

---

# ParK: Sound and Efficient Kernel Ridge Regression by Feature Space Partitions

---

**Luigi Carratino\***

MaLGa - DIBRIS, University of Genova  
luigi.carratino@dibris.unige.it

**Stefano Vigogna\***

MaLGa - DIBRIS, University of Genova  
vigogna@dibris.unige.it

**Daniele Calandriello**

DeepMind Paris  
dcalandriello@google.com

**Lorenzo Rosasco**

MaLGa - DIBRIS, University of Genova  
IIT, CBMM - MIT  
lrosasco@mit.edu

## Abstract

We introduce ParK, a new large-scale solver for kernel ridge regression. Our approach combines partitioning with random projections and iterative optimization to reduce space and time complexity while provably maintaining the same statistical accuracy. In particular, constructing suitable partitions directly in the feature space rather than in the input space, we promote orthogonality between the local estimators, thus ensuring that key quantities such as local effective dimension and bias remain under control. We characterize the statistical-computational tradeoff of our model, and demonstrate the effectiveness of our method by numerical experiments on large-scale datasets.

## 1 Introduction

The development of provably accurate and efficient algorithms for learning is key to tackle modern large-scale applications. Kernel methods [31, 32] provide a natural ground to develop this research direction. On the one hand they have sound statistical guarantees [7, 32, 33], but on the other hand their basic implementations are limited to sample size of only a few tens of thousands of points [32, Chapter 11]. Recent years have witnessed a growing literature introducing algorithmic solutions to improve efficiency, but also theoretical guarantees that quantify how accuracy is affected.

We next recall a few lines of work relevant to our study. A first line of work is based on exploiting ideas from optimization and numerical analysis. This includes for example gradient methods [37], as well as their accelerated [3], stochastic [10], preconditioned [13] and distributed [27] variants. A second line of work is based on using randomized approaches to reduce the size of the problem to be solved. This includes Nyström approximations [36], random features [25] and more generally sketching methods [1]. The theoretical properties of these methods have been recently characterized in terms of sharp statistical bounds [28, 30]. Finally, a third line of work considers different partitioning strategies to divide the estimation step in smaller subproblems. This approach is based on splitting the input space in regions where local estimators are defined [23, 35, 34, 11, 24, 5]. In this context, the emphasis is typically on allowing the estimation of larger classes of functions. Another form of partitioning, called divide-and-conquer, is instead based on randomly splitting the training data to then obtain a global estimator by averaging [38, 20, 14]. In this approach, the focus is primarily on computational saving. Notably, a number of works have considered combinations of these ideas, see for example [6, 29, 8, 24, 19].

---

\*equal contribution

In this paper we propose and study a local kernel algorithm, called ParK, combining partitioning with iterative optimization and sketching. Our goal is to provide an efficient and accurate approximation to a global kernel ridge regression estimator. The main novelty in ParK is in the form of the considered partition, that we define in the feature space, rather than in the input space as in traditional partitioning methods. This allows to promote orthogonality between the local estimators, and thus to control the local effective dimension and the local bias. Given a partition, local kernel ridge estimators are computed using sketching and preconditioned conjugate gradient iterations [29]. From a theoretical point of view, our main contribution is characterizing the statistical properties of ParK, in terms of conditions on the partition and the choice of the hyper-parameters. Borrowing ideas from subspace clustering [12], we show that the minimal angle between suitable subspaces induced by the partition plays a crucial role. Indeed, our analysis shows that, if such an angle is sufficiently large, ParK can achieve the same accuracy as global kernel ridge regression estimators, with only a fraction of computations. Our theoretical results are complemented with numerical experiments on very large datasets, which show that ParK can indeed provide excellent performances, on par and often better than the best available large-scale kernel methods.

The rest of the paper is organized as follows. In Section 2 we state the problem and recall the basics of kernel ridge regression. In Section 3 we illustrate our algorithm. In Section 4 we analyze the prediction error of our method. In Section 5 we present the results of our numerical experiments. In Section 6 we draw some conclusions and report the main limitations of our work. Additional proofs and details are collected in Appendix A.

## 2 Background

Let  $(x_i, y_i)$  with  $i \in [n] = \{1, \dots, n\}$  be  $n$  pairs of points in  $\mathcal{X} \times \mathcal{Y}$ , where  $\mathcal{X} \subseteq \mathbb{R}^d$  with  $d \in \mathbb{N}$  and  $\mathcal{Y} \subseteq \mathbb{R}$ . We assume the relation between input points  $x_i$  and output points  $y_i$  to be determined by the noisy evaluations of an unknown function  $f_* : \mathcal{X} \rightarrow \mathcal{Y}$  as

$$y_i = f_*(x_i) + \varepsilon_i \quad i \in [n]. \quad (1)$$

Based on the samples  $(x_i, y_i)$ , we want to estimate the function  $f_*$ , searching for solutions in a suitable hypothesis space  $\mathcal{H}$  as detailed below.

Let  $\mathcal{H}$  be a reproducing kernel Hilbert space (RKHS), that is, a Hilbert space of functions with inner product  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$  and symmetric positive definite kernel  $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  such that  $K_x = K(x, \cdot) \in \mathcal{H}$  and  $f(x) = \langle f, K_x \rangle_{\mathcal{H}}$  for all  $f \in \mathcal{H}, x \in \mathcal{X}$ . We recall that, for every RKHS  $\mathcal{H}$ , there exist a Hilbert feature space  $\mathcal{F}$  and a feature map  $\phi : \mathcal{X} \rightarrow \mathcal{F}$  such that  $K(x, x') = \langle \phi(x), \phi(x') \rangle_{\mathcal{F}}$  for all  $x, x' \in \mathcal{X}$ . The feature map is not unique; in particular, one may take, as we do in all that follows,  $\mathcal{F} = \mathcal{H}$  and  $\phi(x) = K_x$ , in which case  $\mathcal{H} = \overline{\text{span}} \phi(\mathcal{X})$ , where  $\phi(\mathcal{X}) = \{\phi(x) : x \in \mathcal{X}\}$ .

Kernel ridge regression (KRR) corresponds to minimizing

$$\min_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n |f(x_i) - y_i|^2 + \lambda \|f\|_{\mathcal{H}}^2, \quad (2)$$

where  $\lambda > 0$  and  $\|f\|_{\mathcal{H}}^2 = \langle f, f \rangle_{\mathcal{H}}$ . By the representer theorem [31], the (unique) solution to problem (2) can be written as

$$\hat{f}_\lambda(x) = \sum_{i=1}^n \alpha_i K(x_i, x), \quad \alpha = (K_n + \lambda n I)^{-1} Y, \quad (3)$$

where  $\alpha = [\alpha_1, \dots, \alpha_n]^\top$ ,  $Y = [y_1, \dots, y_n]^\top \in \mathbb{R}^n$ , and  $K_n \in \mathbb{R}^{n \times n}$  is the kernel matrix defined by  $(K_n)_{i,j} = K(x_i, x_j)$  for  $i, j \in [n]$ . As a consequence, the estimator (3) can be derived restricting the minimization problem (2) to the finite-dimensional subspace  $\mathcal{H}_n = \text{span}\{\phi(x_i) \mid i \in [n]\}$ . Computing (3) for large  $n$  is prohibitively expensive, as space and time complexities are, respectively,  $O(n^2)$  and  $O(n^3)$ . The goal of this paper is to provide an algorithm to compute an efficient approximation to (3).

## 3 Algorithm

Our method combines diverse techniques, including partitioning, sketching and preconditioned iterative optimization. We begin focusing on partitioning. While classical partitioning methods

construct partitions in the input space, the main novelty of our approach is that we construct partitions in the feature space. Note that, in the case of a universal kernel on a compact input space, every feature map is injective [32, Lemma 4.55], hence every partition of the input space defines a corresponding partition of the feature space. Thus, we may see our approach as a generalization of classical input space partitioning approaches. As will become apparent from our analysis, the performance of a partitioned kernel estimator depends crucially on two main quantities: the local biases and the local effective dimensions. Since both quantities are strictly related to the RKHS of choice, constructing partitions in feature space allows for a more direct control. In particular, promoting orthogonality in the RKHS metric will generate feature space partitions which tend to minimize both the local biases and the local effective dimensions. In the next section we start discussing how such partitions can be defined.

### 3.1 Learning on feature space partitions

For  $Q \in \mathbb{N}$ , we define a partition of  $\phi(\mathcal{X})$  as a family  $\{V_q\}_{q \in [Q]}$  of subsets  $V_q \subseteq \phi(\mathcal{X})$  such that

$$\phi(\mathcal{X}) = \bigcup_{q=1}^Q V_q \quad V_q \cap V_k = \emptyset \quad q \neq k. \quad (4)$$

The partition (4) induces a local subsampling of the training set and a local hypothesis space. Namely, we define

$$[n]_q = \{i \in [n] : \phi(x_i) \in V_q\}, \quad \mathcal{H}_q = \overline{\text{span}}\{V_q\}.$$

Also, we denote by  $n_q = \#[n]_q$  the local subsampling rate.

**Voronoi partitions.** Notice that so far  $V_q$  is an arbitrary subset of  $\mathcal{H}$ , and therefore computing the set  $[n]_q$  can be arbitrarily difficult (e.g.,  $V_q$  could be defined using an infinite number of constraints and be non-computable). For this reason, although our theoretical analysis holds for any partition defined as in (4), we focus on the special case of Voronoi partitions, where the subsets (also called cells) are induced by a set of  $Q$  centroids  $\{\phi(c_q)\}_{q=1}^Q$  with  $c_q \in \mathcal{X}$  points in the input space. Then, each cell  $V_q$  is uniquely defined as

$$V_q = \{\phi(x) : q = \arg \min_{k \in [Q]} \|\phi(x) - \phi(c_k)\|_{\mathcal{H}}^2\},$$

with ties broken arbitrarily (e.g., by assigning the point to the cell with the smaller  $q$ ). It is now possible to identify the set of indices  $[n]_q$  using the RKHS distance

$$\|\phi(x) - \phi(x')\|_{\mathcal{H}}^2 = K(x, x) + K(x', x') - 2K(x, x'), \quad (5)$$

computing the distance to each centroid and taking the minimum.

We remark that our approach based on directly partitioning the feature space has quite different implications compared to previous approaches that partition the input space. For example, a Voronoi partition of the feature space is very different from a Voronoi partition of the input space, since the pre-image  $\{x \in \mathcal{X} : \phi(x) \in V_q\}$  does not need to follow any Voronoi shape. Moving from input to feature space partitions also opens new computational challenges. For example, we choose to explicitly represent the cell centroid as  $\phi(c_q)$  so that computing the distance and the assignment of each point to a centroid is a  $O(1)$  operation. If instead we chose a more complex centroid, such as a cluster barycenter generated by kernel  $k$ -means, or an eigenvector computed by kernel PCA, this complexity might be much larger. As an example, the barycenter of a cluster of  $m$  points in  $\mathcal{H}$  might not correspond to any single point in  $\mathcal{X}$ , and therefore cannot be explicitly represented, but only implicitly as an average of  $m$  points in  $\mathcal{H}$ . Therefore, computing a distance to such a centroid would be an  $O(m)$  operation rather than a  $O(1)$ . These and more subtle pitfalls appear only when we consider the more flexible framework of feature space partitions.

**Minimal principal angle.** Focusing on partitions of Voronoi type, constructing a good partition is equivalent to choosing a set of centroids that preserve the learning accuracy as much as possible. As we rigorously show in Section 4, this can be guaranteed by choosing centroids that maximize the minimal principal angle between subspaces. This quantity frequently appears in the analysis of

subspace clustering [12], and will be important for us to control both the bias and the variance of our estimator. The first principal angle between two linear subspaces  $U$  and  $W$  of a Hilbert space of inner product  $\langle \cdot, \cdot \rangle$  and norm  $\| \cdot \|$  is defined as

$$\angle(U, W) = \min\{\arccos(\langle u, w \rangle) : u \in U, w \in W, \|u\| = \|w\| = 1\}.$$

We call  $\theta$  the minimal first principal angle between the subspaces  $\mathcal{H}_q$ , that is,

$$\theta = \min_{q \neq k} \angle(\mathcal{H}_q, \mathcal{H}_k). \quad (6)$$

Once again, for computational reasons we cannot use direct optimization of this quantity in  $\mathcal{H}$  as our objective, since the optimal centroid placement might be impossible to express using points from  $\mathcal{X}$ . Instead, to promote large principal angles and obtain centroids that are computationally easy to handle, we consider the following greedy iterative procedure. Let  $X = \{x_i : i \in [n]\}$ . Then

$$c_1 = \arg \max_{c \in X} K(c, c), \quad c_{q+1} = \arg \max_{c \in X \setminus \{c_1, \dots, c_q\}} \text{SC}_q(c), \quad (7)$$

where

$$\text{SC}_q(c) = K(c, c) - [K(c, c_1), \dots, K(c, c_q)]^\top K_q^{-1} [K(c, c_1), \dots, K(c, c_q)]$$

is the Schur complement of a new candidate centroid  $c$  with respect to the  $q$  already selected centroids  $\{c_1, \dots, c_q\}$ , and  $K_q \in \mathbb{R}^{q \times q}$  is defined by  $(K_q)_{i,j} = K(c_i, c_j)$ . Note that the inversion of  $K_q$  can be efficiently computed using rank-1 updates. This strategy has been originally proposed by [9] with the goal of maximizing the volume spanned by the points in the feature space, which is achieved when the angle between all points selected is large as required by our condition. Crucially, it is also easy to apply to RKHS's, since computing Schur complements involves only inner products. Beyond promoting large volume and orthogonality, the Schur complement also has important links with uncertainty estimation and spectral approximation. In particular,  $\text{SC}_q(c)$  is also equivalent to the posterior variance of  $c$  in a Gaussian process [26], and to the leverage score of  $c$  w.r.t. the already selected point in the context of randomized linear algebra [21].

### 3.2 Learning local KRR estimators by sketched preconditioned conjugate gradient

For each cell of a partition, a local estimator  $\hat{f}_q$  can be defined as the solution to the local KRR problem

$$\min_{f \in \mathcal{H}_q} \frac{1}{n_q} \sum_{i \in [n]_q} |f(x_i) - y_i|^2 + \lambda_q \|f\|_{\mathcal{H}}^2 \quad (8)$$

with  $\lambda_q > 0$ . Given the local estimators  $\hat{f}_q$ , we then define a global estimator  $\bar{f}$  by

$$\bar{f}(x) = \hat{f}_q(x) \quad \text{if } \phi(x) \in V_q. \quad (9)$$

Note that the evaluation of the global estimator at a point needs only one local estimator.

Guidance on how to pick the values  $\lambda_q$  in (8) will follow from our theoretical analysis. Meanwhile, we focus on how to efficiently solve the minimization problems (8). Let  $X_q = \{x_i \in X : i \in [n]_q\} \in \mathbb{R}^{n_q \times d}$  and  $Y_q = \{y_i \in Y : i \in [n]_q\} \in \mathbb{R}^{n_q}$  be the local subsets of input/output points, and let  $K_{n_q} \in \mathbb{R}^{n_q \times n_q}$  be the local kernel matrix with entries  $(K_{n_q})_{i,j} = K(x_i, x_j)$  for  $i, j \in [n]_q$ . Following the same ideas to derive (3), one could compute  $\hat{f}_q$  by

$$\hat{f}_q(x) = \sum_{i \in [n]_q} (\alpha_q)_i K(x_i, x), \quad \alpha_q = (K_{n_q} + \lambda_q n_q I)^{-1} Y_q. \quad (10)$$

This would already result in a smaller computational burden compared to the vanilla KRR estimator (3): the space and time complexities are now  $O(\max_{q \in [Q]} n_q^2)$  and  $O(\sum_{q \in [Q]} n_q^3)$ , potentially with  $n_q \ll n$ . Moreover, an additional saving in time can be obtained by distributing each task (10) over  $Q$  different machines, leading to  $O(\max_{q \in [Q]} n_q^3)$  time complexity. However, the scaling in  $n_q$  is still quadratic and cubical. To improve these dependencies, we solve (10) only approximately, using the FALKON algorithm proposed in [29]. To this end, we first need to introduce several key ingredients. While the following constructions hold in general for any set of points, here we adapt them to the partition setting outlined in the previous section.

**Local Nyström subsampling.** For each  $q \in [Q]$ , we consider a subset of  $m_q \leq n_q$  points

$$\tilde{X}_q = \{\tilde{x}_{q,1}, \dots, \tilde{x}_{q,m_q}\} \subseteq X_q \quad (11)$$

sampled uniformly at random from  $X_q$ . We then define  $K_{m_q} \in \mathbb{R}^{m_q \times m_q}$  by  $(K_{m_q})_{i,j} = K(\tilde{x}_{q,i}, \tilde{x}_{q,j})$  for  $i, j \in [m_q]$ , and  $K_{n_q m_q} \in \mathbb{R}^{n_q \times m_q}$  by  $(K_{n_q m_q})_{i,j} = K(x_i, \tilde{x}_{q,j})$  for  $i \in [n]_q, j \in [m_q]$ .

**Local Preconditioner.** For each  $q \in [Q]$ , we define the local (sketched) preconditioner  $B_q \in \mathbb{R}^{m_q \times m_q}$  as

$$B_q B_q^\top = \left( \frac{n_q}{m_q} K_{m_q}^2 + \lambda_q n_q K_{m_q} \right)^{-1}.$$

**Conjugate gradient descent.** We let  $\tilde{\beta}_{q,t} \in \mathbb{R}^{m_q}$  be the  $t$ -th iteration of conjugate gradient minimizing

$$\mathcal{L}_q(\beta) = \frac{1}{n_q} \|K_{n_q m_q} B_q \beta - Y_q\|^2 + \lambda_q \beta^\top (B_q^\top K_{m_q} B_q) \beta. \quad (12)$$

Finally, we define the local FALKON estimator

$$\tilde{f}_{q,t}(x) = \sum_{i=1}^{m_q} (B_q \tilde{\beta}_{q,t})_i K(\tilde{x}_i, x) \quad q \in [Q]. \quad (13)$$

### 3.3 ParK

We are now ready to present ParK. Let  $(\phi(c_q))_{q=1}^Q$  with  $c_q \in X$  be the centroids of the cells  $(V_q)_{q=1}^Q$  selected greedily according to (7). We define the ParK estimator as

$$\bar{f}_t(x) = \tilde{f}_{q,t}(x) \quad \text{if } \phi(x) \in V_q. \quad (14)$$

The algorithm to train the above estimator (see Algorithm 1) consists of three main parts. The first one greedily identifies the representative points  $(c_q)_{q=1}^Q$  such that  $(\phi(c_q))_{q=1}^Q$  are the centroids of the cells; the second one identifies the subsets of points  $X_q, Y_q$  associated to each cell; the third one uses the FALKON algorithm to solve the local minimization problem (12) for each  $X_q, Y_q$  with  $q \in [Q]$ , thus deriving the  $Q$  local estimators (13). At prediction time, the algorithm first identifies to which cell the test point belongs, and then proceeds using the local estimator of the selected cell to predict the output (see Algorithm 2). Note that the RKHS distances in line 4 of Algorithm 1 and line 1 of Algorithm 2) are computed using the polarization identity (5).

---

#### Algorithm 1 ParK: Train

---

**Require:** Training set  $X = (x_i)_{i=1}^n \in \mathbb{R}^{n \times d}, Y = (y_i)_{i=1}^n \in \mathbb{R}^n$ , numbers of local Nyström centers  $\{m_q\}_{q=1}^Q \in \mathbb{N}^Q$ , local regularization parameters  $\{\lambda_q\}_{q=1}^Q \in \mathbb{R}_+^Q$ , number of local iterations  $\{t_q\}_{q=1}^Q \in \mathbb{N}^Q$ .

- 1: Initialize  $[n]_q = \{\}$  for all  $q \in [Q]$
  - 2: Greedily select  $(\phi(c_q))_{q=1}^Q$  according to (7)
  - 3: **for**  $i = 1, \dots, n$  **do**
  - 4:   Compute  $\bar{q} = \operatorname{argmin}_{q \in [Q]} \|\phi(x_i) - \phi(c_q)\|_{\mathcal{H}}^2$
  - 5:   Update  $[n]_{\bar{q}} = [n]_{\bar{q}} \cup \{i\}$
  - 6: **end for**
  - 7: **for**  $q = 1, \dots, Q$  **do**
  - 8:   Select  $X_q = \{x_i \in X : i \in [n]_q\}$  and  $Y_q = \{y_i \in Y : i \in [n]_q\}$
  - 9:   Compute  $\tilde{f}_{q,t_q}$  as in eq. (13) using  $X_q, Y_q$
  - 10: **end for**
  - 11: Collect the local estimators  $\tilde{f}_{q,t_q}$  and return the ParK estimator  $\bar{f}_t$  as in eq. (14)
-

---

**Algorithm 2** ParK: Predict

---

**Require:** Test point  $x \in \mathbb{R}^d$ , local estimators  $\{\tilde{f}_{q,t}\}_{q=1}^Q$ , representatives of partition  $C = (c_q)_{q=1}^Q$

- 1: Select  $\bar{q} = \operatorname{argmin}_{q \in [Q]} \|\phi(x) - \phi(c_q)\|_{\mathcal{H}}^2$
  - 2: Evaluate  $\tilde{f}_{\bar{q},t}(x)$
- 

The time complexity of training ParK is  $O(Q^2 n \log(n))$  to compute the centroids,  $O(Qn)$  to compute the indices  $[n]_q$ , and  $O(t_q m_q n_q)$  to compute each local estimator. Putting these quantities together we get  $O(Q^2 n \log(n) + \sum_{q \in [Q]} t_q m_q n_q)$  in time, and  $O(\max_{q \in [Q]} m_q^2)$  in space. If we parallelize the training of the local estimators over  $Q$  machines, the time complexity further reduces to  $O(Q^2 n \log(n) + \max_{q \in [Q]} t_q m_q n_q)$ . In many practical scenarios, we can think  $Q$  as  $O(1)$ . For example, in all our experiments we take  $Q = 32$  (see Section 5). We compare the complexity of several KRR solver in Table 1.

Table 1: Computational complexity of some KRR solvers (up to constants). For D&C and ParK, we report the time complexity on  $Q$  parallel machines and the space requirement for each machine.

	naive	iterative [37]	Nyström/RF [36, 25]	FALKON [29]	D&C [38]	ParK
space	$n^2$	$n^2$	$m^2$	$m^2$	$(n/Q)^2$	$\max_q m_q^2$
time	$n^3$	$tn^2$	$m^2 n$	$tmn$	$(n/Q)^3$	$Q^2 n \log(n) + \max_q t_q m_q n_q$
test	$n$	$n$	$m$	$m$	$n$	$Q + \max_q m_q$

### Space partitioning vs data splitting

We conclude this section commenting on a different yet related distributed approach. As briefly recalled in the introduction, a straightforward way to decompose the KRR problem is by a simple split of the training data. For example, one can divide the samples uniformly at random into  $Q$  disjoint subsets of cardinality  $n_q = n/Q$ . Methods performing such a step are known as divide-and-conquer [38, 20, 14]. Consisting essentially in a block diagonal approximation of the kernel matrix, the resulting final estimator is an average of globally subsampled models. Divide-and-conquer methods are appealing due to the extreme simplicity of the splitting procedure and the direct control of the subsampling rates  $n_q$ . However, they can suffer from worse approximation error (see discussion in [34]), and be expensive at test and evaluation time. On the other hand, partitions present several potential benefits. First, data splitting is a byproduct of a geometric partition. This opens to the opportunity of exploiting the structure of the space, for instance enforcing notions of locality or orthogonality. Consistently, the final estimator is a union of local estimators, as opposed to an average of global ones. Hence, partitioning may enhance the approximation power of the model, capturing relevant local correlations [23]. As another consequence, at evaluation time only one local estimator, instead of the average of all estimators, needs to be called, yielding further computational saving. These nice properties have motivated a fruitful line of research, notably [23, 34, 24], where the advantage in the partitioning approach has been studied both in statistical and in computational terms. In this paper we concentrate on the computational aspects, expanding on theoretical tradeoffs outlined in [34, 24] and developing [24] with new algorithmic ideas.

## 4 Theory

To simplify the analysis and better highlight the new ideas in play, we consider the problem (1) in a fixed design setting [2, 15], where the  $x_i$  are deterministic and the  $\varepsilon_i$  are independent and identically distributed random variables.

Let  $L^2 = L^2(\rho)$  with  $\rho = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ . We may identify  $L^2$  with  $\mathbb{R}^n$  endowed with the inner product  $\langle u, w \rangle_{L^2} = \frac{1}{n} u^\top w$ . We define the excess risk of an estimate  $\hat{f}$  of  $f_*$  in problem (1) as

$$\mathcal{R}(\hat{f}) = \|\hat{f} - f_*\|_{L^2}^2 = \frac{1}{n} \sum_{i=1}^n |\hat{f}(x_i) - f_*(x_i)|^2. \quad (15)$$

We are interested in studying the performance of the estimator  $\bar{f}_t$  defined in (14) given a partition (4), as measured by the excess risk (15). Our theory will suggest how to construct the partition and tune the regularization in order to get the best learning rate.

#### 4.1 Definitions and assumptions

We start by defining some relevant operators in global and local variants. In view of (3) (and the fixed design setting), we assume without loss of generality that  $\mathcal{H} = \mathcal{H}_n$ . We define the covariance operator  $T : \mathcal{H} \rightarrow \mathcal{H}$  as  $T = \frac{1}{n} \sum_{i \in [n]} \phi(x_i) \otimes \phi(x_i)$ , where, for  $v, w \in \mathcal{H}$ ,  $v \otimes w$  denotes the operator  $u \in \mathcal{H} \mapsto \langle u, v \rangle_{\mathcal{H}} w \in \mathcal{H}$ . The operator  $T$  is standard in the analysis of kernel methods [7]. We now define the local version of the covariance operator conditioned on the partitioning (4). Thanks to (10) (and the fixed design setting), we can assume without loss of generality that  $\mathcal{H}_q = \text{span}\{\phi(x_i) : i \in [n]_q\}$ . The local covariance operator  $T_q : \mathcal{H} \rightarrow \mathcal{H}$  is defined as  $T_q = \frac{1}{n_q} \sum_{i \in [n]_q} \phi(x_i) \otimes \phi(x_i)$ . We denote with  $P_q : \mathcal{H} \rightarrow \mathcal{H}$  the orthogonal projection onto the subspace  $\mathcal{H}_q$ . For all  $q \in [Q]$ , we let  $\rho_q = n_q/n$ . Recall that we denote by  $\theta$  the minimal principal angle between the subspaces  $\mathcal{H}_q$ , as defined in (6).

To measure the capacity of the hypothesis spaces, we will use the standard notion of effective dimension [7].

**Effective dimension.** The (global) effective dimension of the space  $\mathcal{H}$  is given by

$$\mathcal{N}(\lambda) = \text{Tr}((T + \lambda)^{-1}T) \quad \lambda > 0.$$

Consistently, we define the local effective dimension of each space  $\mathcal{H}_q$  as

$$\mathcal{N}_q(\lambda_q) = \text{Tr}((T_q + \lambda_q)^{-1}T_q) \quad \lambda_q > 0.$$

We also define the local maximal degrees of freedom [2] as

$$\mathcal{N}_{\infty,q}(\lambda_q) = \sup_{x \in X_q} \langle \phi(x), (T_q + \lambda_q)^{-1} \phi(x) \rangle_{\mathcal{H}} \quad \lambda_q > 0,$$

which gives the bound  $\mathcal{N}_q(\lambda_q) \leq \mathcal{N}_{\infty,q}(\lambda_q) \leq \lambda_q^{-1} \sup_x K(x, x)$ . The effective dimension is related to the spectrum decay of the covariance operator, and thus it provides a way to quantify how many important eigenfunctions the RKHS contains. In this sense, it serves as an implicit number of parameters for the nonparametric model represented by the RKHS. The interplay between global and local effective dimensions, hence between global and local model complexity, will play a major role in our analysis. Similarly, there exist an interplay between a local and a global version of the maximal degrees of freedom  $\mathcal{N}_{\infty}(\lambda)$ , which is also connected to the coherence of the  $T$  operator, and to the concept of maximal leverage score [2].

We will need a few basic assumptions.

**Assumption 1.**  $f_* \in \mathcal{H}$ .

**Assumption 2.**  $\kappa^2 = \sup_{x \in \mathcal{X}} K(x, x) < \infty$ .

**Assumption 3.** The noise variables  $\varepsilon_i$  are i.i.d. sub-Gaussian of variance proxy  $\sigma^2 < \infty$ ,  $i \in [n]$ .

Assumption 3 is standard in the analysis of any regression model. In particular, sub-Gaussianity allows to control the tails of the noise, and therefore to establish bounds in high probability. Bounded and Gaussian noise are examples, but any variable with sub-Gaussian tail is covered. Assumptions 1 and 2 are instead typical of kernel methods. With Assumption 1, we suppose that the RKHS is a well specified model. We stick to Assumption 1 for simplicity, but we could easily relax it assuming the existence of a function in the RKHS with same excess risk as  $f_*$ , or considering the excess risk with respect to the best in class. Assumption 2 allows to provide explicit bound for kernel related quantities, and ensures in particular that functions in the RKHS are bounded.

#### 4.2 Main results

Our first proposition generalizes the classical bias-variance tradeoff of KRR estimators incorporating iterative optimization, random projections and feature partitioning. The result is a high probability bound for the excess risk of our ParK estimator.

**Proposition 1.** Let  $\delta \in (0, 1)$ . Under the regression model (1) and the assumptions of Section 4.1, let  $\bar{f}_{t_q}$  be the ParK estimator as defined in (14). If for each  $q \in [Q]$ ,  $0 < \lambda_q \leq \kappa^2$ ,

$$m_q \geq 5[1 + 14\mathcal{N}_{\infty,q}(\lambda_q)] \log\left(\frac{8\kappa^2}{\lambda_q\delta}\right), \quad t_q \geq 2 \log\left(4\sigma^2 \left(\|P_q f_*\|_{\mathcal{H}}^2 \lambda_q\right)^{-1/2}\right),$$

then, with probability at least  $1 - 4\delta$ ,

$$\mathcal{R}(\bar{f}_{t_q}) \leq 16 \sum_{q=1}^Q \|P_q f_*\|_{\mathcal{H}}^2 \lambda_q \rho_q + \sigma^2 \sum_{q=1}^Q \frac{\mathcal{N}_q(\lambda_q) + \sqrt{\mathcal{N}_q(\lambda_q) \log(1/\delta)} + 2 \log(1/\delta)}{n}.$$

The proof of Proposition 1 is given in Appendix A.2. The bound consists of a bias and a variance term. The bias term is an average of local biases, measured by the projection of the target function onto the local hypothesis spaces, regularized by a local penalization. The variance term is essentially the ratio between the sum of local effective dimensions and the global sample size. We are going to control bias and variance in the next two propositions, whose proof is postponed to Appendix A.3. For the bias, we prove the following generalized Bessel inequality.

**Proposition 2.** With the definitions of Section 4.1, we have

$$\sum_{q=1}^Q \|P_q f_*\|_{\mathcal{H}}^2 \leq (1 + Q^2 \cos(\theta)) \|f_*\|_{\mathcal{H}}^2.$$

Proposition 2 bounds the possible redundancy of the local projections by the minimal principal angle between the local subspaces. In particular, if the local subspaces are an orthogonal decomposition of the global space, the partitioned estimator has no additional local bias. On the other hand, lack of orthogonality results in a larger bias. Turning to the variance, we obtain the following bound on the local effective dimensions.

**Proposition 3.** With the definitions of Section 4.1, for  $\lambda_q = \lambda \rho_q^{-1}$  we have

$$\sum_{q=1}^Q \mathcal{N}_q(\lambda_q) \leq \left(1 + \kappa^2 \frac{\cos^2(\theta)}{\lambda}\right) \mathcal{N}(\lambda).$$

Once again, the minimal principal angle controls the ratio between local and global quantities. Where there is perfect orthogonality, splitting the hypothesis space does not increase the effective dimension; otherwise, a price proportional to the minimal principal angle is paid. With the above results in hand, we can now control the excess risk of the ParK estimator in terms of the global norm of the target function and the global effective dimension. This allows to compare the performance of our partitioned method to that of a typical global method.

**Theorem 4.** Let  $\delta \in (0, 1)$ . Under the same assumptions of Proposition 1, let  $\bar{q} = \operatorname{argmin}_{q \in [Q]} \rho_q$ , for  $0 < \lambda \leq \rho_{\bar{q}} \kappa^2$  when  $\lambda_q = \lambda \rho_q^{-1}$  for each  $q \in [Q]$ , with probability at least  $1 - 4\delta$ ,

$$\mathcal{R}(\bar{f}_t) \leq 16(1 + Q^2 \cos(\theta)) \|f_*\|_{\mathcal{H}}^2 \lambda + \frac{4\sigma^2}{n} \left(1 + \kappa^2 \frac{\cos^2(\theta)}{\lambda}\right) \mathcal{N}(\lambda) \log\left(\frac{1}{\delta}\right).$$

If we consider a model where  $\mathcal{H}$  is the orthogonal sum of the subspaces  $\mathcal{H}_q$ , as in [24], then  $\cos(\theta) = 0$ , and the bound of Theorem 4 simplifies to  $\mathcal{R}(\bar{f}_t) = O(\lambda \|f_*\|_{\mathcal{H}}^2 + \mathcal{N}(\lambda) n^{-1})$ . In particular, setting  $\lambda = O(1/\sqrt{n})$ , we obtain the learning rate  $O(1/\sqrt{n})$ . This is known to be the optimal rate, in the minimax sense, for global KRR models [7]. Note that, in the orthogonal case, the constraint  $\lambda \lesssim \rho_q$  translates to the minimal local point requirement  $n_q \gtrsim \sqrt{n}$  for all  $q$ , and hence to a bound on the partition size, namely  $Q \lesssim \sqrt{n}$ . On the other hand, when the subspaces  $\mathcal{H}_q$  are not perfectly orthogonal, our bound manifests a statistical-computational tradeoff, which is however quantified by the minimal principal angle. Further, the constraints on the local number of Nyström centers  $m_q$ , iterations  $t_q$  of Proposition 1, and the choice  $\lambda_q = \lambda \rho_q^{-1}$  with  $\lambda = 1/\sqrt{n}$  to achieve the minimax rate, allows to recover a time complexity of  $O\left(Q^2 n \log(n) + \sum_{q \in [Q]} n_q \sqrt{\frac{n_q}{\sqrt{n}}} \log\left(\frac{n_q}{\sqrt{n}}\right)\right)$ .

Analyses of (input space) partitioned kernel estimators have been conducted within different models, such as Gaussian SVM’s on Voronoi partitions [23], general kernels on clusters [34], and block-diagonal kernels on arbitrary partitions [24]. In these works, the bounds are established in random design, for plain [23, 34] or Nyström [24] local KRR estimators. Our result is in fixed design, but compared to [24] incorporates the additional algorithmic ingredient of iterative optimization. For a perfectly orthogonal model ( $\cos \theta = 0$ ), we recover the result in [24] as a special case (although in fixed design). In [34], the bias is controlled choosing same  $\lambda_q$  for all  $q$ , while the crucial bound of Proposition 3 is made as an assumption. Note however that, at least in our proof of Proposition 3, it is important to choose a differently scaled  $\lambda_q$  for each cell. Furthermore, our analysis and numerical tests motivate that partitioning the feature space is key to control both local bias and local effective dimension. Rather than on computational aspects, [23] focuses on extending statistical optimality for functions of local smoothness. This theme is also explored in [34, 24]. However, since the proposed partitioning step is either unsupervised [23, 34] or unspecified [24], improved rates can be obtained only under oracle assumptions, that is, assuming that the smoothness of the target function is localized right on the cells of the chosen partition. Partitions adapting to the unknown local smoothness of the target function can arguably be learned only in a supervised manner. This has been done for piecewise polynomial regression drawing ideas from multi-resolution analysis [4, 18]. An application of these ideas for kernel methods is not straightforward due to the usual computational constraints, but could be subject of future work.

## 5 Experiments

In this section we study the performance of ParK on some large-scale datasets ( $n \approx 10^6, 10^7, 10^9$ ). In particular we consider dataset where at the moment, because of their cardinality, only a few solvers can efficiently learn from. For this reason we compare to the global large-scale kernel method FALKON which has so far being the method that performs the best in terms of time and accuracy on these datasets [22]. A standard divide-and-conquer method can not run on these datasets (for the high space complexity), for this reason we compare with a version where each local estimator is a sketched KRR estimator computed with FALKON. We run two different versions of this algorithm, D&C-FALK(v1) and D&C-FALK(v2), that differ only in their hyper-parameters choices as specified later in this section. We also consider a second version of ParK where the centroids of the partition’s cells are chosen as  $\{\phi(c_q)\}_{q=1}^Q$  with  $c_q$  selected uniformly at randomly from the training data  $X$  (referred to as ParK-Uni). For each experiment we report mean and standard deviation on 10 trials. The experiments are implemented in python using pytorch and the FALKON library [22]. The experiments run on a machine with 2 Intel Xeon Silver 4116 CPUs and 1 GPU NVIDIA Titan Xp. The ram of the machine is 256 GB. We perform experiments on the four large-scale datasets TAXI ( $n \approx 10^9$ ,  $d = 9$ , regression), HIGGS ( $n \approx 10^7$ ,  $d = 28$ , classification), AIRLINE ( $n \approx 10^6$ ,  $d = 8$ , regression), AIRLINE-CLS ( $n \approx 10^6$ ,  $d = 8$ , classification) with the same pre-processing and same random train/test split used in [22]. We do not cross validate hyper-parameters of the local estimators of ParK. Instead we use the same used by FALKON in the paper [22] with the following exceptions: let  $\lambda$  be the global regularization parameters of FALKON and  $m$  the number of the Nyström points, the local estimators of ParK use regularization  $\lambda_q = \lambda \rho_q^{-1}$  and  $m_q = m \rho_q$  as suggested by the theory. D&C-FALK(v1) also follows the same rule for setting the hyper-parameters of its local estimators, while D&C-FALK(v2) uses the same of the (v1) version except the number of Nyström centers which are  $3m_q$  in AIRLINE and AIRLINE-CLS,  $5m_q$  in HIGGS, and  $6m_q$  in TAXI. The number of centroids used by ParK and D&C-FALK is  $Q = 32$  for all experiments. Performance for different  $Q$  values remains almost identical but worsen in time for higher values. Further, note that the local estimators of ParK and D&C-FALK are learned sequentially. We report in Table 2 the errors and times. In particular, for ParK(-Uni) we report the initialization time that include the greedy algorithm to select the centroids (not for ParK-Uni) and the assignment of the training points to the corresponding cell, the sequential training times of the local estimators, and the total time of this pipeline.

We can see that ParK can match the accuracy of the global FALKON estimator with a smaller computational cost. ParK-Uni requires further less time, at the expense of some loss in accuracy, confirming that a worse partition can affect generalization, as suggested by our theory. The reason of the ParK-Uni speedup is twofold. First, the initialization step requires only to assign points to a set of randomly selected centroids, and second, the local subsets of points have uniform cardinality (which is usually not the case for normal ParK). D&C-FALK(v1) is the algorithm with with smallest

Table 2: Accuracy and running time comparison on large-scale datasets.

	TAXI $n \approx 10^9$				HIGGS $n \approx 10^7$			
	ERROR (RMSE)	TIME (MIN.)			ERROR (1-AUC)	TIME (SEC.)		
		INIT	TRAIN	TOTAL		INIT	TRAIN	TOTAL
PARK	312.0±0.2	25±1	39±13	64±13	0.182±0.001	30±2	474±172	504±172
PARK-UNI	315.7±0.6	5±1	13±1	18±1	0.192±0.000	3±1	67±7	70±7
FALKON	311.7±0.1	-	-	120±1	0.180±0.001	-	-	715±6
D&C-FALK(v1)	356.2±0.2	-	-	14±1	0.212±0.000	-	-	50±1
D&C-FALK(v2)	327.4±0.1	-	-	29±1	0.195±0.000	-	-	288±2

  

	AIRLINE $n \approx 10^6$				AIRLINE-CLS $n \approx 10^6$			
	ERROR (MSE)	TIME (SEC.)			ERROR (C-ERR)	TIME (SEC.)		
		INIT	TRAIN	TOTAL		INIT	TRAIN	TOTAL
PARK	0.760±0.005	6±1	71±9	77±10	31.5±0.2 %	9±1	55±6	64±6
PARK-UNI	0.766±0.006	1±1	32±3	33±3	31.6±0.2 %	2±1	22±2	24±2
FALKON	0.758±0.005	-	-	334±2	31.5±0.2 %	-	-	391±5
D&C-FALK(v1)	0.834±0.005	-	-	27±1	33.2±0.1 %	-	-	20±1
D&C-FALK(v2)	0.799±0.005	-	-	96±1	32.2±0.1 %	-	-	73±1

training time but achieve significantly worse performance using the same rule to choose the number of Nyström points of ParK. For this reason, in D&C-FALK(v2) we increase the number of centroid to improve the performance, but the error of the method still results higher than the others with a training time now higher than ParK.

## 6 Conclusions and limitations.

In this paper we have proposed a new algorithm for large scale kernel ridge regression. Our method integrates and jointly exploits three previously uncombined algorithmic strategies, namely partitions, sketching and (preconditioned) iterative optimization. Distinctively from traditional partitioned methods, we have introduced the idea of partitioning the feature space, which allows to directly control and resolve the localization of the kernel model. We have presented a simple analysis that characterizes the statistical-computational trade-off of a partitioned kernel estimator by the interplay of intuitive quantities. Moreover, we have demonstrated that our algorithm performs favourably against a state-of-the-art large scale global method.

The main theoretical limitation of our work is the lack of a result connecting the proposed partitioning algorithm to the properties of the resulting partition. This seems to be a common gap in the literature of partitioned kernel methods, where partitions are often assumed to be given or, if explicitly constructed, are not statistically characterized. While our construction is theoretically motivated by the analysis and practically validated by the experiments, an actual guarantee is missing. In particular, one could try to prove that the proposed greedy procedure would actually find a maximally orthogonal decomposition of the hypothesis space, under suitable assumptions. From an algorithmic point of view, we point out that the computational cost of the greedy algorithm limits the choice of the partition size. Indeed, large partitions accelerate the training step, but increase the initialization time. We remark, however, that our model is flexible enough to include cheaper partitioning options. For example, our experiments show that uniformly chosen partitions can still produce good results.

## Acknowledgments and Disclosure of Funding

The authors thank Nicole Mücke for her useful feedback. This material is based upon work supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216, and the Italian Institute of Technology. We gratefully acknowledge the support of NVIDIA Corporation for the donation of the Titan Xp GPUs and the Tesla k40 GPU used for this research. L. R. acknowledges the financial support of the European Research Council (grant SLING 819789), the AFOSR projects FA9550-18-1-7009, FA9550-17-1-0390 and BAA-AFRL-AFOSR-2016-0007 (European Office of Aerospace Research and Development), and the EU H2020-MSCA-RISE project NoMADS - DLV-777826.

## References

- [1] H. Avron, V. Sindhwani, and D. Woodruff. “Sketching structured matrices for faster nonlinear regression”. In: *Advances in neural information processing systems*. 2013, pp. 2994–3002.
- [2] F. Bach. “Sharp analysis of low-rank kernel matrix approximations”. In: *Proceedings of the 26th Conference on Learning Theory (COLT)* PMLR 30 (2013), pp. 185–209.
- [3] F. Bauer, S. Pereverzev, and L. Rosasco. “On regularization algorithms in learning theory”. In: *Journal of Complexity* 23.1 (2007), pp. 52–72.
- [4] P. Binev, A. Cohen, W. Dahmen, and R. A. DeVore. “Universal Algorithms for Learning Theory Part II: Piecewise Polynomial Functions”. In: *Constructive Approximation* 26.2 (2007), pp. 127–152.
- [5] I. Blaschzyk and I. Steinwart. *Improved Classification Rates for Localized SVMs*. arXiv:1905.01502. 2019.
- [6] R. Camoriano, T. Angles, A. Rudi, and L. Rosasco. “Nytro: When subsampling meets early stopping”. In: *Artificial Intelligence and Statistics*. 2016, pp. 1403–1411.
- [7] A. Caponnetto and E. De Vito. “Optimal rates for the regularized least-squares algorithm”. In: *Foundations of Computational Mathematics* 7.3 (2007), pp. 331–368.
- [8] L. Carratino, A. Rudi, and L. Rosasco. “Learning with SGD and Random Features”. In: *Advances in Neural Information Processing Systems*. Ed. by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett. Vol. 31. Curran Associates, Inc., 2018.
- [9] L. Chen, G. Zhang, and E. Zhou. “Fast Greedy MAP Inference for Determinantal Point Process to Improve Recommendation Diversity”. In: *Advances in Neural Information Processing Systems*. Ed. by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett. Vol. 31. Curran Associates, Inc., 2018.
- [10] A. Dieuleveut and F. Bach. “Nonparametric stochastic approximation with large step-sizes”. In: *The Annals of Statistics* 44.4 (2016), pp. 1363–1399.
- [11] F. Dumpert and A. Christmann. “Universal consistency and robustness of localized support vector machines”. In: *Neurocomputing* 315 (2018), pp. 96–106.
- [12] E. Elhamifar and R. Vidal. “Sparse Subspace Clustering: Algorithm, Theory, and Applications”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35.11 (2013), pp. 2765–2781.
- [13] A. Gonen, F. Orabona, and S. Shalev-Shwartz. “Solving ridge regression using sketched preconditioned svrg”. In: *International Conference on Machine Learning*. 2016, pp. 1397–1405.
- [14] Z.-C. Guo, S.-B. Lin, and D.-X. Zhou. “Learning theory of distributed spectral algorithms”. In: *Inverse Problems* 33.7 (2017), p. 074009.
- [15] D. Hsu, S. M. Kakade, and T. Zhang. “Random Design Analysis of Ridge Regression”. In: *Proceedings of the 25th Annual Conference on Learning Theory (COLT)* PMLR 23 (2012), pp. 9.1–9.24.
- [16] D. Hsu, S. Kakade, T. Zhang, et al. “A tail inequality for quadratic forms of subgaussian random vectors”. In: *Electronic Communications in Probability* 17 (2012).
- [17] T. Klock, A. Lanteri, and S. Vigogna. “Estimating multi-index models with response-conditional least squares”. In: *Electronic Journal of Statistics* 15.1 (2021), pp. 589–629.
- [18] W. Liao, M. Maggioni, and S. Vigogna. “Multiscale regression on unknown manifolds”. In: *Mathematics in Engineering* 4.4 (2022), pp. 1–25.
- [19] J. Lin and V. Cevher. “Kernel Conjugate Gradient Methods with Random Projections”. In: *Applied and Computational Harmonic Analysis* (2021).
- [20] S.-B. Lin, X. Guo, and D.-X. Zhou. “Distributed Learning with Regularized Least Squares”. In: *Journal of Machine Learning Research* 18.92 (2017), pp. 1–31.
- [21] M. W. Mahoney and P. Drineas. “CUR matrix decompositions for improved data analysis”. In: vol. 106. 3. 2009, pp. 697–702.
- [22] G. Meanti, L. Carratino, L. Rosasco, and A. Rudi. “Kernel Methods Through the Roof: Handling Billions of Points Efficiently”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin. Vol. 33. Curran Associates, Inc., 2020, pp. 14410–14422.

- [23] M. Meister and I. Steinwart. “Optimal Learning Rates for Localized SVMs”. In: *Journal of Machine Learning Research* 17.194 (2016), pp. 1–44.
- [24] N. Mücke. “Reducing training time by efficient localized kernel regression”. In: *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*. Vol. 89. 2019, pp. 2603–2610.
- [25] A. Rahimi and B. Recht. “Random features for large-scale kernel machines”. In: *Advances in neural information processing systems*. 2008, pp. 1177–1184.
- [26] C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2006. ISBN: 026218253X.
- [27] D. Richards and P. Rebeschini. “Optimal Statistical Rates for Decentralised Non-Parametric Regression with Linear Speed-Up”. In: *Advances in Neural Information Processing Systems*. Vol. 32. 2019.
- [28] A. Rudi, R. Camoriano, and L. Rosasco. “Less is more: Nyström computational regularization”. In: *Advances in Neural Information Processing Systems*. 2015, pp. 1657–1665.
- [29] A. Rudi, L. Carratino, and L. Rosasco. “FALKON: An optimal large scale kernel method”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 3891–3901.
- [30] A. Rudi and L. Rosasco. “Generalization Properties of Learning with Random Features”. In: *Advances in Neural Information Processing Systems* 30. 2017, pp. 3215–3225.
- [31] B. Schölkopf and A. J. Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT Press, 2002.
- [32] I. Steinwart and A. Christmann. *Support vector machines*. Springer-Verlag New York, 2008.
- [33] I. Steinwart, D. R. Hush, C. Scovel, et al. “Optimal Rates for Regularized Least Squares Regression.” In: *COLT*. 2009.
- [34] R. Tandon, S. Si, P. Ravikumar, and I. Dhillon. *Kernel Ridge Regression via Partitioning*. arXiv:1608.01976. 2016.
- [35] P. Thomann, I. Blaschzyk, M. Meister, and I. Steinwart. “Spatial Decompositions for Large Scale SVMs”. In: *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. Vol. 54. Proceedings of Machine Learning Research. PMLR, 2017, pp. 1329–1337.
- [36] C. K. Williams and M. Seeger. “Using the Nyström method to speed up kernel machines”. In: *Advances in neural information processing systems*. 2001, pp. 682–688.
- [37] Y. Yao, L. Rosasco, and A. Caponnetto. “On Early Stopping in Gradient Descent Learning”. In: *Constructive Approximation* 26.2 (2007), pp. 289–315.
- [38] Y. Zhang, J. Duchi, and M. Wainwright. “Divide and Conquer Kernel Ridge Regression: A Distributed Algorithm with Minimax Optimal Rates”. In: *Journal of Machine Learning Research* 16.102 (2015), pp. 3299–3340.

## A Appendix

### A.1 Relevant operators

We define operators in global, local and subsampled variants. The global definitions are standard in the analysis of kernel methods [7]. In view of (3), we assume without loss of generality that  $\mathcal{H} = \mathcal{H}_n$ . Recall that  $L^2 = L^2(\rho)$  with  $\rho = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ , and that we identify  $L^2$  with  $\mathbb{R}^n$  with inner product  $\langle u, w \rangle_{L^2} = \frac{1}{n} u^\top w$ .

**Global operators:**

- $S : \mathcal{H} \rightarrow L^2$      $Sf(x) = \langle f, \phi(x) \rangle_{\mathcal{H}}$     the sampling operator
- $S^* : L^2 \rightarrow \mathcal{H}$      $S^*w = \frac{1}{n} \sum_{i \in [n]} w_i \phi(x_i)$     the out-of-sample extension operator
- $T : \mathcal{H} \rightarrow \mathcal{H}$      $T = S^*S = \frac{1}{n} \sum_{i \in [n]} \phi(x_i) \otimes \phi(x_i)$     the covariance operator

We now define local versions of the operators above, conditioned on the partitioning (4). Thanks to (10), we can assume without loss of generality that  $\mathcal{H}_q = \text{span}\{\phi(x_i) : i \in [n]_q\}$ . Let  $L_q^2 = L^2(\rho(\cdot | V_q))$  with  $\rho(\cdot | V_q) = \frac{1}{n_q} \sum_{i \in [n]_q} \delta_{x_i}$ . We identify  $L_q^2$  with  $\mathbb{R}^{n_q}$  endowed with the inner product  $\langle u, w \rangle_{L_q^2} = \frac{1}{n_q} u^\top w$ .

**Local operators:**

- $S_q : \mathcal{H} \rightarrow L_q^2$      $S_q f(x) = \langle f, \phi(x) \rangle_{\mathcal{H}}$
- $S_q^* : L_q^2 \rightarrow \mathcal{H}$      $S_q^* w = \frac{1}{n_q} \sum_{i \in [n]_q} w_i \phi(x_i)$
- $T_q : \mathcal{H} \rightarrow \mathcal{H}$      $T_q = S_q^* S_q = \frac{1}{n_q} \sum_{i \in [n]_q} \phi(x_i) \otimes \phi(x_i)$

The orthogonal projection  $P_q : \mathcal{H} \rightarrow \mathcal{H}$  onto the subspace  $\mathcal{H}_q$  is given by

$$P_q = S_q^+ S_q,$$

where  $^+$  denotes the Moore–Penrose pseudoinverse. Let  $\rho_q = \rho(V_q) = n_q/n$ . We observe that

$$T = \sum_q T_q \rho_q, \tag{16}$$

namely, the global covariance is an average of local covariances. Based on the local subsampling (11), we further introduce the following operators.

**Subsampled local operators:**

- $\tilde{S}_q : \mathcal{H} \rightarrow \mathbb{R}^{m_q}$      $\tilde{S}_q f = \frac{1}{\sqrt{m_q}} (\langle f, \phi(\tilde{x}_{q,i}) \rangle)_{i=1}^{m_q}$
- $\tilde{S}_q^* : \mathbb{R}^{m_q} \rightarrow \mathcal{H}$      $\tilde{S}_q^* w = \frac{1}{\sqrt{m_q}} \sum_{i=1}^{m_q} w_i \phi(\tilde{x}_{q,i})$

### A.2 Controlling the excess risk

In this section we prove Proposition 1. Both the Euclidean norm of vectors and the spectral norm of matrices are denoted by  $\|\cdot\|$ .

**From global to local excess risk.** Note that another way to write the excess risk (15) is

$$\mathcal{R}(\hat{f}) = \|T^{1/2}(\hat{f} - f_*)\|_{\mathcal{H}}^2.$$

Define now a local version of the above risk on the cells of the partition (4) as

$$\mathcal{R}_q(\hat{f}_q) = \|S_q(\hat{f}_q - f_*)\|_{L_q^2}^2. \tag{17}$$

**Lemma 1.** For every  $\bar{f}$  defined as in (9),

$$\mathcal{R}(\bar{f}) = \sum_{q=1}^Q \mathcal{R}_q(\hat{f}_q) \rho_q.$$

*Proof.* We have

$$\begin{aligned}
\mathcal{R}(\bar{f}) &= \|\bar{f} - f_*\|_{L^2}^2 \\
&= \sum_{q \in [Q]} \sum_{i \in [n]_q} |\bar{f}(x_i) - f_*(x_i)|^2 \\
&= \sum_{q \in [Q]} \sum_{i \in [n]_q} |\hat{f}_q(x_i) - f_*(x_i)|^2 \\
&= \sum_{q \in [Q]} \sum_{i \in [n]_q} |S_q \hat{f}_q(x_i) - S_q f_*(x_i)|^2 \\
&= \sum_{q \in [Q]} \|S_q(\hat{f}_q - f_*)\|_{L^2_q}^2. \quad \square
\end{aligned}$$

**Lemma 2.** *The local excess risk (17) can be rewritten as*

$$\mathcal{R}_q(\hat{f}_q) = \|T_q^{1/2}(\hat{f}_q - P_q f_*)\|_{\mathcal{H}}^2.$$

*Proof.* Since  $S_q = S_q S_q^+ S_q = S_q P_q$ , we have

$$\begin{aligned}
\|S_q(\hat{f}_q - f_*)\|_{L^2_q}^2 &= \|S_q(\hat{f}_q - P_q f_*)\|_{L^2_q}^2 \\
&= \langle S_q(\hat{f}_q - P_q f_*), S_q(\hat{f}_q - P_q f_*) \rangle_{L^2_q} \\
&= \langle T_q(\hat{f}_q - P_q f_*), (\hat{f}_q - P_q f_*) \rangle_{\mathcal{H}} \\
&= \|T_q^{1/2}(\hat{f}_q - P_q f_*)\|_{\mathcal{H}}^2. \quad \square
\end{aligned}$$

**From FALKON to Nyström local estimators.** We now control the local excess risk of each local estimator  $\tilde{f}_{q,t}$  as defined in (13) with the exact local Nyström estimator defined by

$$\tilde{f}_q(x) = \sum_{i=1}^{m_q} (\tilde{\alpha}_q)_i K(\tilde{x}_i, x), \quad \tilde{\alpha}_q = \operatorname{argmin}_{\alpha \in \mathbb{R}^{m_q}} \frac{1}{n_q} \|K_{n_q m_q} \alpha - Y_q\|^2 + \lambda_q \alpha^\top K_{m_q} \alpha. \quad (18)$$

Adapting the analysis of [29] to fixed design and local setting we derive the following lemma.

**Lemma 3.** *Let  $\delta \in (0, 1]$ , the Nyström centers in  $\tilde{f}_{q,t}$  be selected uniformly at random from  $X_q$ ,  $n_q, m_q, t \in \mathbb{N}$ . If  $0 \leq \lambda_q \leq \kappa^2$  and*

$$m_q \geq 5[1 + 14\mathcal{N}_{\infty, q}(\lambda_q)] \log\left(\frac{8\kappa^2}{\lambda_q \delta}\right), \quad (19)$$

*then, with probability  $1 - 2\delta$ ,*

$$\mathcal{R}_q(\tilde{f}_{q,t})^{1/2} \leq \mathcal{R}_q(\tilde{f}_q)^{1/2} + 6\sigma\kappa \|P_q f_*\|_{\mathcal{H}} \log\left(\frac{1}{\delta}\right) e^{-\frac{t}{2}}.$$

*Proof.* We follow the proof of Theorem 1 and Lemma 11 of [29], replacing the operators  $S, S^*, C$  in [29] with our local operators  $S_q, S_q^*, T_q$ . Note that in fixed design we do not have population operators, hence we can upper bound deterministically quantities that in random design require concentration arguments. Moreover, we upper bound the quantity  $\frac{\|Y_q\|}{\sqrt{n_q}}$  (our equivalent of  $\hat{\nu}$  in Theorem 1 of [29]) as follows. Recalling (1), we have

$$\frac{\|Y_q\|}{\sqrt{n_q}} = \frac{1}{\sqrt{n_q}} \sqrt{\sum_{i \in [n]_q} (f_*(x_i) + \varepsilon_i)^2} \leq \frac{1}{\sqrt{n_q}} \left( \sqrt{2 \sum_{i \in [n]_q} f_*(x_i)^2} + \sqrt{2 \sum_{i \in [n]_q} \varepsilon_i^2} \right). \quad (20)$$

Exploiting Assumption 1, for every  $x \in X_q$  we have  $f_*(x) = \langle f_*, K_x \rangle = \langle P_q f_*, K_x \rangle$ . Thus, Assumption 2 gives

$$\sup_{x \in X_q} |\langle P_q f_*, K_x \rangle_{\mathcal{H}}| \leq \sup_{x \in \mathcal{X}} \|P_q f_*\|_{\mathcal{H}} \|K_x\|_{\mathcal{H}} \leq \kappa \|P_q f_*\|_{\mathcal{H}}. \quad (21)$$

Let  $\widehat{\varepsilon} = [\varepsilon_i]_{i \in [n]_q} \in \mathbb{R}^{n_q}$ , Then, by Assumption 3, using Lemma 19 of [17] we obtain that, with probability at least  $1 - \delta$ ,

$$\|\widehat{\varepsilon}\| \leq \sigma \sqrt{n_q} \log\left(\frac{1}{\delta}\right). \quad (22)$$

Plugging (21) and (22) in (20) we conclude the proof.  $\square$

We now control the local excess risk of each local exact Nyström estimator  $\widetilde{f}_q$  as defined in (18). Adapting the analysis of [28] locally to fixed design we derive the following lemma.

**Lemma 4.** *Let  $\delta \in (0, 1]$ , the Nyström centers in  $\widetilde{f}_q$  be selected uniformly at random from  $X_q$ ,  $n_q, m_q, t \in \mathbb{N}$ . If  $0 \leq \lambda_q \leq \kappa^2$  and*

$$m_q \geq [2 + 3\mathcal{N}_{\infty, q}(\lambda_q)] \log\left(\frac{8\kappa^2}{\lambda_q \delta}\right),$$

then, with probability  $1 - 2\delta$ ,

$$\mathcal{R}_q(\widetilde{f}_q)^{1/2} \leq 3 \|P_q f_*\|_{\mathcal{H}} \sqrt{\lambda_q} + \frac{\sigma}{\sqrt{n_q}} \sqrt{\mathcal{N}_q(\lambda_q) + \sqrt{\mathcal{N}_q(\lambda_q) \log\left(\frac{1}{\delta}\right) + 2 \log\left(\frac{1}{\delta}\right)}}.$$

*Proof.* We follow the proof of Theorem 2 and Proposition 2 of [28], replacing the operators  $S, S^*, C$  in [28] with our local operators  $S_q, S_q^*, T_q$ . As for the proof of Lemma 3, concentration inequalities for empirical operators are replaced by deterministic bounds. Further we need to control the sample error in Lemma 4 of [28] with a different concentration argument. Let  $\widehat{Y}_q = \frac{\|Y_q\|}{\sqrt{n_q}}$ . The sample error in fixed design is

$$\left\| (T_q + \lambda_q)^{1/2} S_q^* (\widehat{Y}_q - S_q P_q f_*) \right\|_{\mathcal{H}}.$$

In view of (1) we have

$$\left\| (T_q + \lambda_q)^{1/2} S_q^* (\widehat{Y}_q - S_q P_q f_*) \right\|_{\mathcal{H}} = \frac{1}{\sqrt{n_q}} \left\| (T_q + \lambda_q)^{1/2} S_q^* (\widehat{\varepsilon}) \right\|_{\mathcal{H}},$$

where  $\widehat{\varepsilon} = [\varepsilon_i]_{i \in [n]_q} \in \mathbb{R}^{n_q}$ . Now using Assumption 3, Remark 2.2 of [16] and the definition of local effective dimension, we obtain, with probability at least  $1 - \delta$ ,

$$\frac{1}{\sqrt{n_q}} \left\| (T_q + \lambda_q)^{1/2} S_q^* (\widehat{\varepsilon}) \right\|_{\mathcal{H}} \leq \frac{\sigma}{\sqrt{n_q}} \sqrt{\mathcal{N}_q(\lambda_q) + \sqrt{\mathcal{N}_q(\lambda_q) \log\left(\frac{1}{\delta}\right) + 2 \log\left(\frac{1}{\delta}\right)}},$$

which concludes the proof.  $\square$

We are now ready to prove Proposition 1.

**Proof of Proposition 1** From Lemmas 3 and 4 we know that, under their respective assumptions and for a value of  $m_q$  as in (19), with probability  $1 - 4\delta$ ,

$$\begin{aligned} \mathcal{R}_q(\widetilde{f}_{q,t})^{1/2} &\leq 3 \|P_q f_*\|_{\mathcal{H}} \sqrt{\lambda_q} + \frac{\sigma}{\sqrt{n_q}} \sqrt{\mathcal{N}_q(\lambda_q) + \sqrt{\mathcal{N}_q(\lambda_q) \log\left(\frac{1}{\delta}\right) + 2 \log\left(\frac{1}{\delta}\right)}} \\ &\quad + 6\sigma\kappa \|P_q f_*\|_{\mathcal{H}} \log\left(\frac{1}{\delta}\right) e^{-\frac{t}{2}}. \end{aligned}$$

We consider now a number of iterations  $t$  such that  $6\sigma\kappa \|P_q f_*\|_{\mathcal{H}} \log\left(\frac{1}{\delta}\right) e^{-\frac{t}{2}} \leq \|P_q f_*\|_{\mathcal{H}} \sqrt{\lambda_q}$ , that is

$$t \geq 2 \log\left(\frac{6\sigma\kappa \log(1/\delta)}{\sqrt{\lambda_q}}\right).$$

Under the above constraint on  $t$  we can rewrite the upper bound on the risk

$$\mathcal{R}_q(\widetilde{f}_{q,t})^{1/2} \leq 4 \|P_q f_*\|_{\mathcal{H}} \sqrt{\lambda_q} + \frac{\sigma}{\sqrt{n_q}} \sqrt{\mathcal{N}_q(\lambda_q) + \sqrt{\mathcal{N}_q(\lambda_q) \log\left(\frac{1}{\delta}\right) + 2 \log\left(\frac{1}{\delta}\right)}}.$$

We can now collect the local excess risk bounds above for all  $q \in [Q]$  using Lemmas 1 and 2, concluding the proof.  $\square$

### A.3 Controlling the partition

In this section we prove Propositions 2 and 3. With a slight abuse of notation, the operator norm on  $\mathcal{H}$  is denoted by  $\|\cdot\|_{\mathcal{H}}$ .

*Proof of Proposition 2.* We have

$$\sum_q \|P_q f_*\|_{\mathcal{H}}^2 = \sum_q \langle P_q f_*, P_q f_* \rangle_{\mathcal{H}} = \sum_q \langle f_*, P_q f_* \rangle_{\mathcal{H}} = \langle f_*, \sum_q P_q f_* \rangle_{\mathcal{H}} \leq \left\| \sum_q P_q \right\|_{\mathcal{H}} \|f_*\|_{\mathcal{H}}^2.$$

Now, let  $U_q : \mathcal{H} \rightarrow \mathbb{R}^{n_q}$  such that  $U_q^* U_q = P_q$ ,  $U_q U_q^* = I_{n_q}$ , and define

$$U = [U_1, \dots, U_Q]^{\top} : \mathcal{H} \rightarrow \mathbb{R}^n.$$

Then  $\sum_q P_q = U^* U$ , and

$$\left\| \sum_q P_q \right\|_{\mathcal{H}} = \|U^* U\|_{\mathcal{H}} = \|U U^*\|,$$

Let  $W = U U^*$ . Then  $W \in \mathbb{R}^{n \times n}$  is built as

$$W = \begin{bmatrix} U_1 U_1^* & \cdots & U_1 U_Q^* \\ \vdots & \ddots & \vdots \\ U_Q U_1^* & \cdots & U_Q U_Q^* \end{bmatrix} = \begin{bmatrix} I_{n_1} & \cdots & W_{1,Q} \\ \vdots & \ddots & \vdots \\ W_{Q,1} & \cdots & I_{n_Q} \end{bmatrix}.$$

Thus, for  $a = [a_1, \dots, a_Q] \in \mathbb{R}^n$ ,  $a_q \in \mathbb{R}^{n_q}$ , we have

$$\begin{aligned} \|U U^*\| &= \|W\| \\ &= \|I + (W - I)\| \\ &= 1 + \lambda_{\max}(W - I) \\ &= 1 + \max_{\|a\|=1} a^{\top} (W - I) a \\ &= 1 + \max_{\|a\|=1} \sum_q a_q^{\top} (W_{q,q} - I_{n_q}) a_q + \sum_{q,k:q \neq k} a_q^{\top} (W_{q,k} - 0) a_k \\ &= 1 + \max_{\|a\|=1} \sum_q a_q^{\top} 0 a_q + \sum_{q,k:q \neq k} a_q^{\top} W_{q,k} a_k \\ &= 1 + \max_{\|a\|=1} \sum_{q,k:q \neq k} a_q^{\top} W_{q,k} a_k \\ &\leq 1 + \sum_{q,k:q \neq k} \max_{\|a\|=1} a_q^{\top} W_{q,k} a_k. \end{aligned}$$

Now we can bound

$$\sum_{q,k:q \neq k} \max_{\|a\|=1} a_q^{\top} W_{q,k} a_k \leq Q^2 \max_{q,k:q \neq k} \max_{\|a\|=1} a_q^{\top} W_{q,k} a_k,$$

and reparameterizing  $a_q = \beta_q b_q$ ,  $\beta_q \geq 0$ ,  $b_q \in \mathbb{R}^{n_q}$ , we get

$$\begin{aligned} \max_{\|a\|=1} a_q^{\top} W_{q,k} a_k &= \max_{\substack{\|b_1\|=\dots=\|b_Q\|=1 \\ \beta_1^2+\dots+\beta_Q^2=1}} \beta_q \beta_k b_q^{\top} W_{q,k} b_k \\ &= \max_{\beta_1^2+\dots+\beta_Q^2=1} \beta_q \beta_k \max_{\|b_1\|=\dots=\|b_Q\|=1} b_q^{\top} W_{q,k} b_k \\ &= \max_{\beta_1^2+\dots+\beta_Q^2=1} \beta_q \beta_k \max_{\|b_q\|=\|b_k\|=1} b_q^{\top} W_{q,k} b_k \\ &= \max_{\beta_1^2+\dots+\beta_Q^2=1} \beta_q \beta_k \cos(\angle(\mathcal{H}_q, \mathcal{H}_k)) \\ &\leq \cos(\theta). \end{aligned}$$

Putting all together, we finally obtain

$$\left\| \sum_q P_q \right\|_{\mathcal{H}} \leq 1 + Q^2 \cos(\theta),$$

which completes the proof.  $\square$

*Proof of Proposition 3.* Let  $\tilde{T}_q = P_q T P_q \rho_q^{-1}$ , and let  $M_q = (\tilde{T}_q + \lambda_q)^{1/2} (T_q + \lambda_q)^{-1} (\tilde{T}_q + \lambda_q)^{1/2}$ . Then

$$\begin{aligned}
\sum_q \mathcal{N}_q(\lambda_q) &= \sum_q \text{Tr}(S_q (T_q + \lambda_q)^{-1} S_q^*) \\
&= \sum_q \text{Tr}(S_q (\tilde{T}_q + \lambda_q)^{-1/2} M_q (\tilde{T}_q + \lambda_q)^{-1/2} S_q^*) \\
&= \sum_q \text{Tr}(M_q (\tilde{T}_q + \lambda_q)^{-1/2} S_q^* S_q (\tilde{T}_q + \lambda_q)^{-1/2}) \\
&\leq \sup_q \|M_q\|_{\mathcal{H}} \sum_q \text{Tr}((\tilde{T}_q + \lambda_q)^{-1/2} S_q^* S_q (\tilde{T}_q + \lambda_q)^{-1/2}) \\
&= \sup_q \|M_q\|_{\mathcal{H}} \sum_q \text{Tr}(S_q (\tilde{T}_q + \lambda_q)^{-1} S_q^*),
\end{aligned}$$

where in the third and last equalities we used the cyclic property of the trace, and in the fourth step we applied Holder's inequality. We first bound the trace. We have

$$\begin{aligned}
(\tilde{T}_q + \lambda_q)^{-1} &= \lambda_q^{-1} (\tilde{T}_q + \lambda_q - \tilde{T}_q) (\tilde{T}_q + \lambda_q)^{-1} \\
&= \lambda_q^{-1} (I - \tilde{T}_q (\tilde{T}_q + \lambda_q)^{-1}) \\
&= \lambda_q^{-1} (I - P_q S^* S P_q \rho_q^{-1} (P_q S^* S P_q \rho_q^{-1} + \lambda_q)^{-1}) \\
&= \lambda_q^{-1} (I - P_q S^* S P_q (P_q S^* S P_q + \lambda_q \rho_q)^{-1}) \\
&= \lambda_q^{-1} (I - P_q S^* (S P_q S^* + \lambda_q \rho_q)^{-1} S P_q) \\
&\preceq \lambda_q^{-1} (I - P_q S^* (S S^* + \lambda_q \rho_q)^{-1} S P_q).
\end{aligned}$$

where the fifth equality follows from the Woodbury identity. Thus, multiplying by  $S_q$  from the left and by  $S_q^*$  from the right, we get

$$\begin{aligned}
S_q (\tilde{T}_q + \lambda_q)^{-1} S_q^* &\preceq \lambda_q^{-1} (S_q S_q^* - S_q P_q S^* (S S^* + \lambda_q \rho_q)^{-1} S P_q S_q^*) \\
&= \lambda_q^{-1} (S_q S_q^* - S_q S^* (S S^* + \lambda_q \rho_q)^{-1} S S_q^*) \\
&= \lambda_q^{-1} (S_q (I - S^* (S S^* + \lambda_q \rho_q)^{-1} S) S_q^*) \\
&= \lambda_q^{-1} (S_q (I - (T + \lambda_q \rho_q)^{-1} T) S_q^*) \\
&= \lambda_q^{-1} (S_q (T + \lambda_q \rho_q)^{-1} (T + \lambda_q \rho_q - T) S_q^*) \\
&= S_q (T + \lambda_q \rho_q)^{-1} S_q^* \rho_q,
\end{aligned}$$

where again the fourth equality follows from the Woodbury identity. Therefore,

$$\text{Tr}(S_q (\tilde{T}_q + \lambda_q)^{-1} S_q^*) \leq \text{Tr}(S_q (T + \lambda_q \rho_q)^{-1} S_q^* \rho_q) = \text{Tr}((T + \lambda_q \rho_q)^{-1} T_q \rho_q).$$

Setting  $\lambda_q = \lambda \rho_q^{-1}$  and using (16) we obtain

$$\sum_q \text{Tr}(S_q (\tilde{T}_q + \lambda_q)^{-1} S_q^*) \leq \text{Tr}((T + \lambda)^{-1} \sum_q T_q \rho_q) = \text{Tr}((T + \lambda)^{-1} T) = \mathcal{N}(\lambda).$$

We next bound  $\|M_q\|_{\mathcal{H}}$ . The operators  $T_q + \lambda_q$  and  $\tilde{T}_q + \lambda_q$  are invertible, hence  $M_q$  shares the same spectrum as  $(T_q + \lambda_q)^{-1/2} (\tilde{T}_q + \lambda_q) (T_q + \lambda_q)^{-1/2}$ , and in particular

$$\|M_q\|_{\mathcal{H}} = \|(T_q + \lambda_q)^{-1/2} (\tilde{T}_q + \lambda_q) (T_q + \lambda_q)^{-1/2}\|_{\mathcal{H}}.$$

Now, let

$$\bar{T}_q : \mathcal{H} \rightarrow \mathcal{H} \quad \bar{T}_q = \frac{1}{n} \sum_{\phi(x_i) \notin V_q} \phi(x_i) \otimes \phi(x_i).$$

Then  $T = T_q \rho_q + \bar{T}_q$ , and

$$\tilde{T}_q = P_q T_q P_q + P_q \bar{T}_q P_q \rho_q^{-1} = T_q + P_q \bar{T}_q P_q \rho_q^{-1}.$$

Therefore,

$$\begin{aligned}
\|M_q\|_{\mathcal{H}} &= \|(T_q + \lambda_q)^{-1/2}(T_q + \lambda_q + P_q \bar{T}_q P_q \rho_q^{-1})(T_q + \lambda_q)^{-1/2}\|_{\mathcal{H}} \\
&= \|I + \rho_q^{-1}(T_q + \lambda_q)^{-1/2}(P_q \bar{T}_q P_q)(T_q + \lambda_q)^{-1/2}\|_{\mathcal{H}} \\
&\leq 1 + \rho_q^{-1} \|T_q + \lambda_q\|_{\mathcal{H}}^{-1/2} \|P_q \bar{T}_q P_q\|_{\mathcal{H}} \|T_q + \lambda_q\|_{\mathcal{H}}^{-1/2} \\
&\leq 1 + \rho_q^{-1} \lambda_q^{-1/2} \|P_q \bar{T}_q P_q\|_{\mathcal{H}} \lambda_q^{-1/2} \\
&= 1 + \frac{1}{\lambda_q \rho_q} \|P_q \bar{T}_q P_q\|_{\mathcal{H}}.
\end{aligned}$$

For  $\lambda_q = \lambda \rho_q^{-1}$ , we get

$$\|M_q\|_{\mathcal{H}} \leq 1 + \frac{1}{\lambda} \|P_q \bar{T}_q P_q\|_{\mathcal{H}}.$$

Finally,

$$\begin{aligned}
\|P_q \bar{T}_q P_q\|_{\mathcal{H}} &\leq \frac{1}{n} \sum_{\phi(x_i) \notin V_q} \|P_q \phi(x_i)\|_{\mathcal{H}}^2 \\
&= \frac{1}{n} \sum_{\phi(x_i) \notin V_q} \|\phi(x_i)\|_{\mathcal{H}}^2 \cos^2(\angle(\phi(x_i), \mathcal{H}_q)) \\
&\leq \frac{1}{n} \sum_{\phi(x_i) \notin V_q} \sup_i \|\phi(x_i)\|_{\mathcal{H}}^2 \cos^2(\min_{k \neq q} \angle(V_q, \mathcal{H}_k)) \\
&= \sup_i \|\phi(x_i)\|_{\mathcal{H}}^2 \cos^2(\min_{k \neq q} \angle(V_q, \mathcal{H}_k)) \frac{1}{n} (n - n_q) \\
&\leq \sup_i \|\phi(x_i)\|_{\mathcal{H}}^2 \cos^2(\min_{k \neq q} \angle(V_q, \mathcal{H}_k)),
\end{aligned}$$

which leads to the desired bound.  $\square$