
- $\tilde{z}_{\mathbf{i}(i, \ell)} = z_\ell$ is the $\mathbf{i}(i, \ell)$ -th row of the matrix $\tilde{Z} \mathbf{1}_n \otimes Z$, namely the $nm \times d_3$ matrix containing n copies of Z .

Then we have that

$$\sum_{i, j, \ell, h=1}^{n, m} [A_{ij} B_{\ell h} k_{\tilde{\eta}_2}(y_i, y'_\ell) k_{\tilde{\eta}_2}(y_j, y'_h)] k_{\eta_2 + \eta'_2}(\frac{\eta_2 y_i + \eta'_2 y'_\ell}{\eta_2 + \eta'_2}, y) k_{\eta_2 + \eta'_2}(\frac{\eta_2 y_j + \eta'_2 y'_h}{\eta_2 + \eta'_2}, y) \quad (\text{F.46})$$

$$k_{\eta_1}(x_i, x) k_{\eta_1}(x_j, x) k_{\eta_3}(z_\ell, z) k_{\eta_3}(z_h, z) \quad (\text{F.47})$$

$$= \sum_{i, j, \ell, h=1}^{n, m} C_{\mathbf{i}(i, \ell) \mathbf{i}(j, h)} k_{\eta_2 + \eta'_2}(\tilde{y}_{\mathbf{i}(i, \ell)}, y) k_{\eta_2 + \eta'_2}(\tilde{y}_{\mathbf{i}(j, h)}, y) k_{\eta_1}(\tilde{x}_{\mathbf{i}(i, \ell)}, x) k_{\eta_1}(\tilde{x}_{\mathbf{i}(j, h)}, x) \quad (\text{F.48})$$

$$k_{\eta_3}(\tilde{z}_{\mathbf{i}(i, \ell)}, x) k_{\eta_3}(\tilde{z}_{\mathbf{i}(j, h)}, x) \quad (\text{F.49})$$

$$= \sum_{s, t}^{nm} C_{st} k_{\eta_2 + \eta'_2}(\tilde{y}_s, y) k_{\eta_2 + \eta'_2}(\tilde{y}_t, y) k_{\eta_1}(\tilde{x}_s, x) k_{\eta_1}(\tilde{x}_t, x) k_{\eta_3}(\tilde{z}_s, x) k_{\eta_3}(\tilde{z}_t, x) \quad (\text{F.50})$$

$$= f(x, y, z; C, [\tilde{X}, \tilde{Y}, \tilde{Z}], (\eta_1, \eta_2 + \eta'_2, \eta_3)), \quad (\text{F.51})$$

as desired. \square

F.5 Reduction

The reduction operation leverages the structure of the base matrix $X \otimes \mathbf{1}_m$ to simplify the PSD model. To this end, denote by $\tilde{X} = X \otimes \mathbf{1}_m$ and consider again the indexing function $\mathbf{i}(i, \ell) = (i-1)n + \ell$. Then (see also the proof of [Prop. 2](#) we have that the $\mathbf{i}(i, \ell)$ -th row of \tilde{X} is $\tilde{x}_{\mathbf{i}(i, \ell)} = x_i$, the i -th row of X . Therefore we have

$$f(x; A, \tilde{X}, \eta) = \sum_{s, t=1}^{nm} A_{st} k_\eta(\tilde{x}_s, x) k_\eta(\tilde{x}_t, x) \quad (\text{F.52})$$

$$= \sum_{i, j, \ell, h=1}^{n, m} A_{\mathbf{i}(i, \ell) \mathbf{i}(j, h)} k_\eta(\tilde{x}_{\mathbf{i}(i, \ell)}, x) k_\eta(\tilde{x}_{\mathbf{i}(j, h)}, x) \quad (\text{F.53})$$

$$= \sum_{i, j, \ell, h=1}^{n, m} A_{\mathbf{i}(i, \ell) \mathbf{i}(j, h)} k_\eta(x_i, x) k_\eta(x_j, x) \quad (\text{F.54})$$

$$= \sum_{i, j=1}^n \left[\sum_{\ell, h=1}^m A_{\mathbf{i}(i, \ell) \mathbf{i}(j, h)} \right] k_\eta(x_i, x) k_\eta(x_j, x) \quad (\text{F.55})$$

$$= \sum_{i, j=1}^n B_{ij} k_\eta(x_i, x) k_\eta(x_j, x), \quad (\text{F.56})$$

where B is a $n \times n$ PSD matrix and each of its entries is the sum of the entries of A corresponding to the repeated rows in $X \otimes \mathbf{1}_m$. Therefore $B = (I_n \otimes \mathbf{1}_m^\top) A (I_n \otimes \mathbf{1}_m)$ as required.

We give here the explicit form for the Markov transition in [Cor. 3](#)

Corollary 3 (Markov Transition). *Let $X \in \mathbb{R}^{n \times d_1}$, $Y \in \mathbb{R}^{n \times d_2}$, $Y' \in \mathbb{R}^{m \times d_2}$, $A \in \mathbb{S}_+^n$, $B \in \mathbb{S}_+^m$ and $\eta_1 \in \mathbb{R}_{++}^{d_1}$, $\eta_2, \eta'_2 \in \mathbb{R}_{++}^{d_2}$. Then*

$$\int f(x, y; A, [X, Y], (\eta_1, \eta_2)) f(y; B, Y', \eta'_2) dy = f(x; C, X, \eta_1), \quad (8)$$

with $C \in \mathbb{S}_+^n$ obtained by applying in order, [Prop. 2](#), [Prop. 1](#) and reduction (7).

Proof. We first multiply the two PSD model to obtain, by [Prop. 2](#)

$$f(x, y; A, [X, Y], (\eta_1, \eta_2)) f(y; B, Y', \eta'_2) = f(x, y; A_1, [\tilde{X}, \tilde{Y}], (\eta, \eta_2 + \eta'_2)), \quad (F.57)$$

with

$$A_1 = (A \otimes B) \circ (\text{vec}(K_{Y, Y', \tilde{\eta}_2}) \text{vec}(K_{Y, Y', \tilde{\eta}_2})^\top) \quad \tilde{\eta}_2 = \frac{\eta_2 \eta'_2}{\eta_2 + \eta'_2}, \quad (F.58)$$

and $\tilde{X} = X \otimes \mathbf{1}_m$ and $\tilde{Y} = \frac{\eta_2}{\eta_2 + \eta'_2} Y \otimes \mathbf{1}_m + \frac{\eta'_2}{\eta_2 + \eta'_2} Y'$. Then, we proceed with marginalization. By [Prop. 1](#), we have

$$\int f(x, y; A_1, [\tilde{X}, \tilde{Y}], (\eta, \eta_2 + \eta'_2)) dy = f(x; A_2, \tilde{X}, \eta), \quad (F.59)$$

where $A_2 = c_{2(\eta_2 + \eta'_2)} A_1 \circ K_{\tilde{Y}, \tilde{Y}, (\eta_2 + \eta'_2)/2}$. Finally, since $\tilde{X} = X \otimes \mathbf{1}_m$, by reduction (7), we have

$$f(x; A_2, \tilde{X}, \eta) = f(x; C, X, \eta) \quad (F.60)$$

with $C = (I_n \otimes \mathbf{1}_m^\top) A_2 (I_n \otimes \mathbf{1}_m)$, which concludes the derivation. \square

F.6 Hidden Markov Models

We conclude this section by providing the derivation of the HMM inference in [Sec. 4.3](#).

Proposition 10 (PSD Hidden Markov Models (HMM)). *Let $X_0 \in \mathbb{R}^{n_0 \times d}$, $X_+, X \in \mathbb{R}^{n \times d}$, $X' \in \mathbb{R}^{m \times d}$, $Y \in \mathbb{R}^{m \times d'}$, $A_0 \in \mathbb{S}_+^{n_0}$, $A \in \mathbb{S}_+^n$, $B \in \mathbb{S}_+^m$ and $\eta_0, \eta, \eta', \eta_+ \in \mathbb{R}_{++}^d$, $\eta_{obs} \in \mathbb{R}_{++}^{d'}$. Let*

$$\hat{\tau}(x_+, x) = f(x_+, x; B, [X_+, X], (\eta_+, \eta)), \quad \hat{\omega}(y, x) = f(y, x; C, [Y, X'], (\eta_{obs}, \eta')), \quad (21)$$

be approximate transition and observation functions. Then, given the initial state probability $\hat{p}(x_0) = f(x_0; A_0, X_0, \eta_0)$, for any $t \geq 1$, the estimate \hat{p} in (20) is a PSD model of the form

$$\hat{p}(x_t | y_{t-1}) = f(x_t; A_t, \tilde{X}, \eta' + \eta_+), \quad (22)$$

where $\tilde{X} = \frac{\eta'}{\eta' + \eta_+} X' \otimes \mathbf{1}_n + \frac{\eta_+}{\eta' + \eta_+} \mathbf{1}_m \otimes X_+$ and A_t is recursively obtained from A_{t-1} as in [Alg. 1](#).

Proof. Let $f(x_t; A_{t-1}, X_{t-1}, \eta_{t-1})$ be the estimate $\hat{p}(x_{t-1} | y_{1:t-1})$ obtained at the previous step, with $A_{t-1} \in \mathbb{S}_+^{n_{t-1}}$ and $X_{t-1} \in \mathbb{R}^{n_{t-1} \times d}$. We then proceed by performing the operations in (20).

Observation $\hat{\omega}_t(x)$. A new observation y_t is received. By (3) we have

$$\hat{\omega}_t(x) = \hat{\omega}(y = y_t, x) = f(y = y_t, x; C, [Y, X'], (\eta_{obs}, \eta')) = f(x; C_t, X', \eta'), \quad (F.61)$$

with

$$C_t = C \circ (K_{Y, y_t, \eta_{obs}} K_{Y, y_t, \eta_{obs}}^\top). \quad (F.62)$$

Product $\hat{\beta}_t(x_+, x) = \hat{\tau}(x_+, x) \hat{p}(x | y_{1:t-1})$. We perform the product between the transition function and the previous state estimation

$$\hat{\beta}_t(x_+, x) = f(x_+, x; B, [X_+, X], (\eta_+, \eta)) f(x; A_{t-1}, X_{t-1}, \eta_{t-1}) \quad (F.63)$$

$$= f(x_+, x; B_t, [X_+ \otimes \mathbf{1}_{n_{t-1}}, \tilde{X}_t], (\eta_+, \eta + \eta_t)), \quad (F.64)$$

with

$$B_t = (B \otimes A_{t-1}) \circ (\text{vec}(K_{X, X_t, \tilde{\eta}_t}) \text{vec}(K_{X, X_t, \tilde{\eta}_t})^\top) \quad \tilde{\eta}_t = \tilde{\eta} \frac{\eta \eta_t}{\eta + \eta_t}, \quad (F.65)$$

and

$$\tilde{X}_t \frac{\eta}{\eta + \eta_t} X \otimes \mathbf{1}_{n_{t-1}} + \frac{\eta_t}{\eta + \eta_t} \mathbf{1}_n \otimes X_t. \quad (\text{F.66})$$

Marginalization (+ Reduction) $\hat{\beta}_t(x_+) = \int \hat{\beta}_t(x_+, x) dx$. We perform marginalization by [Prop. 1](#) to obtain

$$\hat{\beta}_t(x_+) = \int \hat{\beta}_t(x_+, x) dx \quad (\text{F.67})$$

$$= \int f(x_+, x; B_t, [X_+ \otimes \mathbf{1}_{n_{t-1}}, \tilde{X}_t], (\eta_+, \eta + \eta_t)) dx \quad (\text{F.68})$$

$$= f(x_+; D'_t, X_+ \otimes \mathbf{1}_{n_{t-1}}, \eta_+) \quad (\text{F.69})$$

with

$$D'_t = c_{2(\eta + \eta_t)} B_t \circ K_{\tilde{X}_t, \tilde{X}_t, \tilde{\eta}_t/2}. \quad (\text{F.70})$$

Since the PSD model has a redundant base point matrix, we can apply reduction from [\(7\)](#), to obtain

$$\hat{\beta}_t(x_+) = f(x_+; D'_t, X_+ \otimes \mathbf{1}_{n_{t-1}}, \eta_+) \quad (\text{F.71})$$

$$= f(x_+; D_t, X_+, \eta_+), \quad (\text{F.72})$$

where

$$D_t = (I_n \otimes \mathbf{1}_{n_{t-1}}^\top) D'_t (I_n \otimes \mathbf{1}_{n_{t-1}}). \quad (\text{F.73})$$

Multiplication $\hat{\pi}_t(x_+) = \hat{\omega}_t(x_+) \hat{\beta}_t(x_+)$. We now multiply the observation term with the state estimation to obtain

$$\hat{\pi}_t(x_+) = \hat{\omega}_t(x_+) \hat{\beta}_t(x_+) \quad (\text{F.74})$$

$$= f(x_+; C_t, X', \eta') f(x_+; D_t, X_+, \eta_+) \quad (\text{F.75})$$

$$= f(x_+; E_t, \tilde{X}, \eta' + \eta_+), \quad (\text{F.76})$$

with

$$E_t = (C_t \otimes D_t) \circ (\text{vec}(K_{X', X, \tilde{\eta}'}) \text{vec}(K_{X', X, \tilde{\eta}'})^\top) \quad \tilde{\eta}' = \frac{\eta' \eta_+}{\eta' + \eta_+} \quad (\text{F.77})$$

and $\tilde{X} = \frac{\eta'}{\eta' + \eta_+} X' \otimes \mathbf{1}_n + \frac{\eta_+}{\eta' + \eta_+} \mathbf{1}_m \otimes X_+$.

Normalization $\hat{p}(x_t | y_{1:t}) = \hat{\pi}_t(x_+) / \int \hat{\pi}_t(x_+) dx_+$. We finally integrate $\hat{\pi}_t(x_+)$ in order to normalize it. By [\(4\)](#) we have

$$c_t = \int \hat{\pi}_t(x_+) dx_+ = \int f(x_+; E_t, \tilde{X}, \eta' + \eta_+) dx_+ = c_{2(\eta' + \eta_+)} \text{Tr}(E_t K_{\tilde{X}, \tilde{X}, (\eta' + \eta_+)/2}), \quad (\text{F.78})$$

and therefore

$$\hat{p}(x_t | y_{1:t}) = \frac{\hat{\pi}_t(x_+)}{\int \hat{\pi}_t(x_+) dx_+} = f(x_t; A_t, \tilde{X}, \eta' + \eta_+), \quad (\text{F.79})$$

with

$$A_t = E_t / c_t. \quad (\text{F.80})$$

This concludes the proof showing that, at every step, $\hat{p}(x_t | y_{1:t})$ has always same base point matrix \tilde{X} and parameters $\eta' + \eta_+$. Note that the proof above also recovers explicitly the steps in [Alg. 1](#). \square