

ANNALES
HENRI LEBESGUE

STEFFEN GRÜNEWÄLDER

COMPRESSED EMPIRICAL MEASURES (IN FINITE DIMENSIONS)

MESURES EMPIRIQUES COMPRESSÉES (EN
DIMENSIONS FINIES)

ABSTRACT. — We study approaches for compressing the empirical measure in the context of finite dimensional reproducing kernel Hilbert spaces (RKHSs). In this context, the empirical measure is contained within a natural convex set and can be approximated using convex optimization methods. Such an approximation gives rise to a coresets of data points. A key quantity that controls how large such a coresets has to be is the size of the largest ball around the empirical measure that is contained within the empirical convex set. The bulk of our work is concerned with deriving high probability lower bounds on the size of such a ball under various conditions and in various settings: we show how conditions on the density of the data and the kernel function can be used to infer such lower bounds; we further develop an approach that uses a lower bound on the smallest eigenvalue of a covariance operator to provide lower bounds on the size of such a ball; we extend the approach to approximate covariance operators and we show how it can be used in the context of kernel ridge regression. We also derive compression guarantees when standard algorithms like the conditional gradient method are used and we discuss variations of such algorithms to improve the runtime of these standard algorithms. We conclude with a construction of an infinite dimensional RKHS for which the compression is poor, highlighting some of the difficulties one faces when trying to move to infinite dimensional RKHSs.

Keywords: Kernel methods, coresets, compressed empirical measure.

2020 Mathematics Subject Classification: 68Q32, 62G05, 46E22.

DOI: <https://doi.org/10.5802/ahl.271>

RÉSUMÉ. — Nous étudions des méthodes de compression de la mesure empirique dans le cadre des espaces de Hilbert à noyau reproduisant de dimension finie (RKHS). Dans ce contexte, la mesure empirique est contenue dans un ensemble convexe naturel et peut être approximée à l'aide de méthodes d'optimisation convexe. Une telle approximation donne lieu à un petit ensemble de points de données, à savoir un « coreset ». Une quantité clé qui contrôle la taille d'un tel coreset est la taille de la plus grande boule autour de la mesure empirique qui est contenue dans l'ensemble convexe empirique. L'essentiel de notre travail consiste à dériver des bornes inférieures à forte probabilité sur la taille d'une telle boule dans diverses conditions et dans divers contextes : nous démontrons comment les conditions sur la densité des données et sur le noyau peuvent être utilisées afin de déduire de telles bornes inférieures. Nous développons en outre une approche qui utilise une borne inférieure sur la plus petite valeur propre d'un opérateur de covariance pour fournir des bornes inférieures sur la taille d'une telle boule ; nous étendons cette approche à l'approximation des opérateurs de covariance et démontrons comment elle peut être utilisée dans le contexte de la régression à noyau. Nous dérivons également des garanties de compression lorsque des algorithmes standards comme la méthode du gradient conditionnel sont utilisés et nous discutons des variantes de ces algorithmes pour améliorer le temps d'exécution de ces algorithmes standards. Nous concluons avec une construction d'un RKHS de dimension infinie pour lequel la compression est médiocre, soulignant certaines des difficultés auxquelles on est confronté lorsqu'on essaie de passer à des RKHS de dimension infinie.

CONTENTS

| | |
|---|-----|
| 1. Introduction | 473 |
| 1.1. Lower bounding the width of C | 476 |
| 1.2. Locating \mathfrak{m} within C | 478 |
| 1.3. Convergence of C_n to C | 478 |
| 1.4. Assumptions on the spectrum of the covariance operator | 479 |
| 1.5. Adapting the approach to concrete statistical problems | 480 |
| 1.5.1. Simultaneous approximation | 481 |
| 1.6. Implications for algorithms | 482 |
| 1.7. Slow rate of convergence in infinite dimensions | 483 |
| 1.8. Literature | 485 |
| 1.9. Preliminaries | 488 |
| 1.9.1. Separable processes and Rademacher complexities | 489 |
| 1.9.2. Bochner integrals and $\mathcal{L}^p(\mu, \mathcal{H})$ | 489 |
| 1.9.3. Tensor products | 490 |
| 1.9.4. Covariance operators | 491 |
| Acknowledgments | 492 |
| 2. Approximating convex sets and locating \mathfrak{m} and \mathfrak{m}_n | 492 |
| 2.1. Approximation based on ε -nets | 493 |
| 2.2. Empirical convex sets | 494 |
| 2.2.1. VC-theory | 495 |
| 2.2.2. Rademacher complexities | 496 |
| 2.2.3. Examples | 498 |
| 2.3. Width of the convex set C | 503 |
| 2.3.1. Outline | 504 |
| 2.3.2. Some comments regarding infinite dimensional \mathcal{H} | 505 |
| 2.3.3. Interpolation spaces | 505 |
| 2.3.4. Adding and removing constants | 510 |
| 2.3.5. Lower bounds on the approximation error in finite dimensions | 511 |
| 2.3.6. Quantifying the width of the empirical convex set C_n | 516 |

| | |
|--|-----|
| 2.4. Locating \mathbf{m} | 519 |
| 2.4.1. Assumptions on the density | 520 |
| 2.4.2. Assumptions on the covariance operator | 522 |
| 2.4.3. Data attaining values in a subset | 525 |
| 2.5. Locating \mathbf{m}_n within the empirical convex set | 528 |
| 2.6. Rescaling the kernel function does not affect compression | 532 |
| 3. Related approximation problems | 533 |
| 3.1. Assumptions | 535 |
| 3.2. Covariance operators | 535 |
| 3.2.1. Lower bound on the width of C_{\circlearrowleft} | 537 |
| 3.2.2. Lower bounds on the fourth moments | 538 |
| 3.3. Weighted mean embedding | 539 |
| 3.3.1. First setting: bounded Y_i 's & $Y_i f(X_i) = \langle f, Y_i k(X_i, \cdot) \rangle$ | 540 |
| 3.3.2. Second setting: bounded Y_i 's & Eq. (3.3) | 548 |
| 3.3.3. Third setting: unbounded Y_i 's & Eq. (3.3) | 550 |
| 3.4. Simultaneous compression | 556 |
| 3.4.1. Compressing the covariance and weighted mean embedding simultaneously | 557 |
| 3.4.2. Linearly dependent spaces | 561 |
| 3.4.3. Simultaneous least-squares risk approximation for unbounded Y | 563 |
| 4. Compression using the CGM and related approaches | 568 |
| 4.1. Kernel herding and subset selection | 569 |
| 4.1.1. Avoiding the explicit computation of s | 572 |
| 4.2. Better compression with the CGM | 575 |
| 4.2.1. Compression for kernel regression | 577 |
| 4.2.2. Avoiding the explicit computation of s | 579 |
| 5. Applications | 584 |
| 5.1. Two Sample Test | 584 |
| 5.2. Kernel ridge regression | 585 |
| 5.3. Kernel PCA | 587 |
| 6. Example: Slow rate of convergence in infinite dimensions | 589 |
| The construction idea | 589 |
| Interlacing | 590 |
| List of Notation | 602 |
| References | 603 |

1. Introduction

Many methods in machine learning and statistics make use of the empirical measure which is effectively a representation of the data. Reducing the number of points on which the empirical measure is supported, while preserving most of the information that is necessary for inference, can result in a significant speed-up of algorithms without sacrificing accuracy. We study the question of how to compress the empirical measure while preserving information in the context of *finite dimensional Reproducing Kernel Hilbert Spaces (RKHSs)*. To give an overview of our results it is useful to introduce the key objects of our investigation. We are generally concerned with data taking values in some set \mathcal{X} . Often we will assume this set to be compact. We then look at a kernel function k defined on \mathcal{X} and the corresponding RKHS \mathcal{H} . For various results, it is useful to assume that the functions in \mathcal{H} are continuous or even Lipschitz-continuous. Our main interest lies in the unknown distribution

P of data X_1, \dots, X_n where we assume throughout that X_1, \dots, X_n are independent and identically distributed. We adopt a common convention from the empirical process theory literature and will denote by Pf the integral $\int f(x) dP(x)$ whenever $f \in \mathcal{L}^1(\mathcal{X}, P)$. Since P is unknown it is common to use the empirical measure P_n as a surrogate, where $P_n f = (1/n) \sum_{i=1}^n f(X_i)$. There is a useful interplay between the measures P and P_n and RKHSs. Whenever $k(X_1, \cdot)$ is Bochner-integrable with respect to P we can define $\mathbf{m} = \int k(x, \cdot) dP(x) \in \mathcal{H}$ and it follows that

$$\langle \mathbf{m}, h \rangle = Ph, \quad \text{for all } h \in \mathcal{H}.$$

Similarly, by defining $\mathbf{m}_n = (1/n) \sum_{i=1}^n k(X_i, \cdot)$ we have that $\langle \mathbf{m}_n, h \rangle = P_n h$ for all $h \in \mathcal{H}$. Our aim in this paper is to find an element $\bar{\mathbf{m}}_n$ such that

$$\|\bar{\mathbf{m}}_n - \mathbf{m}_n\| \approx \|\mathbf{m}_n - \mathbf{m}\|$$

to guarantee that $\|\bar{\mathbf{m}}_n - \mathbf{m}\|$ is of the same order as $\|\mathbf{m}_n - \mathbf{m}\|$ and $\bar{\mathbf{m}}_n$ can be used in place of \mathbf{m}_n without sacrificing significant accuracy in applications.

To gain such an approximation $\bar{\mathbf{m}}_n$, we make use of another fortunate circumstance. The element \mathbf{m} does not only lie in \mathcal{H} but within the convex set

$$C = \text{cch}\{k(x, \cdot) : x \in \mathcal{X}\},$$

where cch denotes the closed convex hull. This is useful because the extremes of C are contained within the set $\{k(x, \cdot) : x \in \mathcal{X}\}$ and often we can reduce the study of C to studying the interaction between $k(x, \cdot)$ and functions $h \in \mathcal{H}$. For instance, the width of C in a direction $h \in \mathcal{H}$, $\|h\| = 1$, is

$$\text{width}_h(C) = \sup_{x \in \mathcal{X}} \langle k(x, \cdot), h \rangle - \inf_{x \in \mathcal{X}} \langle k(x, \cdot), h \rangle = \sup_{x \in \mathcal{X}} h(x) - \inf_{x \in \mathcal{X}} h(x).$$

The set $\{k(x, \cdot) : x \in \mathcal{X}\}$ is usually infinite and not directly useful for algorithms. However, when using \mathbf{m}_n , we have another convex set in \mathcal{H} that is usable, that is the empirical convex set $C_n = \text{ch}\{k(X_i, \cdot) : i \leq n\}$ which contains \mathbf{m}_n . The extremes of C_n are contained within the finite set $\{k(X_i, \cdot) : i \leq n\}$.

Standard techniques like the conditional gradient method or the kernel herding algorithm are directly applicable to approximate \mathbf{m}_n by convex combinations of $\{k(X_i, \cdot) : i \leq n\}$. The kernel herding algorithm generates an approximation of the form $(1/l) \sum_{i=1}^l k(X_{\iota_i}, \cdot)$, where $\iota : \{1, \dots, l\} \rightarrow \{1, \dots, n\}$ is some selection of data points and $l \leq n$. The data points $X_{\iota(1)}, \dots, X_{\iota(l)}$ themselves can be seen as a coresets for the data set. This approach is visualized in Figure 1.1 (i). The conditional gradient method does not provide such an average but an arbitrary convex combination of the points $k(X_1, \cdot), \dots, k(X_n, \cdot)$ and cannot be used directly to find a coresets. However, a coresets is often not necessary and many algorithms can work directly with an approximation of \mathbf{m} or related quantities; we demonstrate this in Section 1.5 and Section 5. The advantage of the conditional gradient method when compared to the kernel herding algorithm is that it usually leads to a vastly superior compression of the data.

Crucially, the performance of these techniques depends on the size of the largest ball in C_n that can be centered at \mathbf{m}_n . The existence of such a ball is in itself already of major importance for the performance of the techniques and is known as *Slater's condition*. In this paper, our main focus lies on the derivation of high probability




Figure 1.1. (i) The figure depicts how a subset or coresets of the sample is selected: the data is embedded in \mathcal{H} by using the kernel function of \mathcal{H} . An approximation algorithm is then applied to the convex polytope in \mathcal{H} to find an approximation of \mathbf{m} that uses only a few extremes of the convex polytope. The pre-images of these extremes are the sample points that are selected as the coresets. (ii) For most statistical problems approximating \mathbf{m} itself is insufficient and one has to approximate closely related quantities. In the case of least-squares regression, one has to approximate the operator $\mathfrak{C}_{y,n} \in \widehat{\mathcal{H} \odot \mathcal{H}}$ (see Sections 1.9 and 3 for the definitions), which is closely related to the empirical covariance operator, and a “weighted” mean embedding $\mathbf{m}_{y,n} \in \mathbb{R}' \otimes \mathcal{H}$. It is often of interest to approximate $\mathfrak{C}_{y,n}$ and $\mathbf{m}_{y,n}$ simultaneously, for instance, when building a coresets for least-squares regression. This can be achieved by considering the direct sum $(\widehat{\mathcal{H} \odot \mathcal{H}}) \oplus (\mathbb{R}' \otimes \mathcal{H})$ and a “direct sum” of the convex polytopes in the two spaces. The relation between the extremes of the convex polytopes is highlighted in the figure through the dotted lines; i.e. an algorithm will select a pair that is connected by a dotted line and by selecting such a pair of extremes the approximation of both the covariance and mean element will change.

lower bounds on the size of such a ball around \mathbf{m}_n within C_n . Figure 1.2 outlines our approach. In (i) the setting is shown with $\mathbf{m}_n \in C_n \subset C$ and the largest ball around \mathbf{m}_n in C_n is drawn. One of the main difficulties is that both \mathbf{m}_n and C_n are stochastic and change with the sample. We sidestep this difficulty by analyzing C and \mathbf{m} , and relating the empirical quantities C_n and \mathbf{m}_n to C and \mathbf{m} . Standard techniques from empirical process theory suffice to control the deviations between the empirical versions and their population limits (Figure 1.2(iv)). There are at least two useful approaches to control the size of the largest ball around \mathbf{m} within C . The first approach is sketched in (ii) and (iii): first, we lower bound the width of C uniformly over a range of “directions” h in \mathcal{H} (Figure 1.2(ii)). Then we determine how centered \mathbf{m} lies within C in each direction h (Figure 1.2(iii)). Combining these two arguments, we can derive a lower bound on the size of the largest ball around \mathbf{m} in C . The second approach is quite different in that it does not try to control the width of the set C explicitly. Instead, it uses the spectrum of the covariance operator to derive lower bounds on the largest ball around \mathbf{m} in C . In particular,

a simple argument using the Paley–Zygmund inequality goes a long way and leads to lower bounds that are controlled by the smallest non-zero eigenvalue of the centered covariance operator.




Figure 1.2. The figure summarizes some of the key questions we address in this paper: (i) This is the central question in this paper; “how large a ball exists within the empirical convex set C_n around \mathbf{m}_n ?” (ii) This question can be addressed by first controlling the width of C itself in different directions $h_1, h_2, \dots \in \mathcal{H}$. The width in such a direction h is the size of the projection of C on the span of the function $h \in \mathcal{H}$. Lower bounds on the width that hold simultaneously for all relevant h translate to the existence of a ball in C ; furthermore, the size of the ball is directly related to the lower bounds on the width. (iii) We need not just any ball in C but one that is centered at \mathbf{m} . Now, generally, \mathbf{m} can lie close to the boundary and no large ball around it might exist. However, under certain natural conditions, it can be ruled out that \mathbf{m} will lie too closely to the boundary. In particular, under these conditions, we can control the ratio of a/b for the segments shown in the figure. Controlling this ratio for all relevant $h \in \mathcal{H}$ allows us to show that there exists a ball around \mathbf{m} in C . (iv) To translate this back to C_n and \mathbf{m}_n we are making use of empirical process theory to control the convergence of $C_n \rightarrow C$ and $\mathbf{m}_n \rightarrow \mathbf{m}$ which allows us to lower bound the size of a ball around \mathbf{m}_n in C_n with high probability. Similarly to (ii) we control the convergence per direction h and then use high probability guarantees that hold simultaneously for all relevant h .

1.1. Lower bounding the width of C

When trying to control the width of C the first thing one notices is that we seem to know relatively little about C . Even the RKHS \mathcal{H} itself is usually only accessed through k and we do not have easy access to a basis of \mathcal{H} . So it might come as a surprise that there is a relatively simple way to access the width of C . The key to bounding the width is that

$$\text{width}_h(C) = 2 \inf_{c \in \mathbb{R}} \|h - c\mathbf{1}\|_\infty,$$

where $h \in \mathcal{H}$, $\|h\| = 1$, and $\mathbf{1}$ is the constant function that is equal to 1 everywhere. This relationship holds because

$$\sup_{g \in C} \langle h, g \rangle - \inf_{g \in C} \langle h, g \rangle = \sup_{x \in \mathcal{X}} \langle h, k(x, \cdot) \rangle - \inf_{x \in \mathcal{X}} \langle h, k(x, \cdot) \rangle = \sup_{x \in \mathcal{X}} h(x) - \inf_{x \in \mathcal{X}} h(x).$$

The relevance of this equality is that it reduces the problem of measuring the width to the problem of measuring how well constant functions can be approximated by functions in the RKHS. The question of how well certain functions can be approximated by RKHS functions is well understood when the RKHS is infinite dimensional. In particular, the K-functional is a common tool to control the approximation quality, and results about the K-functional can be brought to bear to provide bounds on the width of C . However, in the finite dimensional setting, these results are of limited use. We develop for this case a simple approach to measure how well constant functions can be approximated: if the constant functions do not lie in the RKHS \mathcal{H} then we can construct a new RKHS \mathcal{H}^+ by introducing the kernel function $k^+ = k + \mathbf{1} \otimes \mathbf{1}$, where k is the kernel of \mathcal{H} . The RKHS \mathcal{H}^+ then contains the constant functions and $\mathcal{H} \subset \mathcal{H}^+$. In fact, we have an isometric embedding of \mathcal{H} into \mathcal{H}^+ . Now, in \mathcal{H}^+ it is easy to measure how well constant functions can be approximated by functions in the unit sphere of \mathcal{H} . In detail,

$$\inf_{h \in \mathcal{H}, \|h\|=1} \inf_{c \in \mathbb{R}} \|h - c\mathbf{1}\|_{\mathcal{H}^+} = 1.$$

There are different ways to move from the norm of \mathcal{H}^+ to $\|\cdot\|_\infty$ which we summarize in Lemma 2.3 p. 507. One of these approaches applies if k^+ is a Mercer kernel and $\tilde{\lambda}_{d+1} > 0$ is the smallest eigenvalue in the series expansion. In this case

$$2\tilde{\lambda}_{d+1}^{1/2} \leq \text{width}_h(C),$$

for all $h \in \mathcal{H}$, $\|h\| = 1$.

If the constant functions lie already in \mathcal{H} then a different approach is necessary. Let us mention that we only need to control the width of C within the affine subspace that is spanned by it. Since $\langle k(x, \cdot), \mathbf{1} \rangle = 1$ for all $x \in \mathcal{X}$ we can observe that the space spanned by $\mathbf{1}$ is perpendicular to the affine subspace of C . To get a lower bound on $\text{width}_h(C)$ for functions h in the affine subspace we can consider the kernel $k^- = k - \|\mathbf{1}\|^2 \mathbf{1} \otimes \mathbf{1}$ and the corresponding RKHS \mathcal{H}^- . The constant functions do not lie in \mathcal{H}^- and \mathcal{H}^- can be isometrically embedded in \mathcal{H} . Most importantly the functions $h \in \mathcal{H}^-$ of norm $\|h\|_{\mathcal{H}^-} = 1$ are exactly the directions in which we need to bound the width of C . Now, with an approach analogous to the one involving \mathcal{H} and \mathcal{H}^+ we get a lower bound of the form

$$2\tilde{\lambda}_d^{1/2} \leq \text{width}_h(C),$$

for all $h \in \mathcal{H}^-$, $\|h\|_{\mathcal{H}^-} = 1$, and with $\tilde{\lambda}_d$ being the smallest eigenvalue of the Mercer decomposition of the kernel k . Proposition 2.10 p. 514 contains these results and results about related approaches to bound the width.

1.2. Locating \mathbf{m} within C

Controlling the width of C alone is insufficient since \mathbf{m} might lie in the boundary of C . We need to complement the lower bounds on the width with results that tell us how centered \mathbf{m} lies. This can be achieved by controlling the ratio a/b and b/a of the segments along any function h from \mathbf{m} to the boundary as depicted in Figure 1.2 (iii). An observation that is useful in this context is the following: if we have a probability measure on \mathbb{R} which has a mean value of zero and there exists some measurable set B with $\inf B \geq \epsilon > 0$ and $P(B) > 0$, then there will be probability mass on the negative axis since otherwise

$$0 = \int_{\mathbb{R}} x dP = \int_{[0, \infty)} x dP \geq \int_B x dP \geq \epsilon P(B) > 0.$$

A similar argument applies to C and \mathbf{m} . For instance, if we have a uniform distribution on the boundary of the ellipse shown in Figure 1.2 (iii), then \mathbf{m} cannot lie in the boundary: otherwise, there would exist a function $h \in \mathcal{H}$, $\|h\| = 1$, such that $\langle h, k(x, \cdot) \rangle \geq \langle h, \mathbf{m} \rangle$ for all $x \in \mathcal{X}$ and an $\epsilon > 0$ such that $A = \{x : \langle h, k(x, \cdot) - \mathbf{m} \rangle \geq \epsilon\}$ has non-zero measure. Hence,

$$0 = \int \langle h, k(x, \cdot) - \mathbf{m} \rangle dP(x) \geq \epsilon P(A) > 0.$$

Combining this argument with a Lipschitz assumption on the kernel function and a lower bound on the density allows us to show that \mathbf{m} has to lie away from the boundary. How far away it has to lie is made precise in Proposition 2.12 p. 521.

1.3. Convergence of C_n to C

To transfer the results about C and \mathbf{m} to C_n and \mathbf{m}_n we use VC and Rademacher arguments to bound the difference between C_n and C , and \mathbf{m}_n and \mathbf{m} . For controlling $\|\mathbf{m}_n - \mathbf{m}\|$ a standard argument suffices. However, it is less clear how to best control the difference between C_n and C .

The approach that we are taking is the following. We consider indicators $\chi\{\langle k(X, \cdot) - \mathbf{m}, h \rangle \leq -c\}$ where X is a random variable with the same distribution as X_1, \dots, X_n and c is a constant that we vary. Observe that whenever

$$P\chi\{\langle k(X, \cdot) - \mathbf{m}, h \rangle \leq -c\} > 0$$

then there is a point $x \in \mathcal{X}$, such that $\langle k(x, \cdot) - \mathbf{m}, h \rangle \leq -c$, or in other words, there is a point which lies c away from \mathbf{m} along h . A VC argument allows us to control all these indicators simultaneously over all h in the unit ball of \mathcal{H} and to show that for any such h ,

$$|P_n\chi\{\langle k(X, \cdot) - \mathbf{m}, h \rangle \leq -c\} - P\chi\{\langle k(X, \cdot) - \mathbf{m}, h \rangle \leq -c\}|,$$

is small for sufficiently large n . This allows us to show that C_n converges along h towards C with a certain rate and since we have guarantees that hold uniformly over the unit ball in \mathcal{H} we can derive a rate of convergence of C_n to C . A similar approach works for Rademacher complexities with the main difference being that we have to approximate the indicator functions with continuous functions.

Both approaches rely on a lower bound on the probability that $h(X)$ attains values below a threshold. We use two different approaches to get such lower bounds: the first approach uses an assumption on the density (lower bounded away from zero) and a Lipschitz assumption on the functions in \mathcal{H} . The second approach uses assumptions on the covariance operator. The second approach is more general in the sense that assumptions on the density imply a certain behavior of the covariance operator but our density assumption is certainly not the only way to control the covariance operator. On the other hand, the assumption on the covariance operator is quite abstract while the density assumption is in a sense very concrete.

Combining these different arguments allows us to control the size of the ball around \mathbf{m}_n . In particular, our first theorem combines the Rademacher approach with a density assumption (Theorem 2.16 p. 529). This approach brings together some of the results on the width of C , the location of \mathbf{m} and the convergence results to show that for large enough n there is with high probability a ball of a certain radius around \mathbf{m}_n in C_n . In detail, there exists a ball of size

$$\delta = \frac{c(\tilde{\lambda}_l^{1/2} \wedge L)^{l+1} \beta_l}{(l+1)(2L)^l},$$

where $\mathcal{X} = [0, 1]^l$, L is the Lipschitz constant, $\tilde{\lambda}_d$ the smallest eigenvalue of the Mercer decomposition of k , $\tilde{c} > 0$ is a lower bound on the density of the law of X_1 on \mathcal{X} and β_l is the Lebesgue measure of the l -dimensional unit ball in \mathbb{R}^l .

With probability $q \in (0, 1)$ there then exists a ball of radius $\delta/4$ around \mathbf{m}_n in C_n whenever n is greater than

$$n \geq \left(\frac{\sqrt{2 \log(4/q)} + 96 \|k\|_\infty^{1/2} / \delta}{\tilde{c} \beta_l (\delta/8L)^l} \right)^2 \vee \left(\frac{\sqrt{2} \|k\|_\infty^{1/2} (1 + \sqrt{\log(2/q)})}{\delta/4} \right)^2.$$

We can observe that δ is strongly dependent on the dimension l of the space \mathcal{X} . This stems from our approach: we identify a point $x_0 \in \mathcal{X}$ which corresponds to an element $k(x_0, \cdot) \in \mathcal{H}$ that lies far away from \mathbf{m} . We then identify a second point x_1 such that $k(x_1, \cdot)$ lies in the opposite direction of $k(x_0, \cdot)$ with respect to \mathbf{m} . If the space is low dimensional then $k(x_1, \cdot)$ needs to lie far from \mathbf{m} to counter the mass that is accumulated around $k(x_0, \cdot)$ and, thus, \mathbf{m} lies reasonably centered between $k(x_0, \cdot)$ and $k(x_1, \cdot)$. However, when the space is high dimensional then no single point $k(x_1, \cdot)$ has to lie far away from \mathbf{m} because the mass accumulated around $k(x_0, \cdot)$ can be countered by “many points” that lie close to \mathbf{m} and \mathbf{m} can lie significantly closer to the boundary.

To contrast this worst-case bound with the best-case scenario, observe that there is a point in C such that a ball of radius $2\tilde{\lambda}_d^{1/2}$ lies around it within C . The factor $\tilde{\lambda}_d^{1/2}$ itself is in all likelihood tight and reflects the fact that the convex set C is very small in certain directions.

1.4. Assumptions on the spectrum of the covariance operator

An alternative approach to controlling the width of C in different directions $h \in \mathcal{H}$ and then determining how centered \mathbf{m} lies in each direction is to use assumptions on

the covariance operator. In fact, the argument that involves the covariance operator is considerably simpler: when

$$E(h^2(X)) - E^2(h(X)) \geq \bar{\lambda} > 0$$

for some positive $\bar{\lambda}$ then $|h(X)|$ must attain large enough values with some non negligible probability. Furthermore, when h is a bounded function then both $(h(X) - E(h(X)))^+$ and $(h(X) - E(h(X)))^-$ must be large with a non-negligible probability. A simple argument involving the Paley–Zygmund inequality suffices to make these statements precise. To get a lower bound on the largest ball around \mathbf{m} in C we have to control all h in the unit ball with this approach. In terms of the spectrum of the covariance operator this means that we have to use the smallest non-zero eigenvalue of the covariance operator as $\bar{\lambda}$.

Another advantage of the covariance operator approach is that it adapts nicely to settings where the distribution has support S that is not equal to \mathcal{X} . Effectively, algorithms like the CGM or kernel herding work implicitly with a subspace of \mathcal{H} that is isometric to an RKHS \mathcal{H}_S with kernel $k|_{S \times S}$ (the restriction of k to $S \times S$) and for \mathcal{H}_S we have a covariance operator that has the same non-zero eigenvalues as the covariance operator for \mathcal{H} . Hence, we can use the same $\bar{\lambda}$ for \mathcal{H}_S as for \mathcal{H} and we can control the largest ball around \mathbf{m}_S in \mathcal{H}_S through this argument. Theorem 2.17 p. 531 is based on that argument.

1.5. Adapting the approach to concrete statistical problems

Most methods for inference do not use \mathbf{m} itself but related quantities. For example, in the least squares problem, where we try to fit observations Y_i through $f(X_i)$ with some function f in an RKHS, we have

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n (f(X_i) - Y_i)^2 \\ &= \frac{1}{n} \sum_{i=1}^n \langle f \otimes f, k(X_i, \cdot) \otimes k(X_i, \cdot) \rangle_{\otimes} - \frac{2}{n} \sum_{i=1}^n \langle f, Y_i k(X_i, \cdot) \rangle + \frac{1}{n} \sum_{i=1}^n Y_i^2 \\ &= \langle f \otimes f, \mathfrak{C}_n \rangle_{\mathcal{H} \otimes \mathcal{H}} + 2 \langle f, \mathbf{m}_{y,n} \rangle + \frac{1}{n} \sum_{i=1}^n Y_i^2, \end{aligned}$$

where we denote by $\mathcal{H} \otimes \mathcal{H}$ the tensor space $\mathcal{H} \otimes \mathcal{H}$ when the functions are restricted to the diagonal $\Delta = \{(x, x) : x \in \mathcal{X}\}$, $\mathfrak{C}_n = (1/n) \sum_{i=1}^n k(X_i, \cdot) \otimes k(X_i, \cdot)|_{\Delta}$ and $\mathbf{m}_{y,n} = (1/n) \sum_{i=1}^n Y_i k(X_i, \cdot)$.

There are significant similarities between the problem of compressing \mathbf{m}_n and that of compressing \mathfrak{C}_n or $\mathbf{m}_{y,n}$. We discuss these in Section 3. Let us highlight a few results.

The empirical covariance operator \mathfrak{C}_n can be dealt with quite easily by associating it to the element $(1/n) \sum_{i=1}^n \kappa(X_i, \cdot)$, where $\kappa(x, y) = k^2(x, y)$. This way one can apply all the results we developed for \mathbf{m}_n to \mathfrak{C}_n , one only has to substitute κ for k .

Dealing with the element $\mathbf{m}_{y,n}$ is more challenging and there is a certain degree of freedom of how to phrase the compression problem. A natural and simple choice

is to consider $Y_i k(X_i, \cdot)$ as the random elements which attain values in \mathcal{H} . A first indicator that things are more complicated is that when Y_i is unbounded then we run into serious problems when trying to define a bounded convex set that contains $(1/n) \sum_{i=1}^n Y_i k(X_i, \cdot)$. Things simplify if we assume boundedness of the Y_i and make some natural assumptions about how the data is generated. In particular, if we assume that X_1, \dots, X_n are i.i.d. and $Y_i = f_0(X_i) + \epsilon_i$, where f_0 is some bounded measurable function, the ϵ_i 's are i.i.d., centered, independent of X_1, \dots, X_n and bounded by $|\epsilon_i| \leq b$ a.s., then $\mathbf{m}_{y,n}$ converges to

$$\mathbf{m}_y = \int f_0(X_1) k(X_1, \cdot) dP \in \mathcal{H}$$

and \mathbf{m}_y is contained in the convex set

$$C_y = \text{cch}\{(f_0(x) \pm b)k(x, \cdot) : x \in \mathcal{X}\}.$$

In this setting there is also a simple relationship between the width of C_y and C : consider some $h \in \mathcal{H}$, $\|h\| = 1$, then

$$\text{width}_h(C_y) \geq b \text{width}_h(C)$$

and results on $\text{width}_h(C)$ are applicable. The downside of this approach is that the convergence of the empirical convex set towards C_y can be very slow since the $|\epsilon_i|$ might only have a low probability of attaining values close to b . This problem can be circumvented by using an alternative approach. Instead of considering the convergence of the empirical convex set to a suitable population limit we can directly work with the empirical convex set and analyze how deep the empirical mean element lies within that set. We develop this approach in Section 3.3.1. The discussion in that section cumulates in Proposition 3.1, which provides lower bounds on the radius of a ball that is centered on the empirical mean element $\mathbf{m}_{y,n}$ and which is contained within the empirical convex set.

We extend this approach to the case of unbounded Y_i by using random variables \widehat{Y}_i that are capped at a certain, n dependent, threshold. There are a variety of technical challenges that have to be overcome to make this approach work. In particular, one has to verify that the empirical mean element corresponding to the capped random variables is close to the empirical mean element of the original variables when the threshold of the cap is selected appropriately. Also, one has now to work with a family of covariance operators corresponding to the different thresholds and the corresponding capped random variables. We show that the lowest eigenvalues of these covariance operators are close to the lowest eigenvalue of the original covariance operator if the threshold for the cap is set in the right way. Proposition 3.3 contains the details of that result.

1.5.1. Simultaneous approximation

Up to now we considered the approximation problems in isolation but it also makes sense to try to approximate \mathfrak{C}_n simultaneously to $\mathbf{m}_{y,n}$ by selecting elements $Y_i k(X_i, \cdot)$ that reduce the approximation error for both elements. Quite a different set of techniques are needed to deal with this simultaneous approximation problem.

In Section 3.4 we develop an approach based on direct sums of Hilbert spaces to deal with this problem. The analysis is much more intricate and interesting than for the individual approximation problems. In Figure 1.1 (ii) the high level approach is visualized. The space $\widehat{\mathcal{H} \odot \mathcal{H}}$ is the space of functions $\mathcal{H} \odot \mathcal{H}$ when trivially extended from \mathcal{X} to $\mathbb{R} \times \mathcal{X}$. The space $\mathbb{R}' \otimes \mathcal{H}$, where \mathbb{R}' denotes the dual space of \mathbb{R} , is an RKHS with kernel function $((y_1, x_1), (y_2, x_2)) \mapsto \langle y_1, y_2 \rangle_{\mathbb{R}} k(x_1, x_2)$ which is also defined on $\mathbb{R} \times \mathcal{X}$. The convex sets we introduced above have natural analogues in $\widehat{\mathcal{H} \odot \mathcal{H}}$ and in $\mathbb{R}' \otimes \mathcal{H}$. By taking the direct sum of these spaces we also get a sort of direct sum of these convex sets and we are trying again to control quantities like the width of that set. The particular problem of approximating \mathfrak{C}_n simultaneously to $\mathfrak{m}_{y,n}$ is benefiting from the fact that $\widehat{\mathcal{H} \odot \mathcal{H}} \cap (\mathbb{R}' \otimes \mathcal{H}) = \{0\}$. This allows us to define an RKHS that is isometrically isomorphic to the direct sum. The analysis of the simultaneous approximation problem then breaks down to studying the empirical mean element and the empirical convex set within that RKHS.

The situation that the two Hilbert spaces that we combine through the direct sum are not overlapping is rather special. For instance, if we try to approximate \mathfrak{m} and \mathfrak{C} simultaneously then the Hilbert spaces overlap, which adds another layer of difficulties. We are developing for this case a quotient space approach that factors out the intersection between the two Hilbert spaces. An interesting finding in this context is that the direct sum cannot be related directly to an RKHS but, like in the case of approximating \mathfrak{m} and \mathfrak{C} simultaneously, the affine subspace spanned by the convex set can be isometrically isomorphic to an RKHS which then allows us to use results we developed for RKHSs (see Lemma 3.5 p. 563).

When we apply the conditional gradient method to the above RKHSs then we will not end up with a coreset of data points but with elements in $\widehat{\mathcal{H} \odot \mathcal{H}}, \mathbb{R}' \otimes \mathcal{H}$ or $\widehat{\mathcal{H} \odot \mathcal{H}} \oplus (\mathbb{R}' \otimes \mathcal{H})$. However, that is not a major obstacle and it is for various problems quite easy to adapt the algorithms to deal with these approximations; we highlight that approach for kernel ridge regression in Section 5.

1.6. Implications for algorithms

The various results that we derived to control the size of the largest ball around \mathfrak{m}_n in C_n can be translated directly to results for algorithms like the CGM. In particular, we can give high probability guarantees on the approximation error when the CGM is being run for t iterations and we can give guarantees on the expected size of a coreset when the kernel herding algorithm is used with a stopping criterion that is an error of below $n^{-1/2}$. The corresponding results are contained in Section 4.

One problem with these algorithms is that they require an upfront computation of order $O(n^2)$ which is too high for large-scale data. Standard approaches to scale the CGM to large-scale problems do not seem to yield direct computational advantages but there are some interesting directions to explore. In particular, a divide-and-conquer approach has some intriguing features. The performance of the approach depends to a large extent on the bias of the algorithms (CGM or kernel herding). Section 4 contains a detailed discussion of these ideas.

1.7. Slow rate of convergence in infinite dimensions

It was observed in [BLJO12] that the proof technique used to derive fast rates of convergence for the kernel herding algorithm and the conditional gradient method cannot be applied to compact sets in infinite dimensional RKHSs since compact sets in such spaces do not contain norm balls. It was later found that there are general limits to how well the representer of the empirical measure can be approximated. In particular, [PT20, Theorem 3.1] states that there exists a set of n points x_1, \dots, x_n in \mathbb{R}^d , for large enough d and $n \geq d$, such that for any set of points y_1, \dots, y_l with $l < \sqrt{dn}/2$ it holds that

$$\left\| \frac{1}{l} \sum_{i=1}^l k(y_i, \cdot) - \frac{1}{n} \sum_{i=1}^n k(x_i, \cdot) \right\| > \|k\|_{\infty}^{-1} n^{-1/2},$$

under mild assumptions on the kernel. This implies that for this particular set of elements x_1, \dots, x_n there cannot be any significant compression of the element $\mathbf{m} = (1/n) \sum_{i=1}^n k(x_i, \cdot)$. The argument in [PT20] is not stochastic and is not concerned with draws of samples X_1, \dots, X_n from a distribution, but it seems likely that the argument can be extended to provide restrictions on how well \mathbf{m}_n can be approximated by a core-set of points in high probability (over the values that \mathbf{m}_n attains). Nevertheless, there is hope to circumvent the barriers erected by this theorem. First of all, the construction uses approximations of the form $(1/l) \sum_{i=1}^l k(y_i, \cdot)$ and not arbitrary convex combinations of the elements $k(y_i, \cdot), i \leq l$. A greedy algorithm to find such a core-set of points y_1, \dots, y_l requires generally significantly more points than algorithms that approximate \mathbf{m}_n with convex combinations of elements $k(y_i, \cdot)$. An interesting question is therefore if there exists an inherent limitation for approximating with core-sets that can be avoided by more general convex combinations, or if this difference in performance is simply due to the algorithms (kernel herding vs. CGM). There is a simple argument that hints at the former: consider the set $\mathcal{X} = [0, 1]^{d_{\mathcal{X}}}$ for some positive $d_{\mathcal{X}} \in \mathbb{N}$ and a continuous kernel function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ whose corresponding RKHS is infinite dimensional and separable. The set C is then compact and for any orthonormal basis $\{e_i\}_{i \geq 1}$ of \mathcal{H} it holds that $\sup_{f \in C} \langle e_i, f \rangle - \inf_{g \in C} \langle g, e_i \rangle$ converges to zero as $i \rightarrow \infty$. The rate with which this series converges is in all likelihood of crucial importance for determining how well \mathbf{m}_n can be approximated. Therefore, let us introduce

$$d(P_{U_i}C, C) = \sup_{f \in C} \|P_{U_i}f - f\|,$$

where U_i is the subspace spanned by e_1, \dots, e_i and P_{U_i} is the orthogonal projection onto this subspace. Now, Caratheodory's theorem tell us that for $i \geq 1$ there exists a convex combination $\widehat{\mathbf{m}}_i$ of $i + 1$ elements $k(x_1, \cdot), \dots, k(x_{i+1}, \cdot)$ such that

$$\|\mathbf{m}_n - \widehat{\mathbf{m}}_i\| \leq d(P_{U_i}C, C).$$

In other words, when $d(P_{U_i}C, C)$ is of order $i^{-\alpha}$, for some $\alpha > 0$, and if we are aiming for an approximation error of $n^{-1/2}$ then we need approximately $n^{1/2\alpha}$ many points. Furthermore, when $d(P_{U_i}C, C)$ falls exponentially fast, say with order $\exp(-i)$, then $\log(n)$ many points suffice.




Figure 1.3. The figure shows lower bounds on $|\langle e_n, w \rangle|$ in dependence of the first element a_m that has not yet been chosen. The shaded area is a lower bound on $\|w\|$ when $m = 10^{20}$. The norm of w goes to infinity in m which implies that the kernel herding algorithm converges with a rate that is slower than $1/t$.

Another important aspect of the compression problem that is not captured by the theorem is the dependence of the compression problem on the distribution of the data. For example, in [BLJO12] a lower bound on the density of the distribution was crucial for deriving fast rates of convergence in certain settings. This is because such properties of the density translate directly to geometric properties of the approximation problem (the existence of a ball around \mathbf{m}_n in C_n). Similarly, in this paper, we use the size of the largest non-zero eigenvalue of the covariance operator to control the rate of convergence. One might wonder if such properties also influence the compression performance in infinite dimensional RKHSs. To this end, we provide an example that shows that an assumption on the density alone will in all likelihood be insufficient. The example we construct is *not universal* in the sense that we show that the kernel herding algorithm does not achieve its fast rate of $1/t$ of approximation in this example. As in the example from [PT20], we construct a particular target \mathbf{m} and do not consider the empirical version \mathbf{m}_n . However, our construction incorporates properties of the underlying probability measure and might serve as a starting point for more refined analyses that use properties of the distribution of the data. The counter-example is constructed for the kernel herding algorithm and not the conditional gradient methods since the behavior of the kernel herding algorithm is easier to control but we strongly suspect that similar problems will also occur with the conditional gradient method.

In detail, the example we construct shows that there exists a continuous kernel on $[0, 1]$, a Borel probability measure on $[0, 1]$ which assigns positive measure to open subsets of $[0, 1]$, and an initialization for which the kernel herding algorithm converges with a slower rate than $1/t$ when approximating the representer \mathbf{m} of the probability measure (Theorem 6.1 p. 589). The construction of this example is somewhat involved since we need to gain control over the behavior of the kernel herding algorithm. The basic intuition, however, is rather simple. We start with some

infinite dimensional Hilbert space \mathcal{H} and an orthonormal sequence $\{e_n\}_{n \geq 1}$ in it. The construction is best explained when assuming that $\mathbf{m} = 0$ (we cannot set it exactly to 0 and need later a minor modification). We then construct a compact convex set that contains elements $\{a_n\}_{n \geq 1}$, $\{b_n\}_{n \geq 1}$, where each a_n is a positive multiple of e_n and each b_n is a negative multiple of e_n . Furthermore, b_n is of significantly smaller magnitude than a_n . Consider now an initialization of the algorithm with an element $c \in \mathcal{H}$ which is of small magnitude compared to the a_n and has a positive inner product with each a_n . Because of this positive inner product the different a_n 's will be chosen one by one by the algorithm and because the b_n 's are of small magnitude compared to the a_n 's hardly any weight will be reduced in the directions e_n . This way the element w_t , which measures the approximation error at iteration t , builds up mass in the different directions e_n and its norm grows in t . The construction is more involved than this sketch, but, a suitably adapted version of this approach allows us to show that so much mass will be added to w_t that its norm diverges to infinity. This effect is visualized in Figure 1.3. The figure shows four different w_t as inner products with e_n (n being shown on the x -axis). The shaded area continues past the right end of the plot (the limit of the shaded area is given in the legend: 10^5 for the black line etc.). One can observe that the right limit of the shaded area grows significantly from the black line to the red line, i.e. from 10^5 to 10^{20} . While the right limit grows exponentially the left limit hardly changes. This is due to the small scale of the b_n 's. As a result the overall mass in the shaded area, which corresponds to $\|w_t\|$, diverges to infinity. This implies then directly that the algorithm cannot converge with the fast rate of $1/t$ that is achieved under similar assumptions in the finite dimensional setting. All that then remains to complete the example is to show that there exists a continuous kernel that gives rise to this setup. We construct first a continuous function $\phi : [0, 1] \rightarrow \mathcal{H}$ that goes over all a_n and b_n and we then use this Hilbert space and the continuous function to construct an RKHS with a continuous kernel function.

1.8. Literature

The concept of a *coreset* is known for at least two decades and there is a wide range of literature on its application to machine learning, Bayesian statistics and geometric approximation problems (e.g. [AHPV05, BHPI02, HCB16]). It is natural to apply the conditional gradient method [FW56] in that context (e.g. [HCB16]).

The kernel herding algorithm and the conditional gradient method are greedy approximation algorithm as they choose at each stage t an element that minimizes the remaining error. Greedy algorithms will generally not return the best possible approximation that can be achieved in t steps but they are easy to compute. This is a big advantage since in the large data context computational efficiency is paramount. Greedy algorithms for approximating functions have been popular at least since the late nineties. An overview of the most popular approaches is provided in [Tem11]. The approach is here to make use of a basis of a function space, say of a Sobolev or Besov space, to approximate elements inside these function spaces in a greedy fashion. An important generalization is to use so-called dictionaries which are families

of functions that are not necessarily linearly independent, i.e. there are redundancies in the representation of elements in the function space. These approaches are very natural if one has access to a basis or related families of functions. In contrast to this approach, we are interested in approximating subsets of the function space that are naturally described by point-evaluators, a kernel function, or, more generally, a set of extremes of a convex set. Instead of working then with linear subspaces of the function space we are working with convex subsets of the function spaces and we apply greedy algorithms to approximate elements inside such convex sets.

The methods we are studying compress the sample into a potentially small subset of the original sample while retaining optimal, or nearly optimal, rates of convergence. While our approach is inspired by various optimization methods there are links to *sample compression schemes* as introduced in [FW95, LW86]. Sample compression schemes are concerned with the inference of “concepts”, which are indicators χ_A , A a Borel subset of some topological space X . In this setting, one has given a set of concepts that contains the concept χ_A , or are sufficient to approximate χ_A in a suitable way, and one likes to infer χ_A from observations $(x_1, y_1), \dots, (x_n, y_n)$, $x_i \in \mathcal{X}$, $y_i \in \{0, 1\}$. A sample compression scheme compresses these observations into a subset that is sufficient to reconstruct the original labels y_i for all x_i , $i \leq n$, if the observations are consistent with some concept $\chi_{A'}$, where $\chi_{A'}$ is contained in the predefined set of concepts. Compressibility is directly linked to VC-theory: in [FW95] it is shown that, under some technical conditions, sets with VC-dimension d are d -compressible, meaning that one can always reduce the sample to a sub-sample of size d while still being able to reconstruct the sample in the above sense. Furthermore, it is not possible to compress the sample to less than d -points without losing the reconstructability property. Our aim is quite different in that we do not care about being able to reconstruct the original labels. In that sense our approach is more closely related to sufficient statistics which compress the data to facilitate inference. That being said, there are interesting parallels. For instance, Caratheodory’s theorem tells us that, in our setting, there is a compression of the data down to $d + 1$ -points if we work with a d -dimensional RKHS; such an RKHS has VC-dimension d .

The question of how to construct *coresets* for \mathbf{m} has garnered significant attention in recent years. In [DM21] a good overview is given that covers recent approaches most of which focus on the infinite dimensional setting. In the context of finite dimensional RKHSs it is worth mentioning the paper [HS14] which studies linear kernel functions and shows that under certain conditions they can achieve a compression down to $n^{1/2}$.

Naturally, there are a variety of alternative approaches to deal with large scale data in the RKHS context. In particular, when the RKHS is finite dimensional with dimension d it is straight forward to represent \mathbf{m}_n using a basis expansion: take points X_1, \dots, X_d such that $k(X_1, \cdot), \dots, k(X_d, \cdot)$ are linearly independent and apply the Gram–Schmidt orthogonalization procedure to gain a basis e_1, \dots, e_d of \mathcal{H} then \mathbf{m} can be written as a linear combination of e_1, \dots, e_d which implies that it can be written as a linear combination of $k(X_1, \cdot), \dots, k(X_d, \cdot)$. In more detail, the

coefficients $\alpha_1, \dots, \alpha_d$, such that $\mathbf{m}_n = \sum_{i=1}^d \alpha_i k(X_i, \cdot)$, can be computed recursively by first computing the basis representation through

$$\begin{aligned} \langle e_1, \mathbf{m} \rangle &= \langle k(X_1, \cdot), \mathbf{m}_n \rangle / k(X_1, X_1) \\ &\vdots \\ c_i &= \langle k(X_i, \cdot), \mathbf{m}_n \rangle - \sum_{j=1}^{i-1} \langle k(X_i, \cdot), e_{i-j} \rangle \langle e_{i-j}, \mathbf{m}_n \rangle \\ \langle e_i, \mathbf{m}_n \rangle &= c_i / \left\| k(X_i, \cdot) - \sum_{j=1}^{i-1} \langle k(X_{i-j}, \cdot), e_j \rangle e_j \right\|, \end{aligned}$$

and then to link this back to the coefficients of $k(X_1, \cdot), \dots, k(X_d, \cdot)$. To perform this Gram–Schmidt procedure it is necessary to compute $k(X_i, X_j)$ for all $i, j \leq d$. In other words, we need in the order of d^2 many kernel evaluations. This is a negligible factor when $d \ll n$. Similarly, one can solve concrete statistical problems, like a linear regression problem, by using a $d \times d$ covariance matrix instead of the kernel matrix; one way to gain such a covariance matrix is to use again the Gram–Schmidt procedure. Our aim in this paper is not to compete with these methods in terms of runtime performance in the context of $d \ll n$, but to gain insights into the behavior of greedy algorithms in the absence of complications that arise in infinite dimensional settings.

The above approach is closely related to the *Nyström method* for improving the run-time of kernel methods. The Nyström method was first applied to machine learning problems in [WS00]. It has since seen tremendous success, and there is now a vast literature concerning applications, variations, and theoretical investigations of the method. The most relevant results for this paper concern the compression rates for kernel regression and kernel PCA. For kernel regression, compression rates were developed in [RCR15] for the case that the unknown regression functions lie within the RKHS that is used to construct the estimator. For the more realistic setting, where RKHS functions are used to approximate more complicated functions, results have recently been obtained in [LLW23]. In more detail, [LLW23, Theorem 3.9] provides lower bounds on how many samples the Nyström method requires to preserve the optimal estimation error in kernel regression. This bound depends on the roughness of the target function, which is measured through a coefficient $r \in (0, 1]$, where $r \in [1/2, 1]$ corresponds to the case where the target function lies in the RKHS or is even smoother than RKHS functions. For most applications, the interesting regime is $r \in (0, 1/2)$, which corresponds to the case where the target function is not in the RKHS. There are two more coefficients $\gamma, \alpha \in [0, 1]$, which are related to the effective dimension. Given these parameters, the lower bound on the number of samples is of order $n^{\frac{\alpha}{2r+\gamma}}$ for n data points overall. For example, for $\gamma = \alpha = 1$ and assuming that the unknown regression function lies in the RKHS, one needs of order \sqrt{n} many samples. As the target function gets rougher ($r \rightarrow 0$) the number of samples increases to n and the speed-up due to the Nyström method becomes negligible.

Kernel PCA [SSM97, BBZ07] is another popular kernel algorithm to which the Nyström method can be applied. The most recent results that we are aware of quantify the reconstruction error within the RKHS when using kernel PCA [SSRR20]. This setting is akin to the regression setting, where the true regression function lies within the RKHS. This similarity is reflected in the compression guaranties: when the eigenvalues of the covariance operator decrease quadratically, one can achieve a reconstruction error of order $1/\sqrt{n}$ with \sqrt{n} many samples [SSRR20, Corollary 3]. We are unaware of attempts to move beyond the reconstruction error within the RKHS; such results could come in the form of a pre-processing step for kernel two-sample tests [GBR⁺12], where one compresses the mean embedding before applying it to potentially rougher test functions to quantify differences between measures.

The above results are particularly interesting when the RKHS is infinite dimensional. Since our current results are confined to finite dimensional spaces it is difficult to compare the results. That said, it seems worth to highlight the following: Nyström methods deal very well with small eigenvalues, but struggle when there are many large eigenvalues (roughly speaking, this corresponds to $r < 1/2$ in the regression setting). On the other hand, the coresets methods we investigate deal well with large eigenvalues but struggle with small eigenvalues. In that sense Nyström methods and coresets methods have complementary strengths and one can wonder if there are ways to combine the methods to obtain superior compression techniques.

1.9. Preliminaries

Throughout this paper we will be working with a set \mathcal{X} in which covariates or features attain values and a kernel function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ (see [PR16, Definition 2.12]). Recall that such a kernel function gives rise to an RKHS \mathcal{H} [PR16, Definition 2.14]. While \mathcal{X} does not need a particular structure to define a kernel on, we are interested in integrals involving k and we will assume for most of our results that \mathcal{X} is a measurable space and k is measurable in the sense that $k(x, \cdot) : \mathcal{X} \rightarrow \mathbb{R}$ is measurable for all $x \in \mathcal{X}$. This is equivalent to saying that any $h \in \mathcal{H}$ is a measurable function from \mathcal{X} to \mathbb{R} (see [SC08, Lemma 4.24]). We also use the notation $\phi(x) = k(x, \cdot)$ when this is convenient.

We are making use of empirical process theory in various places and to ease the application we will assume that our underlying probability space corresponds to a product space and the involved random variables are coordinate projections following essentially [Dud14, Section 3.1]. In detail, we will usually have a probability space $(\Omega, \mathcal{A}, \mu)$ with independent and identically distributed random variables X, X_1, X_2, \dots attaining values in $(\mathcal{X}, \mathcal{A}_{\mathcal{X}})$, where \mathcal{X} is a topological space and $\mathcal{A}_{\mathcal{X}}$ is a σ -algebra on \mathcal{X} , which are defined on this probability space. Natural choices for $\mathcal{A}_{\mathcal{X}}$ are the Borel-algebra or the domain of a Radon measure. We usually do not need assumptions on $\mathcal{A}_{\mathcal{X}}$ but at various points we need to guarantee that the support of the law P of X is well defined. In these cases we typically assume that P is a τ -additive topological measure and $\mathcal{A}_{\mathcal{X}}$ is its domain. Alternatively, we could assume that P is a Radon measure which guarantees that P is a τ -additive topological measure (see [Fre01, 411]). This is for a wide range of spaces not a strong assumption.

In particular, if \mathcal{X} is a Polish space then it is a Radon space [Fre01, 434K (b)] and the completion of a Borel measure on \mathcal{X} is a Radon measure [Fre01, 434F (a.iii), 211L]. Ω will usually be the product $\mathcal{X}^{\mathbb{N}}$ and for $\omega \in \Omega$, $X_i(\omega) = \omega_i \in \mathcal{X}$ for all $i \geq 1$, and $X(\omega) = \omega_0$. There are multiple natural choices for the σ -algebra \mathcal{A} . In [Dud14, Section 3.1] \mathcal{A} is the product σ -algebra which is the one that is generated by the cylinder sets, that is the smallest σ -algebra such that all cylinders which are defined by finite many coordinates are measurable. We use in this paper the completion of this σ -algebra as \mathcal{A} . If we have pairs (X_i, Y_i) , where X_i attains values in \mathcal{X} and Y_i in \mathbb{R} then we use the same setting but let $(X_i, Y_i)(\omega) \in \mathcal{X} \times \mathbb{R}$. We reserve P for the law of the random variables, e.g. the law of X , and use \Pr if we want to state probabilities of events in \mathcal{A} . In particular, $h(X) \in \mathcal{L}^1(\Omega, \mu)$ if, and only if, $h \in \mathcal{L}^1(\mathcal{X}, P)$ and, in this case, $\int h(X) d\mu = \int h dP$. The empirical measure P_n is $(1/n) \sum_{i=1}^n \delta_{X_i}$, where $\delta_x(A) = 1$ whenever $A \in \mathcal{B}_{\mathcal{X}}$ and $x \in A$; otherwise $\delta_x(A) = 0$. It is often useful to associate a measure space to P_n to be able to talk about random variables with law P_n . For this purpose we will use the measure space $(\mathcal{X}, \mathcal{B}_{\mathcal{X}})$ and equip it with the random measure P_n . A random variable will be the measurable function $\tilde{X} : \mathcal{X} \rightarrow \mathcal{X}$, $\tilde{X}(x) = x$. If we want to talk about a sequence of independent random variables with law P_n we use the product space with the product measure assigned to it.

1.9.1. Separable processes and Rademacher complexities

There are generally various measurability concerns when working with empirical processes. In this paper these can essentially be avoided by using separability of \mathcal{H} to guarantee that suprema are measurable. In the context of Rademacher complexities we use separability of \mathcal{H} typically in the following way. Assume we have $x_1, \dots, x_n \in \mathcal{X}$, let \mathcal{F} be the unit ball of \mathcal{H} and let $\epsilon_1, \dots, \epsilon_n$ be i.i.d. Rademacher variables. The map $h \mapsto \sum_{i=1}^n \epsilon_i h(x_i)$ is almost surely continuous on \mathcal{H} . In particular, $\sup_{h \in \mathcal{F}} \sum_{i=1}^n \epsilon_i h(x_i)$ is almost surely equal to a supremum over a countable subset of \mathcal{F} and, due to completeness of the probability space, it follows that $\sup_{h \in \mathcal{F}} \sum_{i=1}^n \epsilon_i h(X_i)$ is measurable. In particular, the Rademacher process is a separable stochastic process [GN16, Definition 4.1.2] and we have

$$(1.1) \quad E \left(\sup_{h \in \mathcal{F}} \sum_{i=1}^n \epsilon_i h(X_i) \right) = \sup_{F \subset \mathcal{F}, F \text{ finite}} E \left(\sup_{h \in F} \sum_{i=1}^n \epsilon_i h(X_i) \right).$$

When we have i.i.d. variables X_1, \dots, X_n which are independent of $\epsilon_1, \dots, \epsilon_n$ we will represent this probability space as a product space. It is common to condition wrt. X_1, \dots, X_n and to study $E_\epsilon \left(\sup_{h \in \mathcal{F}} \sum_{i=1}^n \epsilon_i h(X_i) \right)$, where E_ϵ denotes Kolmogorov's conditional expectation with respect to X_1, \dots, X_n . Fubini's theorem guarantees us in this setting that we can express E_ϵ as an integral wrt. the marginal measure corresponding to $\epsilon_1, \dots, \epsilon_n$.

1.9.2. Bochner integrals and $\mathcal{L}^p(\mu, \mathcal{H})$

We need in various places vector valued integrals. In particular, we make use of Bochner integrals and Hilbert-space valued \mathcal{L}^p spaces. Let $(\Omega, \mathcal{A}, \mu)$ be a probability

space and X a random variable that attains values in \mathcal{X} then by $\int f(X) d\mu$, $f : \Omega \rightarrow \mathcal{H}$ Bochner integrable, we mean the Bochner integral of the function $f(X) : \Omega \rightarrow \mathcal{H}$ with respect to the measure μ . The Hilbert space valued \mathcal{L}^p spaces, where $1 \leq p < \infty$, corresponding to this measure space are given by

$$\mathcal{L}^p(\mu; \mathcal{H}) = \left\{ f : \Omega \longrightarrow \mathbb{R} : f \text{ Bochner measurable and } \int \|f(\omega)\|^p d\mu(\omega) < \infty \right\}.$$

The seminorm on $\mathcal{L}^p(\mu; \mathcal{H})$ is $\|f\|_p^p = \int \|f\|^p d\mu$. We use bold fonts for the $\mathcal{L}^p(\mu; \mathcal{H})$ seminorms throughout this paper. As usual there are corresponding spaces L^p of equivalence classes with norms $\|\cdot\|_p$ under which these L^p spaces are complete. The space $L^2(P; \mathcal{H})$ is a Hilbert space with the inner product corresponding to the bi-linear function $\langle \cdot, \cdot \rangle_2$ on $\mathcal{L}^2(\mu; \mathcal{H})$ given by $\langle f, g \rangle_2 = \int \langle f(\omega), g(\omega) \rangle d\mu(\omega)$ whenever $f, g \in \mathcal{L}^2(\mu; \mathcal{H})$. Of particular importance to us is the Bochner integral $\int k(X, \cdot) d\mu \in \mathcal{H}$ which is well defined whenever $k(X, \cdot) \in \mathcal{L}^1(\mu; \mathcal{H})$ and \mathcal{H} is separable. We will denote this integral by \mathbf{m} . Finally, we have the following important relation between the inner product in \mathcal{H} and Bochner integrals: whenever $f \in \mathcal{L}^1(\mu; \mathcal{H})$ and $h \in \mathcal{H}$ then according to [DU77, Theorem 6, p. 47],

$$\langle h, \int f d\mu \rangle = \int \langle f, h \rangle d\mu.$$

In rare occasions we will make statements about equivalence classes and not functions itself. We use the notation f^\bullet to denote the equivalence class corresponding to f , i.e. if $f \in \mathcal{L}^2(\mu)$ then $f^\bullet \in L^2(\mu)$ and, similarly, for Hilbert space valued functions.

1.9.3. Tensor products

In various parts of this paper we make use of the *tensor product of two Hilbert spaces* \mathcal{H}_1 and \mathcal{H}_2 . One way to define this tensor product is to first define an algebraic tensor product of the vector spaces \mathcal{H}_1 and \mathcal{H}_2 ; given that we are only working with Hilbert spaces of functions it is natural to define the algebraic tensor product as

$$\left\{ f : \mathcal{X} \times \mathcal{Y} \longrightarrow \mathbb{R} : f(x, y) = \sum_{i=1}^n g_i(x)h_i(y), g_i \in \mathcal{H}_1, h_i \in \mathcal{H}_2, n \in \mathbb{N} \right\},$$

where we assume that functions in \mathcal{H}_1 map from \mathcal{X} to \mathbb{R} and functions in \mathcal{H}_2 from \mathcal{Y} to \mathbb{R} . That this is a tensor product for \mathcal{H}_1 and \mathcal{H}_2 can be verified by applying [DF93, Criterion 2.3]. Next, we equip the algebraic tensor product with the inner product $\langle g_1 \otimes h_1, g_2 \otimes h_2 \rangle_\otimes = \langle g_1, h_1 \rangle_1 \langle g_2, h_2 \rangle_2$, e.g. [Mur90, Theorem 6.3.1], and complete the resulting pre-Hilbert space. In the case where \mathcal{H}_1 and \mathcal{H}_2 are RKHSs with kernels k_1 and k_2 we have bounded point evaluators for elements in the pre-Hilbert space, i.e. $\langle h_1 \otimes h_2, k_1(x, \cdot) \otimes k_2(y, \cdot) \rangle_\otimes = h_1(x)h_2(y)$ for all $x \in \mathcal{X}, y \in \mathcal{Y}, h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$. Due to [Aro50, 2nd theorem, §4 p. 347] there is then a unique functional completion of the algebraic tensor product and we will use this completion when working with RKHSs. We do not use the algebraic tensor product itself and, in the following, will reserve the notation $(\mathcal{H}_1 \otimes \mathcal{H}_2, \langle \cdot, \cdot \rangle_\otimes)$ for the above defined tensor product of the two Hilbert spaces, that is $\mathcal{H}_1 \otimes \mathcal{H}_2$ is a Hilbert space with inner product $\langle \cdot, \cdot \rangle_\otimes$, and, whenever \mathcal{H}_1 and \mathcal{H}_2 are RKHSs, $\mathcal{H}_1 \otimes \mathcal{H}_2$ is a Hilbert

space of functions. In fact, in the latter case $\mathcal{H}_1 \otimes \mathcal{H}_2$ is an RKHS with kernel $\tilde{k}((x_1, y_1), (x_2, y_2)) = k_1(x_1, y_1)k_2(x_2, y_2)$. See also [PR16, Theorem 5.11].

When X, Y are independent random variables under the measure μ attaining values in $\mathcal{X}_1, \mathcal{X}_2$, k_1, k_2 are kernel functions on \mathcal{X}_1 and \mathcal{X}_2 respectively, $g \in \mathcal{H}_1, h \in \mathcal{H}_2$, and the Bochner integrals $\int k_1(X, \cdot)d\mu, \int k_2(Y, \cdot)d\mu, \int k_1(X, \cdot) \otimes k_2(Y, \cdot)d\mu$ are well defined then

$$\begin{aligned} \left\langle g \otimes h, \int k_1(X, \cdot) \otimes k_2(Y, \cdot)d\mu \right\rangle_{\otimes} &= \int g(X)h(Y)d\mu = \int g(X)d\mu \int h(Y)d\mu \\ &= \left\langle g, \int k_1(X, \cdot)d\mu \right\rangle_1 \left\langle h, \int k_2(Y, \cdot)d\mu \right\rangle_2 \\ &= \left\langle g \otimes h, \int k_1(X, \cdot)d\mu \otimes \int k_2(Y, \cdot)d\mu \right\rangle_{\otimes}. \end{aligned}$$

Since this holds for all $g \otimes h, g \in \mathcal{H}_1, h \in \mathcal{H}_2$,

$$(1.2) \quad \int k_1(X, \cdot) \otimes k_2(Y, \cdot)d\mu = \int k_1(X, \cdot)d\mu \otimes \int k_2(Y, \cdot)d\mu.$$

There is another natural way to define a tensor product for two RKHSs \mathcal{H}_1 and \mathcal{H}_2 that is often of use. Here, we identify the tensor product with a rank one operator mapping from \mathcal{H}_1 to \mathcal{H}_2 . To distinguish it from the above definition we will use $g \widehat{\otimes} h, g \in \mathcal{H}_1, h \in \mathcal{H}_2$, to denote this tensor product. Whenever $f, g \in \mathcal{H}_1, h \in \mathcal{H}_2$, the tensor product is defined by $(g \widehat{\otimes} h)(f) = \langle g, f \rangle_{\mathcal{H}_1} h \in \mathcal{H}_2$. Furthermore, we can define an inner product on this tensor space by letting $\langle f_1 \widehat{\otimes} f_2, h_1 \widehat{\otimes} h_2 \rangle_{\widehat{\otimes}} = \langle f_1, h_1 \rangle_{\mathcal{H}_1} \langle f_2, h_2 \rangle_{\mathcal{H}_2}, f_1, h_1 \in \mathcal{H}_1, f_2, h_2 \in \mathcal{H}_2$. Using Parseval’s identity one can observe that is just the usual inner product of the space $HS(\mathcal{H}_1, \mathcal{H}_2)$ of Hilbert–Schmidt operators and $\text{span}\{g \widehat{\otimes} h : g \in \mathcal{H}_1, h \in \mathcal{H}_2\}$ lies dense in $HS(\mathcal{H}_1, \mathcal{H}_2)$. It is therefore natural to use $HS(\mathcal{H}_1, \mathcal{H}_2)$ as the completion of the algebraic tensor product defined in terms of rank one operators. We will therefore denote the inner product between such tensors by $\langle \cdot, \cdot \rangle_{HS}$.

1.9.4. Covariance operators

A first application of this tensor product leads us to covariance operators. The covariance operator $\tilde{\mathcal{C}} : \mathcal{H} \rightarrow \mathcal{H}$, given by $\langle \tilde{\mathcal{C}}g, h \rangle = E(g(X)h(X))$, is linear ($\langle \tilde{\mathcal{C}}(\alpha f + g), h \rangle = \alpha E(f \times h) + E(g \times h) = \langle \alpha \tilde{\mathcal{C}}(f) + \tilde{\mathcal{C}}(g), h \rangle$ for all $h \in \mathcal{H}$ and, therefore, $\tilde{\mathcal{C}}(\alpha f + g) = \alpha \tilde{\mathcal{C}}(f) + \tilde{\mathcal{C}}(g)$ whenever $f, g \in \mathcal{H}, \alpha \in \mathbb{R}$) and is bounded whenever \mathcal{H} can be continuously embedded in $L^2(\mathcal{X}, P)$, i.e. for some $c > 0, \|h\|_2 \leq c\|h\|$ for all $h \in \mathcal{H}$, since then

$$\begin{aligned} \|\tilde{\mathcal{C}}\|_{\text{op}} &= \sup_{\|f\|=1} \|\tilde{\mathcal{C}}f\| = \sup_{\|f\|=1} \sup_{\|h\|=1} |\langle \tilde{\mathcal{C}}f, h \rangle| \\ &= \sup_{\|f\|=1} \sup_{\|h\|=1} |E(f \times h)| \leq \sup_{\|f\|=1} \sup_{\|h\|=1} \|f\|_2 \|h\|_2 \leq c^2. \end{aligned}$$

In fact it is a Hilbert–Schmidt operator whenever \mathcal{H} is separable and $k(X, X) \in \mathcal{L}^2(\mu)$ because then for any orthonormal basis $\{e_n\}_{n \in \mathbb{N}}$ of $\mathcal{H}, \sum_{n,m \in \mathbb{N}} |\langle \tilde{\mathcal{C}}e_n, e_m \rangle|^2 \leq E((\sum_{n \in \mathbb{N}} |\langle e_n, k(X, \cdot) \rangle|^2)^2) = E(k^2(X, X))$ due to Beppo Levi’s theorem. In this case $\tilde{\mathcal{C}}$ is also self-adjoint and the spectral theorem applies. Furthermore, we can write

the covariance operator as a Bochner-integral of the tensors $k(x, \cdot) \widehat{\otimes} k(x, \cdot)$, i.e. $\widetilde{\mathfrak{C}} = \int k(x, \cdot) \widehat{\otimes} k(x, \cdot) dP$. This Bochner integral is well defined and attains values in $HS(\mathcal{H})$ whenever $\int \|k(x, \cdot) \widehat{\otimes} k(x, \cdot)\|_{HS} dP = \int k(x, x) dP < \infty$ and \mathcal{H} is separable. Separability of \mathcal{H} is important in this context because it implies that $HS(\mathcal{H})$ is separable and Bochner measurability, that is necessary for the Bochner integral above to be well defined, is not a restrictive assumption [DF93, Appendix B12].

Observe that there is a close relationship between the eigen-decomposition of $\widetilde{\mathfrak{C}}$ and the expansion of the integral operator

$$T_k : \mathcal{L}^2(P) \longrightarrow \mathcal{L}^2(P), (T_k f)(y) = \int f(x)k(x, y) dP(x).$$

Whenever \mathcal{H} is infinite dimensional and Mercer’s theorem applies there exists an orthonormal sequence $\{e_i^\bullet\}_{i \geq 1}$ in $L^2(P)$ and corresponding values $\{\widetilde{\lambda}_i\}_{i \geq 1}$ in \mathbb{R} such that e_i are eigenfunctions of T_k with eigenvalues $\widetilde{\lambda}_i$ and $\{\widetilde{\lambda}_i^{1/2} e_i\}_{i \geq 1}$ is an orthonormal basis for \mathcal{H} . Furthermore, $\langle \widetilde{\mathfrak{C}} e_i, e_j \rangle = E(e_i(X)e_j(X)) = \langle e_i, e_j \rangle_{\mathcal{L}^2(P)} = \delta_{ij}$ and $\widetilde{\lambda}_1^{1/2} e_1, \widetilde{\lambda}_2^{1/2} e_2, \dots$ are the eigenvectors of $\widetilde{\mathfrak{C}}$ with corresponding eigenvalues $\widetilde{\lambda}_1, \widetilde{\lambda}_2, \dots$. Also notice that for all $y \in \mathcal{X}$, $(T_k \mathbf{1})(y) = \int k(y, x) dP = \langle k(y, \cdot), \mathbf{m} \rangle = \mathbf{m}(y)$ whenever the Bochner integral $\int k(x, \cdot) dP$ is well defined. Since $T_k \mathbf{1}$ and \mathbf{m} are real valued functions defined on \mathcal{X} that are equal for all $y \in \mathcal{X}$ it follows that $T_k \mathbf{1} = \mathbf{m}$.

The covariance operator as described above is giving us the second moments but not the covariance itself. The centered version $\widetilde{\mathfrak{C}}_c = \widetilde{\mathfrak{C}} - \mathbf{m} \widehat{\otimes} \mathbf{m}$ gives us the covariance itself, i.e. $E((f(X) - E(f(X)))(g(X) - E(g(X)))) = \langle \widetilde{\mathfrak{C}}_c f, g \rangle$ for any $f, g \in \mathcal{H}$. This operator is also self-adjoint under suitable conditions on the kernel and has a spectral decomposition.

Direct sum. Another construction that we need is the *direct sum of two Hilbert spaces* \mathcal{H}_1 and \mathcal{H}_2 . The direct sum $\mathcal{H}_1 \oplus \mathcal{H}_2$ is the Cartesian product $\{(g, h) : g \in \mathcal{H}_1, h \in \mathcal{H}_2\}$ equipped with the inner product $\langle (g_1, h_1), (g_2, h_2) \rangle_\oplus = \langle g_1, g_2 \rangle_1 + \langle h_1, h_2 \rangle_2$ [RS72, p. 40, Example 5]. We do not assume here that $\mathcal{H}_1 \cap \mathcal{H}_2 = \{0\}$.

Acknowledgments

I would like to thank the referee for their careful reading of the manuscript, for identifying typos and minor errors, and for suggesting improvements that enhanced the readability of the paper. The author is also grateful to many colleagues for helpful discussions.

2. Approximating convex sets and locating \mathbf{m} and \mathbf{m}_n

We start this section with a discussion of a simple approach for approximating convex sets using ε -nets. We will find that such an approximation is of very limited use only which motivates the remainder of the paper. In this remainder we analyze a stochastic approach at length where we consider the random convex set which is

induced by the sample. In detail, we control the difference between the empirical convex set C_n corresponding to the sample and its population limit C using VC-theory and Rademacher complexities in Section 2.2. Such tools are not necessary for the finite dimensional setting but the question of convergence of the empirical convex set to its population limit can easily be developed for the infinite dimensional setting. In particular, the approach based on Rademacher complexities applies directly to infinite dimensional RKHSs. In Section 2.3 we study the width of the convex set C . We link here lower bounds on the width of C to how well constant functions can be approximated within the unit ball of the RKHS. Building up on these sections we study how deep \mathbf{m} lies within C in Section 2.4. We also look in this section at an approach based on covariance operators which adapts automatically to the support of the unknown measure. Finally, in Section 2.5 we translate these findings to \mathbf{m}_n and we provide our main theorems in this section which give high probability bounds on the size of balls within the empirical convex set which are centered at \mathbf{m}_n .

2.1. Approximation based on ε -nets

Let \mathcal{H} be an RKHS of real-valued functions acting on $\mathcal{X} = [0, 1]^d$ with kernel function k being bounded by 1. Furthermore, let $\phi : \mathcal{X} \rightarrow \mathcal{H}$ be the map $\phi(x) = k(x, \cdot)$ and $\mathbf{m}_n = \frac{1}{n} \sum_{i=1}^n \phi(x_i)$ for certain points $x_1, \dots, x_n \in X$. For $\varepsilon > 0$ there exists an ε -net for $[0, 1]^d$ that consists of $N_{\varepsilon, d} = \lceil d^{d/2}/\varepsilon^d \rceil$ many closed balls that are centered at points $y_1, \dots, y_{N_{\varepsilon, d}}$ in $[0, 1]^d$. This ε -cover of $[0, 1]^d$ gives rise to a $c\varepsilon^\alpha$ -cover of $\phi[\mathcal{X}] = S$ if ϕ is α -Hölder continuous with Lipschitz constant c . Let $s_i = \phi(x_i)$ for all $i \leq n$ and s'_i the closest point to s_i in $\phi[\{y_1, \dots, y_{N_{\varepsilon, d}}\}]$. Then the approximation $\mathbf{m}'_n = \frac{1}{n} \sum_{i=1}^n s'_i$ of \mathbf{m}_n , which can be written as a sum over at most $N_{\varepsilon, d}$ many terms, achieves an approximation error of

$$\|\mathbf{m}_n - \mathbf{m}'_n\| \leq \frac{1}{n} \sum_{i=1}^n \|s_i - s'_i\| \leq c\varepsilon^\alpha.$$

If we want to achieve an approximation error of at most $n^{-1/2}$ then we need to include $\lceil d^{d/2}(c^2 n)^{d/(2\alpha)} \rceil$ many balls in the cover. If $c \geq 1$ then we can only represent \mathbf{m}_n with less than n -points if $d = 1$ and $\alpha > 1/2$. The Lipschitz constant c is here only of limited help if we choose our kernel independent of n .

We can also observe that a fine cover is necessary for good approximation if we do not impose assumptions on the measure and on \mathbf{m} . For instance, consider again $\mathcal{X} = [0, 1]^d$ and a kernel k such that $k(x, x) = 1$ for all $x \in \mathcal{X}$ and such that $1 - k(x, y) \leq c\|x - y\|$ for some constant $c > 0$ and any $x, y \in \mathcal{X}$. Furthermore, assume that we have a cover centered at l^d points x_1, \dots, x_{l^d} then there exists a point x_0 with $\min_{i \leq l^d} \|x_0 - x_i\| \geq 1/2l$. If we consider now the measure with unit mass on x_0 , i.e. $\mathbf{m} = k(x_0, \cdot)$, then the error, when approximating the expected value of the norm one function $h = k(x_0, \cdot)$, is

$$|\langle \mathbf{m}, h \rangle - \langle \mathbf{m}_n, h \rangle| = \|\mathbf{m}\|^2 - \sum_{i=1}^{l^d} \alpha_i k(x_i, x_0) \geq \frac{c}{2l}.$$

Hence, to attain an approximation error of order $n^{-1/2}$ we need a cover consisting of at least $n^{d/2}$ many points.

2.2. Empirical convex sets

In the following, let \mathcal{H} be a separable RKHS and let $C_n = \text{ch}\{\phi(X_i) : i \leq n\}$ be the set valued random variable determined by X_1, \dots, X_n . The variable C_n attains values in the closed convex subsets $\mathbf{C}(\mathcal{H})$ of \mathcal{H} . There exists various natural topologies on $\mathbf{C}(\mathcal{H})$ (see [Bee93]). We equip $\mathbf{C}(\mathcal{H})$ with the Vietoris topology and the corresponding Borel–Effron σ -algebra. The random variable C_n is then well defined as a measurable map from Ω to $\mathbf{C}(\mathcal{H})$. The random variable C_n tends to $C = \text{cch}\{\phi(x) : x \in \mathcal{X}\}$ as n tends to infinity. We aim to quantify how similar C_n is to C . We do so by framing the question of convergence in the context of empirical process theory. In the following discussion we assume that \mathcal{X} is *compact*, \mathcal{H} is *finite dimensional with dimension d* , and the corresponding *kernel function k is continuous*. In particular, $\|k\|_\infty^{1/2} =: b$ is finite.

Observe that we can reduce the question of convergence of C_n to C to the question of how fast the projection of C_n on some direction $u \in \mathcal{H}$, $\|u\| = 1$, converges to the projection of C on u . More specifically, if we can control the convergence uniformly over all such u then we have control of the convergence of C_n to C . Furthermore, since C_n and C are convex we only need to control the end points of the projections; these points correspond to projections of extremes of C_n and C onto $\text{span}\{u\}$. With this aim in mind, let us introduce the functions $f_{u,c}(x) = \chi\{u(x) \leq c\}$, $f_{u,c} : \mathcal{X} \rightarrow \mathbb{R}$, for $u \in \mathcal{H}$, $\|u\| = 1$, and with c going through the interval $\{\langle u, h \rangle : h \in C\} = \text{ch}\{u(x) : x \in \mathcal{X}\}$ or a superset of this interval.

The importance of the functions $f_{u,c}$ is that $Pf_{u,c} > 0$ if, and only if, there is an element $h \in C$ such that $\langle h, u \rangle \leq c$ (given that there is non-zero mass on that element or the mass of all elements whose projection falls below c is strictly greater than zero). For instance, if C contains the origin then we could vary negative c 's to explore the extension of the projection of C in direction u . Since the extremes of C are a subset of $S := \{\phi(x) : x \in \mathcal{X}\}$ and the probability measure is concentrated on S it is sufficient to work with elements in S instead of all of C . This setup is depicted in parts (i)–(ii) of Figure 2.1.

The situation is similar for the empirical convex set. The empirical convex set will contain an element which lies c away from the origin in direction u if, and only if, $P_n(u(\tilde{X}) \leq c) > 0$, with \tilde{X} being a random variable with law P_n (see the preliminaries in Section 1.9). Notice that the condition $P_n(u(\tilde{X}) \leq c) > 0$ is equivalent to $\min_{i \leq n} u(X_i) \leq c$.

To be able make use of this approach to quantify the difference between C and C_n we need to control the convergence of $P_n f_{u,c}$ to $P f_{u,c}$ simultaneously over all these $f_{u,c}$. Two natural ways to control this difference are VC-theory and Rademacher complexities. We discuss both of these below. Note that the resulting concentration bounds are only useful if the probability $P f_{u,c}$ and related terms are not too small. In the following examples, and in most parts of the paper, we deal with this difficulty by assuming lower bounds on the density of the data generating random variables.

This guarantees that $Pf_{u,c}$ cannot get too small. If one wishes to avoid this assumption, one has to find a way to deal with areas of the input space that have low probability to contain samples. One option to deal with these areas might be to use a threshold θ_n on the values of the density. One could then change the original density p to a truncated version $p_{\theta_n} \approx p \times \chi\{p \geq \theta_n\}$ and adapt the threshold θ_n with sample size n to let it go to zero in a controlled way. The convex set corresponding to p_{θ_n} can then be analyzed with our tools and one only has to control the difference between p and p_{θ_n} , and the corresponding convex sets. While this seems like a natural solution, we do not explore this path in this paper and we focus fully on the case of densities that are bounded away from zero.

2.2.1. VC-theory

Since we are working here with finite dimensional RKHSs it is rather easy to use VC-theory. In detail, whenever $u \in \mathcal{H}$, $\|u\| = 1$, then $|u(x)| \leq b$ and we can use $[-b, b]$ as the interval over which we vary c . Hence, let

$$\mathcal{F} = \{f_{u,c} : u \in \mathcal{H}, \|u\| = 1, -b \leq c \leq b\}.$$

We want to show that \mathcal{F} is a VC-subgraph class of functions. In fact, it is convenient to work with a countable dense subset of \mathcal{F} to sidestep measure theoretic complications relating to the empirical process. To this end, let $\tilde{\mathcal{H}}$ be a countable dense subset of \mathcal{H} such that $\tilde{\mathcal{H}} \cap \{u : \|u\| = 1, u \in \mathcal{H}\}$ lies dense in $\{u : \|u\| = 1, u \in \mathcal{H}\}$ and define the countable set

$$(2.1) \quad \tilde{\mathcal{F}} = \{f_{u,c} : u \in \tilde{\mathcal{H}}, \|u\| = 1, c \in (\mathbb{Q} \cap [-b, b]) \cup \{-b, b\}\} \subset \mathcal{F}.$$

The family \mathcal{F} is a VC-subgraph class and its VC-dimension is upper bounded by $d + 1$: consider the family of function $\mathcal{G} = \text{span}(\mathcal{H} \cup \{c\mathbf{1} : c \in \mathbb{R}\})$. The dimension of \mathcal{G} is at most $d + 1$ and $c - u(x) \in \mathcal{G}$ for every $u \in \mathcal{H}$, $-b \leq c \leq b$. Applying [Dud14, Theorem 4.6] shows that the VC dimension of $\text{Pos}(\mathcal{G}) = \{\text{pos}(g) : g \in \mathcal{G}\}$, where $\text{pos}(g) = \{x : g(x) \geq 0\}$, is at most $d + 1$. Furthermore, the family \mathcal{G}' of sets of the form $\{(x, t) : x \in \text{pos}(g), t \leq 1\}$, $g \in \mathcal{G}$, has the same VC-dimension. But \mathcal{G}' is a family of subgraphs that contains all the subgraphs of functions in \mathcal{F} and the claim follows. Since $\tilde{\mathcal{F}} \subset \mathcal{F}$ it also follows that $\tilde{\mathcal{F}}$ is a VC-subgraph class with VC-dimension at most $d + 1$.

The family $\tilde{\mathcal{F}}$ has the measurable envelope $\chi\mathcal{X}$ and, due to [GN16, Theorem 3.6.9], its covering numbers can be bounded by

$$N(\tilde{\mathcal{F}}, \mathcal{L}^2(Q), \varepsilon) \leq 4(8/\varepsilon^2)^{d+2} \vee \tilde{c},$$

where \tilde{c} can be chosen as $\max\{m \in \mathbb{N}_+ : \log m \geq m^{1/(d+1)(d+2)}\}$ and whenever Q is a probability measure on \mathcal{X} . Be aware that the ν -index as defined in [GN16] is equal to one plus the VC-dimension when using the definition of [Dud14] for the VC-dimension.

Now, applying Hölder's inequality,

$$J(\delta) = \int_0^\delta \sup_Q \sqrt{\log 2N(\tilde{\mathcal{F}}, \mathcal{L}^2(Q), \varepsilon)} d\varepsilon \leq \delta(\log(2\tilde{c}) \vee (1 + 2(d + 2)))^{1/2}.$$

In particular, $J(1) \leq \sqrt{\log 2\tilde{c}} \vee \sqrt{1 + 2(d + 2)}$. By [GN16, Remark 3.5.5 and Theorem 3.5.4] we can conclude that

$$E\left(\sup_{f \in \tilde{\mathcal{F}}} |P_n f - P f|\right) \leq 12J(1)n^{-1/2}.$$

We use now Bousquet’s version of Talagrand’s inequality to move to a high probability bound (e.g. [GN16, Theorem 3.3.9]). For simplicity, we will denote the supremum over u, c , such that $f_{u,c} \in \tilde{\mathcal{F}}$, by $\sup_{u,c}$ in the following. Let

$$S_n = \sup_{u,c} \left| \sum_{i=1}^n (f_{u,c}(X_i) - P f_{u,c}) \right| = n \sup_{u,c} |P_n f_{u,c} - P f_{u,c}|.$$

Observe that $\|P f_{u,c} - f_{u,c}\|_\infty \leq 1$ and $ES_n = nE(\sup_{u,c} |P_n f_{u,c} - P f_{u,c}|)$. Applying Talagrand’s inequality yields

$$e^{-x} \geq \Pr\left(\max_{j \leq n} S_j \geq ES_n + \sqrt{2x(2ES_n + n)} + x/3\right)$$

for all $x \geq 0$. Letting $p = \exp(-x)$ leads to the following lemma.

LEMMA 2.1. — *Let \mathcal{H} be a finite dimensional RKHS with continuous kernel function k defined on a compact set \mathcal{X} . Assume that $\|k\|_\infty^{1/2} = b < \infty$. Let $\tilde{\mathcal{F}}$ be the space associated to \mathcal{H} through Eq. (2.1). With probability at least $1 - p$,*

$$\sup_{f \in \tilde{\mathcal{F}}} |P_n f - P f| \leq 12J(1)n^{-1/2} + n^{-1/2} \sqrt{2 \log(1/p)(24J(1)n^{-1/2} + 1)} + \log(1/p)/3n.$$

2.2.2. Rademacher complexities

As is usually the case with metric entropy based bounds, the constants are loose and n needs to be large to gain useful results. Tighter bounds can often be attained by using *Rademacher complexities* (see [BM03, GN16]). While the resulting bounds are generally tighter it is not possible to work directly with the indicator functions $f_{u,x}$ but we need a continuous approximation of these. Also, in the Rademacher approach that we develop it is beneficial to center the functions $h \in C$ by moving to $C_c = \{h - \mathbf{m} : h \in C\}$. In the following, let $F = \{(u, c) : u \in \mathcal{H}, \|u\| = 1, -b \leq c \leq b\}$ and

$$(2.2) \quad \tilde{F} = \{(u, c) : u \in \tilde{\mathcal{H}}, \|u\| = 1, c \in ([-b, b] \cap \mathbb{Q}) \cup \{-b, b\}\}.$$

Furthermore, consider the function $\psi_\gamma : \mathbb{R} \rightarrow \mathbb{R}$, with $\gamma > 0$, defined by

$$\psi_\gamma(x) = \begin{cases} 1 & x \leq -\gamma, \\ -x/\gamma & -\gamma < x < 0, \\ 0 & 0 \leq x. \end{cases}$$

Then $f_{u,c}(h) = \chi\{\langle u, h \rangle \leq c\} \geq \psi_\gamma(\langle u, h \rangle - c)$ for any $u, h \in \mathcal{H}$ and $-b \leq c \leq b$. The function ψ_γ is depicted in Figure 2.1 (iii). Importantly, $\psi(0) = 0$ and $|\psi_\gamma(x) - \psi_\gamma(y)| \leq |x - y|/\gamma$, that is $\gamma\psi_\gamma(\cdot)$ is a contraction vanishing at zero (see [GN16, Theorem 3.2.1] or [LT91, Theorem 4.12]).




Figure 2.1. (i) The figure show C as a subset of \mathcal{H} . The diagonal (blue) line is the $\text{span}\{u\}$ for some function $u \in \mathcal{H}, \|u\| = 1$, The short lines connecting this line to the ellipse indicate the projection of C on $\text{span}\{h\}$. In particular, the distance between the two short lines is $\text{width}_u(C)$. The long line which is orthogonal to $\text{span}\{u\}$ (red) indicates a threshold; the interest is here if C extends past this threshold. (ii) The question if C extends past the threshold is rephrased in this figure by focusing on $\text{span}\{u\}$ and considering the probability that values $u(x)$ are attained that lie beyond the threshold. In this figure, we assume for simplicity that the measure on C induces a density function $p(y)$ through the projection on $\text{span}\{u\}$, where y goes over the range of u . The threshold is in this figure set to $-c$ and C extends past the threshold if the density function is non-zero to the left of $-c$. (iii) To link this construction to the empirical measure we use the function ψ_γ whose graph is plotted in this figure against $u(x)$. The motivation is here to approximate the indicator function corresponding to the event $u(X) \leq -c$ from below by a continuous function. The parameter γ controls the approximation and for $\gamma \rightarrow 0$ the function ψ_γ converges to the indicator function.

We have that $P_n f_{u,c+\langle u, \mathbf{m} \rangle} \geq P_n \psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c)$. The proof of [GN16, Theorem 3.4.5] gives us a high probability lower bound on the latter term (in the notation of the book, combine $S_n < ES_n + \sqrt{2x/n}$ with $ES_n \leq 2E\tilde{S}_n$). In detail, with probability $1 - p$, simultaneously for all $u \in \tilde{\mathcal{H}}, \|u\| = 1$, and $c \in ([-b, b] \cap \mathbb{Q}) \cup \{-b, b\}$, we have the following lower bound on $P_n \psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c)$,

$$P_n \psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c) - 2E \left(\sup_{(u',c') \in \tilde{F}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i \psi_\gamma(\langle u', \phi(X_i) - \mathbf{m} \rangle - c') \right| \right) - \sqrt{\frac{2 \log(2/p)}{n}},$$

where ϵ_i are i.i.d. Rademacher variables that are independent of X_1, \dots, X_n . Because $\gamma \psi_\gamma$ is a contraction vanishing at zero

$$E \left(\sup_{(u',c') \in \tilde{F}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i \psi_\gamma(\langle u', \phi(X_i) - \mathbf{m} \rangle - c') \right| \right) \leq \frac{2}{\gamma} E \left(\sup_{(u',c') \in \tilde{F}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i (\langle u', \phi(X_i) - \mathbf{m} \rangle - c') \right| \right).$$

Applying [BM03, Theorem 12(7) and Lemma 22] and using that the Rademacher complexity for the constant functions $\langle u', \mathbf{m} \rangle + c'$, $(u', c') \in \tilde{F}$, where $|\langle u', \mathbf{m} \rangle + c'| \leq b + |c'| \leq 2b$, is upper bounded by $4bn^{-1/2}$,

$$\begin{aligned} E \left(\sup_{(u', c') \in \tilde{F}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i (\langle u', X_i \rangle - c') \right| \right) \\ \leq 4bn^{-1/2} + E \left((2/n) \left(\sum_{i=1}^n k(X_i, X_i) \right)^{1/2} \right) \leq 6bn^{-1/2}. \end{aligned}$$

This implies the following lemma.

LEMMA 2.2. — *Let \mathcal{H} be an RKHS with continuous kernel function k defined on a compact set \mathcal{X} . Assume that $\|k\|_\infty^{1/2} = b < \infty$. Let \tilde{F} be the set associated to \mathcal{H} through Eq. (2.2). With probability at least $1 - p$, simultaneously for all $(u, c) \in \tilde{F}$,*

$$P_n \psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c) \geq P \psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c) - \left(\sqrt{2 \log(2/p)} + 24b/\gamma \right) n^{-1/2}.$$

The VC and Rademacher bounds allow us to control the size of the empirical convex set in terms of $Pf_{u,c}$ and $P\psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c)$. In either case we need to get a handle on P to move further. In particular, we need to understand how P concentrates around the extremes of C . We are now looking at a few examples to get a better understanding of how P concentrates and what this implies for the convergence of the empirical convex set to C . Of major importance is how smooth $\phi : \mathcal{X} \rightarrow \mathcal{H}$ is and how the distribution of X_1, \dots, X_n on \mathcal{X} looks like. We start with a couple of simple examples and discuss links to stochastic geometry before addressing typical settings that one faces in practice.

2.2.3. Examples

Example 1 (Unit circle). — Consider the unit circle in \mathbb{R}^2 with the uniform distribution on it. What can we say about the interior of C_n as a function of n ? In particular, what can be said about the size of C_n in direction $u \in \mathbb{R}^2$, $\|u\| = 1$? Due to the symmetry of the unit sphere and because the uniform distribution is used it is sufficient to consider the vector $u = (1, 0)^\top$. The probability that a sample point, when we sample just once, lies to the right of cu , with $c \in [-1, 1]$, is $(1/2\pi) \int_0^{2\pi} \chi\{\langle u, (\cos \theta, \sin \theta)^\top \rangle \geq c\} = \arccos(c)/\pi$, using here that \arccos is a monotonically decreasing function. Similarly, the probability that a sample point lies to the left of cu is $1 - \arccos(c)/\pi$. Furthermore, if we draw n independent samples then the probability to see at least one sample point to the right of cu is $1 - (1 - \arccos(c)/\pi)^n$ and that at least one sample point lies to the left of cu is $1 - (\arccos(c)/\pi)^n$. Moving on to the distribution of the length of the interval, which corresponds to the projection of C_n onto u , that is the distribution of $\text{width}(u, C_n)$, we can observe that $\text{width}(u, C_n)$ attains values in $[0, 2]$ and that $\text{width}(u, C_n) = \max_i \cos \theta_i - \min_i \cos \theta_i$, where we denote with θ_i independent and uniformly distributed random variables on $[0, 2\pi)$. We could now try to calculate the distribution of $\text{width}(u, C_n)$ by controlling the maximum and minimum. Since we are interested in getting a better understanding of

the VC and Rademacher approach we use instead the uniform guarantees on $P_n f_{u,c}$. Let \mathcal{X} be the unit circle and let the kernel function be $k(x, y) = \langle x, y \rangle_{\mathbb{R}^2}$. This way \mathcal{H} becomes the dual space $(\mathbb{R}^2)'$ of \mathbb{R}^2 : recall that a basis of $(\mathbb{R}^2)'$ is given by $\langle e_1, \cdot \rangle_{\mathbb{R}^2}, \langle e_2, \cdot \rangle_{\mathbb{R}^2}$, where e_1, e_2 is the standard basis in \mathbb{R}^2 , and for any $i, j \in \{1, 2\}$, $\langle \langle e_i, \cdot \rangle_{\mathbb{R}^2}, \langle e_j, \cdot \rangle_{\mathbb{R}^2} \rangle_{(\mathbb{R}^2)'} = \langle e_i, e_j \rangle_{\mathbb{R}^2}$. Since e_1, e_2 lie in \mathcal{X} it holds for any $i, j \in \{1, 2\}$, $\langle \langle e_i, \cdot \rangle_{\mathbb{R}^2}, \langle e_j, \cdot \rangle_{\mathbb{R}^2} \rangle_{(\mathbb{R}^2)'} = \langle k(e_i, \cdot), k(e_j, \cdot) \rangle$ and the claim follows.

Associate to $u \in \mathbb{R}^2$ the function $\tilde{u} \in \mathcal{H}$ given by $\tilde{u}(x) = \langle u, x \rangle_{\mathbb{R}^2}$. Let $\tilde{\mathcal{H}}$ be a countably dense subset of \mathcal{H} such that $\{u : u \in \tilde{\mathcal{H}}, \|u\| = 1\}$ lies dense in the unit sphere of \mathcal{H} . Define $\mathcal{F} = \{f_{\tilde{u},c} : u \in \tilde{\mathcal{H}}, \|u\| = 1, -1 \leq c \leq 1\}$ and $\tilde{\mathcal{F}} = \{f_{\tilde{u},c} : u \in \tilde{\mathcal{H}}, \|u\| = 1, -1 \leq c \leq 1\}$, where $f_{\tilde{u},c}(x) = \chi\{\tilde{u}(x) \leq c\} = \chi\{\langle u, x \rangle \leq c\}$. The family of functions $\tilde{\mathcal{F}}$ is a VC-subgraph class and on an event of probability at least p it holds simultaneously for all $f_{\tilde{u},c} \in \tilde{\mathcal{F}}$ that

$$P_n f_{\tilde{u},c} \geq P f_{\tilde{u},c} - n^{-1/2} \xi_n = 1 - \arccos(c)/\pi - n^{-1/2} \xi_n,$$

where $\xi_n = 12J(1) + \sqrt{2 \log(1/p)(24J(1)n^{-1/2} + 1)} + \log(1/p)n^{-1/2}/3$. We use here that the uniform distribution on the unit circle is invariant under rotations, i.e. for a given u let A be the rotation matrix for which $Au = (1, 0)^\top$. Then,

$$P f_{\tilde{u},c} = \frac{1}{2\pi} \int_0^{2\pi} \chi\{\langle Au, A(\cos(\theta), \sin(\theta))^\top \rangle \leq c\} = \frac{1}{\pi} \int_0^\pi \chi\{\cos(\theta) \leq c\}.$$

In other words, on an event of probability p , whenever n is such that $n^{-1/2} \xi_n < 1/2$, and for any $c > \cos((1 - n^{-1/2} \xi_n)\pi)$, there will be a sample point which has an inner product with u which is smaller than c . To be exact, let $c_0 < 0$ be a real number strictly larger than $\cos((1 - n^{-1/2} \xi_n)\pi)$ and let the above event be denoted by B . It holds that $P(B) \geq p$ and for any $\omega \in B$, $\min_{i \leq n} \langle u, X_i(\omega) \rangle \leq c_0$. In fact, the VC-argument shows that B can be chosen such that $P(B) \geq p$ and for all $\omega \in B$,

$$\sup_{u \in \mathcal{H}, \|u\|=1} \min_{i \leq n} \langle u, X_i(\omega) \rangle \leq c_0.$$

From this we can infer that a ball centered at the origin and of radius c_0 is contained in the empirical convex set $C_n(\omega)$, whenever $\omega \in B$: consider without loss of generality the vector $v = (c_0, 0)^\top$ and an element $\omega \in B$. There exist elements X_j such that $X_j(\omega)$ lies on the unit circle and $\langle -v, X_j(\omega) \rangle \leq c_0^2$, that is $\langle v, X_j(\omega) \rangle \geq \|v\|^2$. Let $X_i(\omega)$ be such an element which also attains the maximum of the map $j \mapsto \langle v, X_j(\omega) \rangle$.

Assume that $X_i(\omega)$ does not lie in $\text{span } v$ and that $X_i(\omega)$ lies north of $\text{span } v$, i.e. $\langle X_i(\omega), (0, 1)^\top \rangle > 0$. Consider the lines between $X_i(\omega)$ and the elements $X_j(\omega)$, $j \leq n, j \neq i$. There will be an index $j_0 \leq n, j_0 \neq i$, such that the line between $X_i(\omega)$ and $X_{j_0}(\omega)$ intersects with $\text{span } v$. Consider the vector $w = (0, -c_0)^\top$. There will be a sample point $X_{j_1}(\omega)$ such that $\langle w, X_{j_1}(\omega) \rangle \geq \|w\|^2$ and the line between $X_i(\omega)$ and $X_{j_1}(\omega)$ crosses $\text{span } v$. Order the samples according to how large the inner product between the point of intersection of the line between the sample and $X_i(\omega)$ and v is. Let $X_{j_2}(\omega)$ be the maximum in this ordering. Assume that $\langle X_{j_2}(\omega), v \rangle < \|v\|^2$, that is the intersection lies to the left of v . Let \tilde{v} be the point on the circle with radius c_0 for which the line between $X_i(\omega)$ and $X_{j_2}(\omega)$ is tangent and which lies to the right of the line. There is now a point $X_{j_3}(\omega)$ on the sphere such that $\langle X_{j_3}(\omega), \tilde{v} \rangle \geq \|\tilde{v}\|^2$. The point $X_{j_3}(\omega)$ cannot lie north of v since this would contradict the maximality

of $X_i(\omega)$. However, if $X_{j_3}(\omega)$ lies south of v then the line between $X_{j_3}(\omega)$ and $X_i(\omega)$ crosses span v further to the right than the line between $X_{j_2}(\omega)$ and $X_i(\omega)$ which contradicts the maximality of $X_{j_2}(\omega)$.

Hence, we have either two points to the right of v , one on the north side and one on the south side of the sphere, or the point $(1, 0)^\top$ is contained in the sample. By the same argument, either $(-1, 0)^\top$ is contained in the sample or there are two points left of $-v$, one on the north side and one on the south side. The convex hull of these points is a subset of $C_n(\omega)$ and contains v .

To provide a concrete example, let $p = 0.9$ and observe that the \tilde{c} which appears in the bound of $J(\delta)$ can be chosen as 10^{21} . Then $J(1) \leq 8 \vee 3 = 8$ and $\xi_n \leq 96 + \sqrt{2 \log(10)(192n^{-1/2} + 1)} + n^{-1/2} \log(10)/3$. Hence, a ball of radius 0.2 exists around the origin inside the empirical convex set with probability p for n being about 52000 or larger.

As expected n needs to be large to guarantee the existence of the ball or radius 0.2. Using *Rademacher complexities* we can attain significantly tighter bounds in this setting. Building up on our discussion and using $\mathbf{m} = 0$ we can see that

$$P_n \psi_\gamma(\langle u, \cdot \rangle - c) \geq P \psi_\gamma(\langle u, \cdot \rangle - c) - \left(\sqrt{2 \log(2/p)} + 12/\gamma \right) n^{-1/2}.$$

Finally, by using a rotation of u and with $c_\gamma = (c - \gamma) \vee -1$ it follows that

$$\begin{aligned} & P \psi_\gamma(\langle u, \cdot \rangle - c) \\ &= \frac{1}{\pi} \int_0^\pi \psi_\gamma(\cos(\theta) - c) \\ &= 1 - \frac{\arccos(c_\gamma)}{\pi} + \frac{1}{\pi \gamma} \int_{\arccos(c)}^{\arccos(c_\gamma)} (c - \cos(\theta)) \\ &= 1 - \frac{\arccos(c_\gamma)}{\pi} (1 - c/\gamma) - \frac{c \arccos(c)}{\pi \gamma} + \frac{1}{\pi \gamma} \left(\sqrt{1 - c^2} - \sqrt{1 - c_\gamma^2} \right). \end{aligned}$$

For instance, with $\gamma = 1$ and $c < 0$ this leads to

$$P_n \psi_\gamma(\langle u, \cdot \rangle - c) \geq c(1 - \arccos(c)/\pi) + \frac{\sqrt{1 - c^2}}{\pi} - \left(\sqrt{2 \log(2/p)} + 12 \right) n^{-1/2}.$$

The bound guarantees in this case the existence of a ball of radius 0.2 around the origin within the empirical convex set with probability at least 0.9 when n is about 5000, a 10-fold improvement in the constant over the VC-bound. While the bound is significantly better it does not come close to capture the right magnitude: even a number as small as $n = 10$ suffices in experiments for the empirical convex set to contain a ball of radius 0.2 with high probability.

In this example, we assumed that the distribution is uniform on the unit circle. This assumption is not crucial but helps with the computations. That said, it is important for our argument that the density of the distribution does not approach zero. Otherwise, $P \psi_\gamma(\langle u, \cdot \rangle - c)$ could be arbitrary small and we could not guarantee a positive value for its empirical counterpart.

Example 2 (Polytopes with finitely many extremes). — Let us consider next a simple polytope. Let $\mathcal{H} = \mathbb{R}^d$ with the usual inner product ($\langle x, y \rangle = x^\top y$) and $k(x, y) = \langle x, y \rangle$. Furthermore, consider $C = \text{ch}\{x_i : i \leq m\}$ with $x_1, \dots, x_m \in \mathbb{R}^d$

and such that the random variable X attains values in $\{x_1, \dots, x_m\}$ and $\Pr(X = x_i) \geq \alpha > 0$ for all $i \leq m$. Then for all $u \in \mathcal{H}$, $\|u\| = 1$, $|c| \leq \|k\|_\infty^{1/2}$ either $Pf_{u,c} = 0$ or $Pf_{u,c} \geq \alpha$. Hence, we have that $P_n f_{u,c} > 0$ with probability at least $1 - e^{-\sqrt{n}}$ whenever

$$n \geq \left(12J(1) + \sqrt{2(24J(1) + 1)} + 1/3\right)^2 / \alpha^2.$$

In other words, for n that large the empirical convex set equals C on an event of probability at least $1 - e^{-\sqrt{n}}$.

If each of the x_i is an extreme then we can compare this probability to the probability that in n independent trials all m extremes are drawn: the probability that element i is not drawn in n independent trials is $1 - \Pr(X = x_i)^n$ and the probability that at least one element i is not drawn is upper bounded by

$$\Pr\left(\bigcup_{i \leq m} \bigcap_{j \leq n} \{X_j \neq x_i\}\right) \leq \sum_{i \leq m} (1 - \Pr(X = x_i))^n \leq m(1 - \alpha)^n = m \exp(-\beta n).$$

where $\beta = -\log(1 - \alpha)$. In other words, instead of $1 - e^{-\sqrt{n}}$ we get a probability of $1 - m \exp(-\beta n)$ that the empirical convex set matches the convex set C .

Consider now the special case of the d -dimensional simplex $\text{ch} S$, with $S = \{0, e_1, \dots, e_d\}$ and e_1, \dots, e_d being an orthonormal basis in \mathbb{R}^d . Furthermore, assume that each $x \in S$ has probability $1/(d+1)$ to be sampled. The interior of the empirical convex set C_n is empty unless all points have been sampled. Hence, in this example either $\text{int} C_n = \emptyset$ or $C_n = C$ and the interior of n does not grow slowly in size as n increases but changes abruptly.

As a final example consider a *rhombus* given by $C = \text{ch}\{e_1, -e_1, re_2, -re_2\}$ where e_1, e_2 are orthonormal vectors in \mathbb{R}^2 and $r \in (0, 1)$. Furthermore, let X be uniformly distributed on the boundary of C . We can again consider the functions $f_{u,c}$ to measure the interior of the empirical convex set. However, in contrast to the unit circle the measure in direction u that lies c apart from the origin is not the same for all u but depends strongly on the direction. For instance, for direction $-e_1$ and $c \in (0, 1)$ it holds that $Pf_{-\tilde{e}_1, -c} = 2p(1 - c)\sqrt{1 + r^2}$, where p denotes here the density of the uniform distribution on the boundary, while for $-e_2$ we get $Pf_{-\tilde{e}_2, -c} = 2p(1 - c/r)\sqrt{1 + r^2}$, for $c \in (0, r)$. In particular, for $c = 0.9$, $c' = 0.9r$ the probabilities $Pf_{-\tilde{e}_1, -c}$ and $Pf_{-\tilde{e}_2, -c'}$ are equal and the probability $Pf_{-\tilde{e}_1, -c}$, which is spread out over an interval of length 0.1 in direction e_1 , is contained in an interval of length $0.1r$ in direction e_2 irrespective of how small r is.

As in Example 1, the uniform assumption helps here with the computations but is not crucial for the argument. Since we are in a discrete setting we do not have the same problem as in Example 1: there cannot be a sequence of non-zero probabilities on our m points x_1, \dots, x_m that converges to zero. On the other hand, if there is a probability of exactly zero for some points, for example for attaining value x_1 , then we can simply ignore these points and carry out the same analysis for the points with non-zero probability.

Example 3 (Image of a Lipschitz-continuous kernel function). — Let us go back to the setting that we discussed at the beginning of the section. In detail, let k be a *continuous kernel function on compact set \mathcal{X}* that is upper bounded by b .

Furthermore, let us assume that the corresponding feature map $\phi(x) = k(x, \cdot)$ is L -Lipschitz continuous with Lipschitz constant $L > 0$ and the law of X_1, \dots, X_n has a density on \mathcal{X} which is lower bounded by $b' > 0$. We are now aiming to quantify the extension of the convex set in a direction u after centering the convex set around \mathbf{m} . In detail, for $u \in \mathcal{H}$, $\|u\| = 1$, let $x_u \in \mathcal{X}$ be a point at which $c_u^* := \langle u, \phi(x_u) - \mathbf{m} \rangle = \min_{x \in \mathcal{X}} \langle u, \phi(x) - \mathbf{m} \rangle$. As before, for $c \in \mathbb{R}$, let $f_{u,c} = \chi\{\langle u, \phi(\cdot) - \mathbf{m} \rangle \leq c\}$ and observe that $|\langle u, \phi(x_u) - \phi(x) \rangle| \leq L\|x_u - x\|$. Therefore, with $r_u = (c - c_u^*)/L$ and whenever $c > c_u^*$,

$$Pf_{u,c} \geq \int_{\mathcal{X}} b' \times \chi\{\langle u, \phi(x) - \mathbf{m} \rangle \leq c\} dx \geq \int_{B(x_u, r_u) \cap \mathcal{X}} b' = b' \text{vol}(B(x_u, r_u) \cap \mathcal{X}).$$

For example, when $\mathcal{X} = [0, 1]$, u any element in \mathcal{H} with $\|u\| = 1$ and $c > c_u^*$ such that $x_u - (c - c_u^*)/L \geq 0$, it follows that $\text{vol}(B(x_u, r_u) \cap \mathcal{X}) \geq r_u$ and $Pf_{u,c} \geq b'r_u = b'(c - c_u^*)/L$.

This lower bound on $Pf_{u,c}$ can directly be combined with a metric entropy bound. If we want to use instead a Rademacher complexity bound then we have to apply ψ_γ to $\langle u, \phi(\cdot) - \mathbf{m} \rangle - c$. Under the above Lipschitz assumption for any u , $\|u\| = 1$, and whenever $c_u^* \leq c - \gamma$,

$$\begin{aligned} (2.3) \quad 1 = \psi_\gamma(\langle u, \phi(x) - \mathbf{m} \rangle - c) &\iff c - \gamma \geq \langle u, \phi(x) - \mathbf{m} \rangle \\ &\iff c - \gamma - c_u^* \geq \langle u, \phi(x) - \phi(x_u) \rangle \\ &\iff c - \gamma - c_u^* \geq L\|x - x_u\|. \end{aligned}$$

Also, $\psi_\gamma(\langle u, \phi(x) - \mathbf{m} \rangle - c)$ is strictly positive whenever $L\|x - x_u\| \leq c - c_u^*$.

Let $r_{u,1} = (c - \gamma - c_u^*)^+/L$. For $x \in B(x_u, r_{u,1})$ we have that $\psi_\gamma(\langle u, \phi(x) - \mathbf{m} \rangle - c) = 1$ which gives us right away the following lower bound

$$(2.4) \quad P\psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c) \geq b' \text{vol}(\mathcal{X} \cap B(x_u, r_{u,1})).$$

This bound can now be combined with the Rademacher complexity bounds. However, to say anything concrete about the size of the empirical convex set some knowledge of c_u^* is required. In the Section 2.3 we derive approaches to measure the width of C in any direction, then we derive lower bounds on c_u^* . We combine these bounds in Section 2.5 with the above bound.

The assumption that there is a non-zero lower bound on the density is important. We could weaken it slightly to allow for parts of the density to be exactly zero but we would face major obstacles if we would allow for the density to approach zero. In this example, this problem manifests itself in Eq. (2.4). We could not guarantee a lower bound on the right side which is strictly positive if we allowed the density to approach zero.

Example 4 (Data attaining values in a subset). — Working with the empirical convex set has the advantage that the set is adapted to the support of the distribution: if P has support S on \mathcal{X} then C_n converges to $\text{cch}\{k(x, \cdot) : x \in S\}$ (see [Fre01, Definition 411N] for the definition of support). Instead of proving this under a density and Lipschitz assumption, we are using here an assumption on the covariance operator.

A simple way to deal with S is to consider the space $\mathcal{H}_S = \{h \upharpoonright S : h \in \mathcal{H}\}$ which is again an RKHS with kernel $k_S = k \upharpoonright S \times S$. We discuss \mathcal{H}_S and how covariance operators are naturally adapted to S at length in Section 2.4.3. For the moment it is sufficient to note that for all $h \in \mathcal{H}$,

$$\langle \tilde{\mathfrak{C}}_c h, h \rangle = \langle \tilde{\mathfrak{C}}_c^S h \upharpoonright S, h \upharpoonright S \rangle_{\mathcal{H}_S},$$

where $\tilde{\mathfrak{C}}_c$ is the centered covariance operator and $\tilde{\mathfrak{C}}_c^S$ the corresponding operator for the RKHS \mathcal{H}_S . In Section 2.4.3 we show that \mathcal{H}_S is naturally linked to the affine subspace spanned by $k(x, \cdot), x \in S$, and that statements about the behavior of C_n can be derived by analyzing \mathcal{H}_S . In particular, the eigenfunctions of $\tilde{\mathfrak{C}}_c$ which have eigenvalue zero are almost surely constant on S and are all mapped to the same one dimensional subspace of \mathcal{H}_S . Important for the analysis later on are the eigenfunctions which have non-zero eigenvalues and are therefore not constant on S . Let $u \in \mathcal{H}_S$ be an eigenfunction of $\tilde{\mathfrak{C}}_c^S$ with eigenvalue $\bar{\lambda}$. Let ϕ_S be the feature map corresponding to k_S and \mathbf{m}_S the corresponding mean embedding. The function u has by definition norm one and for $\gamma > 0, c \in \mathbb{R}$, and X a random variable with law P ,

$$P\psi_\gamma(\langle u, \phi_S(\cdot) - \mathbf{m}_S \rangle_{\mathcal{H}_S} - c) \geq \Pr(-(u(X) - E(u(X))) \geq -\gamma - c).$$

In the following, let $Z = -(u(X) - E(u(X)))$ and write $Z = Z^+ - Z^-$ where $Z^+ = Z \times \chi\{Z \geq 0\}, Z^- = Z \times \chi\{Z \leq 0\}$. Since Z has mean zero we have that $E(Z^+) = E(Z^-)$. Whenever $\|k\|_\infty < \infty$ we also have that $E((Z^+)^2) \leq \|k\|_\infty^{1/2} E(Z^+)$ and $E(Z^+) = E(Z^-) \geq E((Z^-)^2)/\|k\|_\infty^{1/2}$. Furthermore,

$$\bar{\lambda} = E(Z^2) = E((Z^+)^2) + E((Z^-)^2) \leq E((Z^+)^2) + \|k\|_\infty^{1/2} E(Z^+) \leq 2\|k\|_\infty^{1/2} E(Z^+).$$

Consider now γ, c such that $0 < -\gamma - c \leq \bar{\lambda}/2\|k\|_\infty^{1/2} \leq E(Z^+)$ then the Paley–Zygmund inequality yields

$$\begin{aligned} \Pr(Z \geq -\gamma - c) &= \Pr(Z^+ \geq -\gamma - c) \geq \frac{(E(Z^+) - (-\gamma - c))^2}{E((Z^+)^2)} \\ &\geq \frac{(\bar{\lambda}/2\|k\|_\infty^{1/2} - (-\gamma - c))^2}{\|k\|_\infty} = \left(\bar{\lambda}/2\|k\|_\infty - (-\gamma - c)/\|k\|_\infty^{1/2}\right)^2. \end{aligned}$$

In particular, when $-\gamma - c = \bar{\lambda}/8\|k\|_\infty^{1/2}$,

$$(2.5) \quad P\psi_\gamma(\langle u, \phi_S(\cdot) - \mathbf{m}_S \rangle_{\mathcal{H}_S} - c) \geq \bar{\lambda}^2/8\|k\|_\infty.$$

2.3. Width of the convex set C

The width of a convex set plays an important role when trying to control the convergence behavior of various convex approximation algorithms. By the width of the convex set $C = \text{cch}\{\phi(x) : x \in \mathcal{X}\}$, where \mathcal{X} is as usual a measurable space and ϕ is a feature map, we mean the size of the projection of C on a function of norm one within the RKHS corresponding to ϕ ,

$$\text{width}_h(C) := \sup_{x \in \mathcal{X}} \langle h, \phi(x) \rangle - \inf_{x \in \mathcal{X}} \langle h, \phi(x) \rangle = \sup_{x \in \mathcal{X}} h(x) - \inf_{x \in \mathcal{X}} h(x),$$

where $h \in \mathcal{H}$, $\|h\| = 1$.

There is a simple relationship between the width of the convex set C in direction h and how close h is to a constant function. In the following, let $\mathbf{1}$ denote the function that is equal to one for all $x \in \mathcal{X}$ and let $\|f\|_\infty = \sup_{x \in \mathcal{X}} |f(x)|$ for any function $f : \mathcal{X} \rightarrow \mathbb{R}$, allowing for $\|f\|_\infty = \infty$. For any $h \in \mathcal{H}$, $\|h\| = 1$,

$$(2.6) \quad \text{width}_h(C) = 2 \inf_{c \in \mathbb{R}} \|h - c\mathbf{1}\|_\infty,$$

In particular, h is a constant function if, and only if, $\text{width}_h(C) = 0$.

Small widths of C in any direction h are a concern when trying to approximate \mathbf{m} because various performance bounds of algorithms discussed in later sections depend on a lower bound on the width; the higher this lower bound the faster the convergence. To be precise, the set C can lie in an affine subspace that is not all of \mathcal{H} and the algorithms we study depend only on the affine subspace. Denote the closure of the affine span of C by $\overline{\text{aff } C}$. In other words, $\overline{\text{aff } C}$ is the closure of $\{\alpha_1 h_1 + \dots + \alpha_n h_n : n \in \mathbb{N}, h_i \in C, \alpha_i \in \mathbb{R} \text{ for all } i \leq n\}$ which is a closed affine subspace. Furthermore, let $U_C = \overline{\text{aff } C} - f$, where f is any element of C , then U_C is a closed subspace of \mathcal{H} . Observe that the dimension of U_C^\perp is at most one since for $h \in U_C^\perp$ it holds that $h(x) = \langle h, \phi(x) \rangle = \langle h, \phi(y) \rangle = h(y)$ for all $x, y \in \mathcal{X}$, and, hence, only constant functions can lie in U_C^\perp .

The key quantity which influences the behavior of the algorithms is

$$\inf_{h \in U_C, \|h\|=1} \text{width}_h(C).$$

If $\mathbf{1}$ lies in the RKHS then $\mathbf{1}$, and all constant functions, lie in U_C^\perp and we do not have to worry about them. The important question is now, how closely can an $h \in U_C$, $\|h\| = 1$, approximate a constant function.

2.3.1. Outline

In this section, we investigate a number of different approaches to control the width of C . Section 2.3.2 summarizes basic observations for infinite dimensional \mathcal{H} and highlights a key difficulty that needs to be addressed if one wants to work with infinite dimensional \mathcal{H} . A key tool in approximation theory to measure how well a function can be approximated is the K -functional. In Section 2.3.3 we use the K -functional to measure how well the constant functions can be approximated by \mathcal{H} and we relate the width of the convex set C to the K -functional and the smallest eigenvalue of a Mercer decomposition of the kernel function. We then aim to complement these results with results that rely on kernel matrices. The first step to achieve this goal is taken in Section 2.3.4 where we introduce the spaces \mathcal{H}^+ and \mathcal{H}^- . The space \mathcal{H}^+ expands a space \mathcal{H} , which does not contain constant functions, by adding constant functions to it. Similarly, \mathcal{H}^- removes the constant functions from a space \mathcal{H} that contains constant functions. In Section 2.3.5 we then make use of the spaces \mathcal{H}^+ and \mathcal{H}^- to provide lower bounds on the width of C that depend on the smallest eigenvalue of suitably chosen kernel matrices. Proposition 2.10 in that section provides a summary of the key results concerning the lower bounds on the width of C that we derived. The results in these sections are for the unknown

convex set C . It seems natural to also consider the empirical counterpart C_n which can actually be accessed. We adapt the approaches to C_n in Section 2.3.6 and discuss how one might get good lower bounds on the width of C_n in practice. Such lower bounds can be used to quantify the quality of a coresets and can be useful for providing stopping criteria for compression algorithms. We do not explore this possibility in this paper and leave this for future research. Also, let us highlight that Section 2.3.6 is not required for the main theorems and can safely be skipped if one wants to get without delay to the main results.

2.3.2. Some comments regarding infinite dimensional \mathcal{H}

Let us recall some topological properties. If \mathcal{X} is compact and ϕ is continuous, then $\phi[\mathcal{X}]$ is compact [Eng89, Theorem 3.1.10]. Due to Mazur's theorem $C = \text{cch } \phi[\mathcal{X}]$ is then also compact [DU77, Theorem 12, p. 51]. This implies, in particular, that there exists no norm ball inside $\phi[\mathcal{X}]$, $\text{ch } \phi[\mathcal{X}]$ or $\text{cch } \phi[\mathcal{X}]$ whenever \mathcal{H} is infinite dimensional because a closed norm ball inside the compact set $\text{cch } \phi[\mathcal{X}]$ would be compact [Eng89, Theorem 3.1.2]. However, closed norm balls in infinite dimensional Hilbert spaces are not compact [Wer02, Section I.2.7]. Similarly, there exists no norm ball B such that $B \cap \overline{\text{aff } C}$ lies inside C .

Furthermore, whenever \mathcal{H} is infinite dimensional, C is compact, $(e_n)_{n \geq 1}$ is an orthonormal sequence in \mathcal{H} and $\epsilon > 0$, it holds that for only finite many of the e_n the width $\text{width}_{e_n}(C)$ can be greater than ϵ . Assume otherwise and let $I : \mathbb{N} \rightarrow \mathbb{N}$ be an enumeration of all the elements e_n for which the width is greater than ϵ . Also assume w.l.o.g. that C is centered in the sense that for all $n \in \mathbb{N}$, $\sup_{u \in C} \langle u, e_{I(n)} \rangle + \inf_{u \in C} \langle u, e_{I(n)} \rangle = 0$. Since C is compact $\sup_{u \in C} \langle u, e_{I(1)} \rangle$ is attained at some point $u_1 \in C$. Inductively, we can select a countably infinite sequence of points $(u_n)_{n \geq 1}$ in C such that $\|u_n - u_m\| \geq \epsilon/4 > 0$ whenever $n \neq m$: given points u_1, \dots, u_n there exists $m \in \mathbb{N}$ such that $\max_{i \leq n} |\langle u_i, e_{I(m')} \rangle| \leq \epsilon/4$ for all $m' \geq m$. Let u_{n+1} be a point in C such that $\epsilon/2 \leq \sup_{u \in C} \langle u, e_{I(m')} \rangle = \langle u_{n+1}, e_{I(m')} \rangle$. Then $\|u_{n+1} - u_i\| \geq \epsilon/4$ for all $i \leq n$. Hence, we have countably infinitely many points with distance of at least $\epsilon/4$ between them. These points give rise to an open cover of C that does not contain a finite sub-cover, contradicting the compactness of C .

This last statement implies that whenever \mathcal{H} is infinite dimensional, C is compact, and $\|k\|_\infty < \infty$ then for any $\epsilon > 0$ there are infinitely many orthonormal elements h_1, h_2, \dots in \mathcal{H} such that for each i , $\sup_{x \in \mathcal{X}} h_i(x) - \inf_{x \in \mathcal{X}} h_i(x) \leq \epsilon$. Furthermore, at most one of the h_i 's can be constant, because if h_i and h_j , $i \neq j$, were both constants then they clearly would not be orthogonal.

2.3.3. Interpolation spaces

Interpolation spaces are useful when trying to quantify the width of C because we can use them to measure how accurately constant functions can be approximated. Consider $\mathbf{1}$ as an element of $C(\mathcal{X})$ and let \mathcal{H} be an RKHS that is continuously embedded in $C(\mathcal{X})$; for simplicity, we will treat \mathcal{H} as a subset of $C(\mathcal{X})$. Furthermore, define for $\theta \in (0, 1)$ the interpolation space $\mathcal{H}_\theta := (C(\mathcal{X}), \mathcal{H})_\theta = \{f : \|f\|_\theta < \infty\}$,

where $\|f\|_\theta = \sup_{t>0} K(f, t)/t^\theta$ and $K : C(\mathcal{X}) \times (0, \infty) \rightarrow \mathbb{R}$ is the K-functional defined by

$$K(f, t) = \inf_{h \in \mathcal{H}} (\|f - h\|_\infty + t\|h\|).$$

If $\mathbf{1} \in \mathcal{H}_\theta$ then for any $r > 0$ there exists an element $h \in \mathcal{H}$, $\|h\| \leq r$, such that $\|\mathbf{1} - h\|_\infty \leq \|\mathbf{1}\|_\theta^{1/(1-\theta)} r^{-2\theta/(1-\theta)}$. In particular, for any $\epsilon > 0$ there exists an r and $h \in \mathcal{H}$, $\|h\| \leq r$, such that $\|\mathbf{1} - h\|_\infty < \epsilon$. Therefore, with $c = 1/\|h\|$ and $h^* = h/\|h\|$, i.e. $\|h^*\| = 1$, it holds that $\|c\mathbf{1} - h^*\|_\infty < \epsilon$ and $\text{width}_{h^*}(C) \leq 2\epsilon$. If $\mathbf{1}$ itself does not lie in \mathcal{H} then h^* lies in the affine span of C and is a problematic direction.

In the finite dimensional case the situation is simpler. If \mathcal{H} is finite dimensional and if the constant function is not in \mathcal{H} then it is also not in any of the interpolation spaces since \mathcal{H}_θ is a subset of the closure of \mathcal{H} in $C(\mathcal{X})$. But because \mathcal{H} is finite dimensional the closure of \mathcal{H} is equal to \mathcal{H} , i.e. $\mathcal{H}_\theta = \mathcal{H}$ for all $\theta \in (0, 1)$. The K-functional can be used in this case to quantify how well $\mathbf{1}$ can be approximated.

The K-functional has a few useful properties with regard to the constant function. Observe that $K(\mathbf{1}, 1) \leq \|\mathbf{1} - 0\|_\infty = 1$, which does not need any conditions on the kernel function. When $\|k\|_\infty \leq 1$ then we also have for any $h \in \mathcal{H}$ that

$$\|h - \mathbf{1}\|_\infty + \|h\| \geq (1 - \|h\|) + \|h\| = 1$$

since $\|h\|_\infty \leq \|k\|_\infty^{1/2} \|h\| \leq \|h\|$. Hence, $K(\mathbf{1}, 1) = 1$ whenever $\|k\|_\infty \leq 1$. It is straight forward to generalize this to any $c \in \mathbb{R}$ whenever $\|k\|_\infty < \infty$, i.e.

$$(2.7) \quad K(c\mathbf{1}, \|k\|_\infty^{1/2}) = c.$$

Also, for any $c \in \mathbb{R}$, $t > 0$ we have the trivial bound $K(c\mathbf{1}, t) \leq c$. For $t < \|k\|_\infty^{1/2}$ the value $K(c\mathbf{1}, t)$ can be smaller than c . If $K(c\mathbf{1}, t) < c$ then for any $\epsilon > 0$ there exists a function $h \in \mathcal{H}$, $h \neq 0$, such that

$$K(c\mathbf{1}, t) + \epsilon \geq \|c\mathbf{1} - h\|_\infty + t\|h\|,$$

and the norm of such an element h is bounded by

$$\frac{c - K(c\mathbf{1}, t) - \epsilon}{\|k\|_\infty^{1/2}} \leq \|h\| \leq \frac{K(c\mathbf{1}, t) + \epsilon}{t}.$$

Furthermore,

$$K(c\mathbf{1}, t) = |c| \inf_{h \in \mathcal{H}} (\|\mathbf{1} - h/c\|_\infty + t\|h/c\|) = |c|K(\mathbf{1}, t)$$

and a minimizer exists for $K(\mathbf{1}, t)$ if, and only if, there exists a minimizer for $K(c\mathbf{1}, t)$. The relation between these minimizers is straight forward: h^* is a minimizer of $K(\mathbf{1}, t)$ if, and only if, ch^* is a minimizer of $K(c\mathbf{1}, t)$.

When \mathcal{H} is finite dimensional and $\|k\|_\infty < \infty$ then there exists a minimizer of the K-functional. For any $c \in \mathbb{R}$, $t > 0$,

$$K(c\mathbf{1}, t) = \inf_{h \in \mathcal{H}} \|c\mathbf{1} - h\|_\infty + t\|h\| = \min_{h \in A} \|c\mathbf{1} - h\|_\infty + t\|h\|,$$

where $A = \{h : h \in \mathcal{H}, \|h\| \leq (c/t) \wedge (1 + K(c\mathbf{1}, t)/t)\}$. This holds because A is compact and $h \mapsto \|c\mathbf{1} - h\|_\infty + t\|h\|$ is continuous whenever $\|k\|_\infty < \infty$. The norm of such a minimizer $h_{t,c}^*$ is bounded by $(c \wedge K(c\mathbf{1}, t))/t$ (the additional one in the

definition of A is, in fact, unnecessary as the above argument shows that the infimum is attained). Hence, we have that

$$\frac{1 - K(\mathbf{1}, t)}{\|k\|_\infty^{1/2}} \leq \frac{\|h_{t,c}^*\|}{c} \leq \frac{1 \wedge K(\mathbf{1}, t)}{t}.$$

In fact, we can say more about the norm of $h_{t,c}^*$ in the finite dimensional case. Notice first that $\|h_{t,c}^*\|_\infty \leq 2c$ since otherwise 0 would be a better approximation of $c\mathbf{1}$. Since the RKHS is finite dimensional this implies an upper bound on the RKHS-norm of $h_{t,c}^*$ as the next lemma shows. The lemma is actually of major importance in this paper and we develop it further than what is needed for the current discussion. In particular, the second part of the lemma is concerned with the relation between $\|h\|_\infty$ and $\|h\|$ when Mercer's theorem (e.g. [SC08, Theorem 4.49]) applies. Recall that Mercer's theorem provides us under certain conditions with orthonormal elements $e_1^\bullet, \dots, e_d^\bullet$ in $L^2(\mathcal{X}, \mu)$, μ being a Borel measure on \mathcal{X} , where e_1, \dots, e_d are continuous functions and such that $\tilde{e}_i = \tilde{\lambda}_i e_i$ for all $i \leq d$, where $\tilde{\lambda}_1 \geq \dots \geq \tilde{\lambda}_d > 0$, lie in the RKHS \mathcal{H} and are an orthonormal basis of \mathcal{H} . The kernel function has to be continuous for Mercer's theorem to hold. There are various forms of Mercer's theorem together with a variety of assumptions for the theorems to hold. Instead of making such assumptions the following lemma assumes directly in its second part that the e_1, \dots, e_d exist and have the above properties.

LEMMA 2.3. — *Let \mathcal{X} be a set, k a kernel on \mathcal{X} such that the corresponding RKHS \mathcal{H} is d -dimensional. For any $c \in \mathbb{R}$, $\{h : \|h\|_\infty \leq c\}$ is a compact subset of \mathcal{H} . Furthermore, for $h \in \mathcal{H}$ and any points x_1, \dots, x_d for which $k(x_1, \cdot), \dots, k(x_d, \cdot)$ are linearly independent,*

$$(\lambda_d/d)^{1/2} \|h\| \leq \|h\|_\infty,$$

where λ_d is the smallest eigenvalue of the kernel matrix for the points x_1, \dots, x_d .

Whenever \mathcal{X} is a topological space, k is a continuous kernel function on \mathcal{X} and there exist continuous functions $e_i : \mathcal{X} \rightarrow \mathbb{R}$, $i \leq d$, and a Borel probability measure μ on \mathcal{X} such that $e_1^\bullet, \dots, e_d^\bullet$ are orthonormal in $L^2(\mathcal{X}, \mu)$, and $\{\tilde{e}_i\}_{i \leq d}$ is an orthonormal basis of \mathcal{H} where $\tilde{e}_i = \tilde{\lambda}_i^{1/2} e_i$, for all $i \leq d$, and $\tilde{\lambda}_1 \geq \tilde{\lambda}_2 \geq \dots \geq \tilde{\lambda}_d > 0$, then

$$\tilde{\lambda}_d^{1/2} \|h\| \leq \|h\|_\infty.$$

Proof.

(a). — For the first statement let x_1, \dots, x_d be such that $k(x_1, \cdot), \dots, k(x_d, \cdot)$ are linearly independent. Observe that such points always exist: assume that $d' < d$ points $x_1, \dots, x_{d'}$ exist such that any $k(x, \cdot)$ lies in the span of $k(x_1, \cdot), \dots, k(x_{d'}, \cdot)$. Now any $h \in \mathcal{H}$ of the form $\sum_{j=1}^N \alpha_j k(z_j, \cdot)$ with coefficients α_j and $z_j \in \mathcal{X}$ can be written as a sum $\sum_{i=1}^{d'} \beta_i k(x_i, \cdot)$ with suitable coefficients β_i . The family of functions h that can be written this way lies dense in \mathcal{H} , that is, $\text{span}\{k(x_i, \cdot) : i \leq d'\}$ is a dense subspace of \mathcal{H} . But this subspace is closed and therefore equal to \mathcal{H} . Hence, \mathcal{H} is d' -dimensional contradicting our assumption about \mathcal{H} .

Consider the linear operator $A : \mathcal{H} \rightarrow \mathbb{R}^d$, defined for any $f \in \mathcal{H}$ by

$$Af = (f(x_1), \dots, f(x_d))^\top = (\langle f, k(x_1, \cdot) \rangle, \dots, \langle f, k(x_d, \cdot) \rangle)^\top.$$

The operator is bounded since

$$\|Af\|_{\mathbb{R}^d}^2 \leq \|f\|^2 \sum_{i=1}^d k(x_i, x_i) \quad \text{and} \quad \|A\|_{\text{op}}^2 \leq \sum_{i=1}^d k(x_i, x_i).$$

A is also injective. One way to see this is by means of Gram–Schmidt orthogonalization through which we gain an orthonormal basis e_1, \dots, e_d of \mathcal{H} from $k(x_1, \cdot), \dots, k(x_d, \cdot)$ and for any $f, g \in \mathcal{H}$ it holds that $f = g$ if, and only if, $\langle e_i, f \rangle = \langle e_i, g \rangle$ for all $i \leq d$ if, and only if, $\langle k(x_i, \cdot), f \rangle = \langle k(x_i, \cdot), g \rangle$ for all $i \leq d$.

Since A is injective and the dimension of \mathcal{H} is d it follows that A is surjective and invertible. By the open mapping theorem A^{-1} is continuous and $A^{-1}[\{v : v \in \mathbb{R}^d, \|v\|_\infty \leq c\}]$ is a compact subset of \mathcal{H} .

(b). — Let K be the kernel matrix corresponding to the points x_1, \dots, x_d . The rows of the kernel matrix are linearly independent since they are the images of the linearly independent elements $k(x_1, \cdot), \dots, k(x_d, \cdot)$ under the isomorphism A . Hence, K is invertible and for any $y \in \mathbb{R}^d$, with $\alpha = K^{-1}y$,

$$A\left(\sum_{i=1}^d \alpha_i k(x_i, \cdot)\right) = \sum_{i=1}^d \alpha_i (k(x_i, x_1), \dots, k(x_i, x_d))^\top = K\alpha = y.$$

In particular, for $f = \sum_{i=1}^d \alpha_i k(x_i, \cdot)$, with $\alpha_i \in \mathbb{R}$, it follows that $\alpha = K^{-1}A(f)$. We have a useful inner product on \mathbb{R}^d given by $\langle x, y \rangle_{K^{-1}} = x^\top K^{-1}y$. For arbitrary $f, g \in \mathcal{H}$ with $f = \sum_{i=1}^d \beta_i k(x_i, \cdot)$ and $g = \sum_{i=1}^d \alpha_i k(x_i, \cdot)$,

$$\langle f, g \rangle = \beta^\top K\alpha = (K^{-1}Af)^\top K(K^{-1}Ag) = \langle Af, Ag \rangle_{K^{-1}}.$$

Applying this to h ,

$$\begin{aligned} \|h\|^2 &= (Ah)^\top K^{-1}(Ah) = \text{tr}(K^{-1}(Ah)(Ah)^\top) \\ &\leq \|K^{-1}\|_{\text{op}}(Ah)^\top(Ah) \leq d\|K^{-1}\|_{\text{op}}\|h\|_\infty^2. \end{aligned}$$

(c). — Now assuming that k is continuous and the e_1, \dots, e_d have the assumed properties, we can write any $h \in \mathcal{H}$ as $h = \sum_{i=1}^d \alpha_i \tilde{e}_i$, $\sum_{i=1}^d \alpha_i^2 = \|h\|^2$, and

$$\|h\|_2^2 = \sum_{i=1}^d \alpha_i^2 \tilde{\lambda}_i \geq \tilde{\lambda}_d \|h\|^2.$$

Since $\|h\|_2 \geq \tilde{\lambda}_d \|h\|$ and μ is a probability measure, there has to be some point $x \in \mathcal{X}$ at which $|h(x)| \geq \tilde{\lambda}_d \|h\|$. □

Example 2.4. — Consider the space $\mathcal{X} = \{1, \dots, d\}$ with kernel function $k(x, y) = 1$ if $x = y$ and zero otherwise. Then $\|h\|^2 = \sum_{i=1}^d |h(i)|^2$ and if $h(i) = c > 0$ for all $i \leq d$ then $\|h\| = \sqrt{d}\|h\|_\infty$ which matches the bound if we use $x_1 = 1, \dots, x_d = d$.

Coming back to the case of \mathcal{H} being d -dimensional, $\|k\|_\infty < \infty$ and $x_1, \dots, x_d \in \mathcal{X}$ be any points such that $k(x_1, \cdot), \dots, k(x_d, \cdot)$ are linearly independent and the kernel matrix is full rank. Furthermore, let λ_d be the smallest eigenvalue of the kernel matrix. Consider the map $\psi(h) = \|\mathbf{1} - h\|_\infty$. By a similar argument as above we can infer that *there exists a minimizer of ψ* . However, *the minimizer is usually not unique*. Consider, for example, $\mathcal{X} = [-1, 1]$ and the RKHS consisting of linear and quadratic functions such that $x \mapsto x$ and $x \mapsto x^2$ both have norm 1. Then

both of these functions minimize the distance to $\mathbf{1}$ as does 0. Any minimizer h of ψ has norm $\|h\|_\infty \leq 2$ and, therefore, according to Lemma 2.3, it has an RKHS norm $\|h\| \leq (4d/\lambda_d)^{1/2} =: r$. In particular, all minimizers of ψ are included in the compact ball $B = \{h : h \in \mathcal{H}, \|h\| \leq r\}$. Let A be the set of all minimizers of ψ then A is a compact set: if A is finite then this follows right away. Otherwise, take a convergent sequence $\{h_n\}_{n \in \mathbb{N}}$ in A and denote the limit by h . Since for all $x \in \mathcal{X}$, $|h(x) - 1| = \lim_{n \rightarrow \infty} |h_n(x) - 1| \leq \min_{h \in \mathcal{H}} \|\mathbf{1} - h\|_\infty$ and $h \in A$. Finally, consider the norm as a function on A . The norm is continuous and the image of the compact set A under the norm is a compact subset in \mathbb{R} . Hence, there exists an element h^* in A of maximal norm. Let us assume first that $h^* \neq 0$. For such an element h^* let $b = 1/\|h^*\|$ and note that $\|b\mathbf{1} - bh^*\|_\infty = \inf_{c \in \mathbb{R}} \min_{\|h\|=1} \|c\mathbf{1} - h\|_\infty$. Otherwise, there is an element \tilde{h} , $\|\tilde{h}\| = 1$, and a c such that $\|c\mathbf{1} - \tilde{h}\|_\infty < \|b\mathbf{1} - bh^*\|$. The constant c cannot be equal to b since then $\|\mathbf{1} - \tilde{h}/b\|_\infty < \|\mathbf{1} - h^*\|$ in contradiction to our assumption on h^* . It also cannot be larger than b because then $\|\mathbf{1} - \tilde{h}/c\|_\infty < (b/c)\|\mathbf{1} - h^*\|_\infty < \|\mathbf{1} - h^*\|_\infty$ which is again in contradiction to h^* being a best approximation of $\mathbf{1}$. But c can also not be smaller than b ; whenever $\|c\mathbf{1} - \tilde{h}\|_\infty$ is minimal it follows that $\|\mathbf{1} - \tilde{h}/c\|_\infty$ is minimal and equal to $\|\mathbf{1} - h^*\|_\infty$. However, $\|\tilde{h}/c\| = 1/c > 1/b = \|h^*\|$ in contradiction to the assumption that $\|h^*\|$ has maximal norm within A . Therefore,

$$\lim_{t \rightarrow 0} K(b\mathbf{1}, t) = \|b\mathbf{1} - bh^*\|_\infty = \inf_{c \in \mathbb{R}} \min_{\|h\|=1} \|c\mathbf{1} - h\|_\infty = (1/2) \inf_{\|h\|=1} \text{width}_h(C)$$

and, since $K(b\mathbf{1}, t) = bK(\mathbf{1}, t) \geq (1/r)K(\mathbf{1}, t)$, it follows that

$$(2.8) \quad \left(\frac{\lambda_d}{d}\right)^{1/2} \lim_{t \rightarrow 0} K(\mathbf{1}, t) \leq \inf_{\|h\|=1} \text{width}_h(C).$$

If $h^* = 0$ then $\lim_{t \rightarrow 0} K(\mathbf{1}, t) = \|\mathbf{1}\|_\infty = 1$ but also for any $c \in \mathbb{R}$, $h \in \mathcal{H}$, $\|c\mathbf{1} - h\|_\infty > \|c\mathbf{1}\|_\infty$ since otherwise h/c would be a minimizer of norm greater than zero, contradicting the assumption that $h^* = 0$ is the minimizer with the largest norm. For $h \in \mathcal{H}$, there is a sequence of points x_1, x_2, \dots such that $\lim_{n \rightarrow \infty} h(x_n)$ converges and $|h(x_n)| \rightarrow \|h\|_\infty$. Fix one such sequence and let $\sigma(h)$ be the sign of all but finitely many elements of this sequence $h(x_1), h(x_2), \dots$, e.g. if $\sigma(h)$ is positive and h attains maxima then there is a point x such that $h(x) = \|h\|_\infty$. By another application of Lemma 2.3 it follows for any $h \in \mathcal{H}$, $\|h\| = 1$, that

$$\text{width}_h(C) = \|\sigma(h)\|h\|_\infty \mathbf{1} - h\|_\infty > \|h\|_\infty \geq \left(\frac{\lambda_d}{d}\right)^{1/2} \lim_{t \rightarrow 0} K(\mathbf{1}, t)$$

and

$$\inf_{\|h\|=1} \text{width}_h(C) \geq \left(\frac{\lambda_d}{d}\right)^{1/2} \lim_{t \rightarrow 0} K(\mathbf{1}, t).$$

We can set in the above derivation r to $(4/\tilde{\lambda}_d)^{1/2}$ when Mercer's theorem applies, where $\tilde{\lambda}_d$ is the d^{th} eigenvalue of T_k . The bound then becomes

$$\tilde{\lambda}_d^{1/2} \lim_{t \rightarrow 0} K(\mathbf{1}, t) \leq \inf_{\|h\|=1} \text{width}_h(C).$$

These results are only meaningful if $\mathbf{1}$ is not in the RKHS. In the next section we discuss an approach to remove constants from an RKHS which allows us, among other things, to extend these results to RKHSs that contain constants.

2.3.4. Adding and removing constants

It is sometimes useful to be able to remove constant functions from an RKHS or to add constant functions to an RKHS. There is an efficient way to do this by manipulating the kernel function.

In the following let \mathcal{X} be some topological space and consider the p.s.d. functions $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ that lie in $\mathcal{L}^2(\mathcal{X} \times \mathcal{X})$ and denote these by \mathcal{K} . Furthermore, consider the partial order on \mathcal{K} given by $k \succeq l$ if, and only if, $k - l$ is p.s.d. where $k, l \in \mathcal{K}$. Also note that \mathcal{K} is not a lattice, i.e. for $k, l \in \mathcal{K}$ the infimum $k \wedge l$ and the supremum $k \vee l$ will generally not be defined.

For a function $f : \mathcal{X} \rightarrow \mathbb{R}$ we let $f \otimes f$ be the function that maps (x, y) to $f(x)f(y)$ for any $x, y \in \mathcal{X}$. There is a simple criterion which tells us if $f \in \mathcal{H}_k$ for a kernel function $k \in \mathcal{K}$. Assume that $f \otimes f \in \mathcal{L}^2(\mathcal{X} \times \mathcal{X})$, then $f \in \mathcal{H}_k$ if, and only if, there exists a $c > 0$ with $c^2 k \succeq f \otimes f$. In case that $f \in \mathcal{H}_k$ it holds that $\|f\|_k = \inf\{c : c^2 k \succeq f \otimes f\}$.

This observation motivates the following definitions.

DEFINITION 2.5. — For an RKHS \mathcal{H} with kernel k that does not contain $\mathbf{1}$ define

$$(2.9) \quad k^+ := k + \mathbf{1} \otimes \mathbf{1} \quad \text{and} \quad \mathcal{H}^+ := \mathcal{H}_{k^+}.$$

The function k^+ is a kernel function being the sum of the kernel functions k and $\mathbf{1} \otimes \mathbf{1}$ and \mathcal{H}^+ is well defined. We denote the norm of \mathcal{H}^+ by $\|\cdot\|_+$ and we can observe that

$$\|\mathbf{1}\|_+ = \inf\{c : c^2(k + \mathbf{1} \otimes \mathbf{1}) \succeq \mathbf{1} \otimes \mathbf{1}\} \leq 1.$$

In fact, $\|\mathbf{1}\|_+ = 1$ because otherwise there exists a $c < 1$ such that

$$c^2 k \succeq (1 - c^2)\mathbf{1} \otimes \mathbf{1} \implies \left(\frac{c}{\sqrt{1 - c^2}}\right)^2 k \succeq \mathbf{1} \otimes \mathbf{1} \implies \mathbf{1} \in \mathcal{H}.$$

We also have that $\mathcal{H} \subset \mathcal{H}^+$, since $k \preceq k^+$, and for $h \in \mathcal{H}$,

$$(2.10) \quad \|h\|_+ \leq \|h\|.$$

When the RKHS \mathcal{H} is finite dimensional then $\|h\|_+$ is actually equal to $\|h\|$. To show this we make use of the following lemma which is a simple extension of [PR16, Section 5.3].

LEMMA 2.6. — Let h_1, \dots, h_d be linearly independent functions mapping from some topological space \mathcal{X} to \mathbb{R} and let $a_1, \dots, a_d > 0$ then $\kappa = \sum_{i=1}^d a_i h_i \otimes h_i$ is a kernel function, the functions h_i lie in \mathcal{H}_κ and are orthogonal in \mathcal{H}_κ . Furthermore, the dimension of \mathcal{H}_κ is d and $\|h_i\|_\kappa = 1/\sqrt{a_i}$.

Now, let $d < \infty$ be the dimension of \mathcal{H} , choose orthogonal functions $h_1, \dots, h_d \in \mathcal{H}$, $h_1, \dots, h_d \neq 0$, and define the kernel $\kappa = \sum_{i=1}^d (1/\|h_i\|^2) h_i \otimes h_i$. Then $k = \kappa$. This follows because, according to the above lemma, both spaces consist of $\text{span}\{h_1, \dots, h_d\}$, $\|h_i\|_\kappa = \|h_i\|_k$, for all $i \leq d$, and the h_i 's are orthogonal in both spaces, i.e. $H_k = H_\kappa$ which implies that $k = \kappa$. The importance of this statement is that it shows that we can write the kernel as a finite sum of weighted tensor products.

From this description of κ we also gain that $k^+ = \sum_{i=1}^d a_i h_i \otimes h_i + \mathbf{1} \otimes \mathbf{1}$ and, because $\mathbf{1}$ is not in the original RKHS \mathcal{H} , it follows that $\mathbf{1}$ is linearly independent of h_1, \dots, h_d which implies that $\mathbf{1}$ is orthogonal to h_1, \dots, h_d in \mathcal{H}^+ .

Consider now one of the h_i 's. We like to show that $\|h_i\|_+ \geq \|h_i\|$ which then implies, together with (2.10), that $\|h_i\|_+ = \|h_i\|$ and $\|h\|_+ = \|h\|$ for all $h \in \mathcal{H}$; the h_i 's are orthogonal in both \mathcal{H} and \mathcal{H}^+ . Let $l = k^+ - (1/\|h_i\|^2) h_i \otimes h_i$ so that $h_i \notin \mathcal{H}_l$. Furthermore, consider any c such that $0 < c < \|h_i\|$. If $\|h_i\|_+ = c$ then

$$c^2 k^+ \succeq h_i \otimes h_i \implies c^2 l \succeq (1 - c^2/\|h_i\|^2) h_i \otimes h_i \implies \frac{c^2 \|h_i\|^2}{\|h_i\|^2 - c^2} l \succeq h_i \otimes h_i \implies h_i \in \mathcal{H}_l,$$

which is impossible and, therefore, $\|h_i\|_+ \geq \|h_i\|$.

For an RKHS that contains already the constant functions it is useful to be able to remove these from the RKHS.

DEFINITION 2.7. — *Let \mathcal{H} be an RKHS that does contain $\mathbf{1}$ and is not of dimension 1. Define*

$$(2.11) \quad k^- := k - c^2 \mathbf{1} \otimes \mathbf{1}, \quad \text{where } c = \inf\{\tilde{c} : \tilde{c}^2 k \succeq \mathbf{1} \otimes \mathbf{1}\}, \quad \text{and } \mathcal{H}^- := \mathcal{H}_{k^-}.$$

It follows right away that $\mathbf{1} \notin \mathcal{H}_-$ and because, $k^- \preceq k$ we know that $\mathcal{H}^- \subset \mathcal{H}$ and $\|h\|_- \leq \|h\|$ for all $h \in \mathcal{H}^-$. Next, notice that we can write $k = \sum_{i=1}^{d-1} a_i h_i \otimes h_i + c^2 \mathbf{1} \otimes \mathbf{1}$ where $h_1, \dots, h_{d-1}, \mathbf{1}$ are orthogonal in \mathcal{H} and $a_1, \dots, a_{d-1} > 0$. Due to the orthogonality it follows that the h_1, \dots, h_{d-1} are linearly independent elements in \mathcal{H}^- and \mathcal{H}^- is $d - 1$ dimensional. Lemma 2.6 tells us furthermore that h_1, \dots, h_{d-1} are orthogonal in \mathcal{H}^- . Finally, for all $i \leq d - 1$ we have that $\|h_i\|_- = \|h_i\|$; assume $c = 1$ and observe that in this case $(\mathcal{H}^-)^+ = \mathcal{H}$ and due to the above results for \mathcal{H}^+ we can conclude that $\|h_i\|_- = \|h_i\|_{-+} = \|h_i\|$. The above argument for \mathcal{H}^+ does not rely on $\|\mathbf{1}\| = c = 1$ and we can generalize this result right away to any $c > 0$. Because the norm of the h_i does not change and since the h_i are orthogonal we can conclude that $\|h\|_- = \|h\|$ for all $h \in \mathcal{H}_-$. We summarize these results in the following lemma.

LEMMA 2.8. — *If \mathcal{H} is a finite dimensional RKHS with dimension d , kernel $k \in \mathcal{K}$, and which does not contain $\mathbf{1}$ then \mathcal{H}^+ , is $d + 1$ dimensional, $\mathcal{H} \subset \mathcal{H}^+$, $\mathbf{1} \in \mathcal{H}^+$ with $\|\mathbf{1}\|_+ = 1$, $\langle g, h \rangle_+ = \langle g, h \rangle$ for all $g, h \in \mathcal{H}$, and $\mathbf{1}$ is orthogonal in \mathcal{H}^+ to all $h \in \mathcal{H}$. Similarly, if \mathcal{H} is a finite dimensional RKHS with dimension $d > 1$, kernel $k \in \mathcal{K}$, and which does contain $\mathbf{1}$ then \mathcal{H}^- , is $d - 1$ dimensional, $\mathcal{H}^- \subset \mathcal{H}$, $\mathbf{1} \notin \mathcal{H}^-$, $\langle g, h \rangle_- = \langle g, h \rangle$ for all $g, h \in \mathcal{H}$ which are orthogonal to $\mathbf{1}$.*

2.3.5. Lower bounds on the approximation error in finite dimensions

In finite dimensions we can now provide lower bounds on the approximation error of any function $f : \mathcal{X} \rightarrow \mathbb{R}$. Before specializing to constant functions we take a short

detour and discuss the general technique. The approach to get lower bounds is the following: let $k = \sum_{i=1}^d a_i h_i \otimes h_i$ for linearly independent h_1, \dots, h_d and $a_i > 0$. If f is linearly dependent on the h_i 's then $f \in \mathcal{H}$. Otherwise, we can move to the kernel function $k' = \sum_{i=1}^d a_i h_i \otimes h_i + f \otimes f$ and the corresponding RKHS \mathcal{H}' . The function f is orthogonal to h_1, \dots, h_d in \mathcal{H}' . That means that the lowest approximation error, when approximating f by functions in the subspace corresponding to \mathcal{H} , is given by the projection onto this subspace. Due to the orthogonality the projection of f onto this subspace is just the origin and the approximation error is $\|f\|_{\mathcal{H}'} = 1$ when measured in the RKHS norm of \mathcal{H}' . If we consider the constraint that the approximation has to lie in \mathcal{H} and has to have norm $\|h\| = 1$ then the best approximation error of f is $\sqrt{2}$, i.e.

$$\inf_{h \in \mathcal{H}, \|h\|=1} \|f - h\|_{\mathcal{H}'} = \sqrt{2}.$$

To gain a lower bound on the approximation error in $\|\cdot\|_\infty$ we use Lemma 2.3 which shows that

$$\inf_{h \in \mathcal{H}, \|h\|=1} \|f - h\|_\infty \geq \sqrt{2} \left(\frac{\lambda_{d+1}}{d+1} \right)^{1/2},$$

where we get $d + 1$ since we use the RKHS \mathcal{H}' which has dimension $d + 1$. The constant λ_{d+1} is the smallest eigenvalue of a kernel matrix corresponding to points x_1, \dots, x_{d+1} such that $k'(x_1, \cdot), \dots, k'(x_{d+1}, \cdot)$ are linearly independent. Notice, that this approximation error depends implicitly on the particular function f through the kernel matrix and the smallest eigenvalue. The bound can become loose when $\|f\|_\infty$ is significantly larger than $\|h_i\|_\infty$, but observe that we can always replace f by cf for some constant $c < 1$ to rescale the infinity norm. In the following, let $k'' = \sum_{i=1}^d a_i h_i \otimes h_i + (cf) \otimes (cf)$ and treat \mathcal{H} as a subset of $\mathcal{H}'' := \mathcal{H}_{k''}$. Such a rescaling leads to a problem in the constraint $\|h\| = 1$ because

$$\inf_{h \in \mathcal{H}, \|h\|=1} \|cf - h\|_\infty = c \inf_{h \in \mathcal{H}, \|h\|=1/c} \|f - h\|_\infty.$$

We can compensate for this by using the constraint $\|h\| = c$. Since $\|cf\|_{\mathcal{H}''} = 1$,

$$\inf_{h \in \mathcal{H}, \|h\|=1} \|f - h\|_\infty = \frac{1}{c} \inf_{h \in \mathcal{H}, \|h\|=c} \|cf - h\|_\infty \geq \frac{\sqrt{1+c^2}}{c} \left(\frac{\lambda_{d+1}}{d+1} \right)^{1/2},$$

where λ_{d+1} is again the smallest eigenvalue of a kernel matrix but now for the kernel k'' .

Example 2.9. — Let $\mathcal{X} = \{0, 1\}$ and $h : \mathcal{X} \rightarrow \mathbb{R}$ be given by $h(0) = 1, h(1) = 0$, and let $f : \mathcal{X} \rightarrow \mathbb{R}$ be defined by $f(0) = 0, f(1) = r$ for $r > 0$. Let \mathcal{H} be the RKHS with kernel $h \otimes h$ which consists of $\text{span}\{h\}$. The smallest approximation error of f by elements in \mathcal{H} which have norm 1 is attained by $-h$ and h and is equal to $\|h - f\|_\infty = r \vee 1$. Considering now the bound: let the kernel of the RKHS \mathcal{H}' be $k = h \otimes h + f \otimes f$. Consider $x_1 = 0, x_2 = 1$ and the corresponding kernel matrix

$$K = \begin{pmatrix} h(0)^2 & 0 \\ 0 & f(1)^2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & r^2 \end{pmatrix}$$

which has minimal eigenvalue $1 \wedge r^2$. The corresponding lower bound is

$$\inf_{g \in \mathcal{H}, \|g\|=1} \|g - f\|_\infty \geq \sqrt{2} \left(\frac{1 \wedge r^2}{2} \right)^{1/2} = 1 \wedge r$$

which is exact when $\|f\|_\infty = 1$ but degrades for r away from 1.

Scaling f by $c = 1/\|f\|_\infty = 1/r$ gives us the kernel $k' = h \otimes h + (1/r)^2 f \otimes f$ and a kernel matrix

$$K' = \begin{pmatrix} h(0)^2 & 0 \\ 0 & (1/r)^2 f(1)^2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

which has minimal eigenvalue 1. The bound becomes

$$\inf_{g \in \mathcal{H}, \|g\|=1} \|g - f\|_\infty \geq \sqrt{\frac{1+r^2}{2}} \geq 1 \wedge r.$$

Coming back to the approximation of constant functions. When \mathcal{H} does not contain the constant functions then an approach to calculate lower bounds is to use the kernel k_+ and the corresponding RKHS \mathcal{H}^+ . The norm of $c\mathbf{1}$ in this RKHS, where $c \in \mathbb{R}$, is $|c|$ and for any such c ,

$$\inf_{h \in \mathcal{H}, \|h\|=1} \|h - c\mathbf{1}\|_\infty \geq \frac{\sqrt{1+|c|^2}}{|c|} \left(\frac{\lambda_{d+1}}{d+1} \right)^{1/2} \geq \left(\frac{\lambda_{d+1}}{d+1} \right)^{1/2}$$

with d being the dimension of \mathcal{H} and λ_{d+1} the lowest eigenvalue of a kernel matrix corresponding to points x_1, \dots, x_{d+1} for the kernel k^+ . Using the right hand side as the lower bound has the advantage that we only deal with one RKHS, i.e. with \mathcal{H}^+ , and we only need λ_{d+1} for that kernel. Scaling of the function $\mathbf{1}$ in dependence of which constant $c\mathbf{1}$ we want to approximate might improve the lower bounds but then λ_{d+1} has to be calculated for the individual scalings.

When Mercer's theorem applies we gain the bound

$$\inf_{c \in \mathbb{R}} \inf_{h \in \mathcal{H}, \|h\|=1} \|h - c\mathbf{1}\|_\infty \geq \tilde{\lambda}_{d+1}^{1/2},$$

where $\tilde{\lambda}_{d+1}$ is the $(d+1)^{\text{th}}$ eigenvalue of T_{k^+} . For Mercer's theorem to apply it is important that k^+ is continuous. But when k is continuous then so is k^+ .

If \mathcal{H} already contains the constant functions then we are interested in determining the width of the convex set in the affine subspace spanned by C . In particular, because $\langle k(x, \cdot), \mathbf{1} \rangle = 1$ for all $x \in \mathcal{X}$, there exists a subspace S of \mathcal{H} that is orthogonal to $\mathbf{1}$ and a $c \neq 0$ such that $\text{aff } C = \text{aff}\{k(x, \cdot) : x \in \mathcal{X}\} = c\mathbf{1} + S$. In fact, $c = \arg \min_{c' \in \mathbb{R}} \|k(x, \cdot) - c'\mathbf{1}\|$, where we can use an arbitrary $x \in \mathcal{X}$ and $S = \mathcal{H}^-$. This is exactly the same situation that we faced above with \mathcal{H}^+ and a lower bound on the width of the convex set in the affine space spanned by it can be gained through

$$\inf_{c \in \mathbb{R}} \inf_{\|h\|=1} \|h - c\mathbf{1}\|_\infty \geq \left(\frac{\lambda_d}{d} \right)^{1/2},$$

where d is the dimension of \mathcal{H} and λ_d the smallest eigenvalue of any kernel matrix

for kernel k . If we can use Mercer’s theorem then we also gain the lower bound

$$\inf_{c \in \mathbb{R}} \inf_{\|h\|_-=1} \|h - c\mathbf{1}\|_\infty \geq \tilde{\lambda}_d^{1/2},$$

where $\tilde{\lambda}_d$ is the d^{th} eigenvalue of T_k .

We can also extend the results from Section 2.3.3 on the application of K -functionals. We summarize in the following proposition these results together with a variety of results on the width of C that we derived up to now. We use the notation $K_-(\mathbf{1}, t)$ for the K -functional corresponding to \mathcal{H}_- . We hope that the use of the letter K for both the K -functional and the kernel matrix does not lead to confusion. To streamline the statement of the following proposition let us say that k has a *Mercer decomposition with lowest eigenvalue* $\tilde{\lambda}_d$ if k is a continuous kernel function on \mathcal{X} and there exist continuous functions $e_i : \mathcal{X} \rightarrow \mathbb{R}, i \leq d$, and a Borel probability measure μ on \mathcal{X} such that $e_1^\bullet, \dots, e_d^\bullet$ are orthonormal in $L^2(\mathcal{X}, \mu)$, $\{\tilde{e}_i\}_{i \leq d}$ is an orthonormal basis of \mathcal{H} , where $\tilde{e}_i = (\tilde{\lambda}_i)^{1/2} e_i$, for all $i \leq d$, and $\tilde{\lambda}_1 \geq \tilde{\lambda}_2 \geq \dots \geq \tilde{\lambda}_d > 0$. Notice that the Mercer decomposition based results in the following proposition do not seem to have a dependence on d beyond the eigenvalue λ_d but this is somewhat misleading as the discussion in Section 2.3.6 demonstrates.

PROPOSITION 2.10. — *Let \mathcal{X} be a measurable set and $k \in \mathcal{K}$ a kernel function defined on \mathcal{X} . The following holds.*

- (1) *If \mathcal{H} is infinite dimensional, \mathcal{X} is compact and k is continuous, then for every $\epsilon > 0$ there exist infinitely many orthonormal elements $(e_n)_{n \geq 1}$ in \mathcal{H} such that $\sup_{n \geq 1} \text{width}_{e_n}(C) < \epsilon$.*

If \mathcal{H} is finite dimensional with dimension $1 \leq d$ then the following hold.

- (2) *If $\mathbf{1} \in \mathcal{H}_\theta$ for some $\theta \in (0, 1)$ then there exists $h \in \mathcal{H}, \|h\| = 1$, such that $\text{width}_h(C) = 0$.*
- (3) *If $\mathbf{1} \notin \mathcal{H}$ then for any $x_1, \dots, x_{d+1} \in \mathcal{X}$ and corresponding kernel matrix $K^+ = (k^+(x_i, x_j))_{i,j \leq d+1}$ with smallest eigenvalue λ_{d+1} ,*

$$\inf_{\|h\|=1} \text{width}_h(C) \geq 2 \left(\frac{\lambda_{d+1}}{d+1} \right)^{1/2}.$$

- (4) *If $\mathbf{1} \notin \mathcal{H}, \|k\|_\infty < \infty$, then for any $x_1, \dots, x_d \in \mathcal{X}$ and corresponding kernel matrix $K = (k(x_i, x_j))_{i,j \leq d}$ with smallest eigenvalue λ_d ,*

$$\inf_{\|h\|=1} \text{width}_h(C) \geq \left(\frac{\lambda_d}{d} \right)^{1/2} \lim_{t \rightarrow 0} K(\mathbf{1}, t).$$

- (5) *If $\mathbf{1} \in \mathcal{H}$ and $2 \leq d$, then for any $x_1, \dots, x_d \in \mathcal{X}$ with corresponding kernel matrix $K = (k(x_i, x_j))_{i,j \leq d}$ and with the smallest eigenvalue of K being λ_d ,*

$$\inf_{\|h\|_-=1} \text{width}_h(C) \geq 2 \left(\frac{\lambda_d}{d} \right)^{1/2}.$$

- (6) If $\mathbf{1} \in \mathcal{H}$, $2 \leq d$, then for any $x_1, \dots, x_{d-1} \in \mathcal{X}$ with corresponding kernel matrix $K^- = (k(x_i, x_j))_{i,j \leq d-1}$ and with the smallest eigenvalue of K^- being λ_{d-1} ,

$$\inf_{\|h\|_{-1}} \text{width}_h(C) \geq \left(\frac{\lambda_{d-1}}{d-1} \right)^{1/2} \lim_{t \rightarrow 0} K_-(\mathbf{1}, t).$$

In the following, let \mathcal{X} be a compact space and k a continuous kernel function on \mathcal{X} . The following hold.

- (7) If k^+ has a Mercer decomposition with smallest eigenvalue $\tilde{\lambda}_{d+1}$ and $\mathbf{1} \notin \mathcal{H}$ then

$$\inf_{\|h\|=1} \text{width}_h(C) \geq 2\tilde{\lambda}_{d+1}^{1/2}.$$

- (8) If k has a Mercer decomposition with smallest eigenvalue $\tilde{\lambda}_d$ and $\mathbf{1} \notin \mathcal{H}$ then

$$\inf_{\|h\|=1} \text{width}_h(C) \geq \tilde{\lambda}_d^{1/2} \lim_{t \rightarrow 0} K(\mathbf{1}, t).$$

- (9) If k has a Mercer decomposition with smallest eigenvalue $\tilde{\lambda}_d$ and $\mathbf{1} \in \mathcal{H}$ then

$$\inf_{\|h\|_{-1}} \text{width}_h(C) \geq 2\tilde{\lambda}_d^{1/2}.$$

- (10) If \mathcal{H} is $d \geq 2$ dimensional, k^- has a Mercer decomposition with smallest eigenvalue $\tilde{\lambda}_{d-1}$ and $\mathbf{1} \notin \mathcal{H}$ then

$$\inf_{\|h\|_{-1}} \text{width}_h(C) \geq \tilde{\lambda}_{d-1}^{1/2} \lim_{t \rightarrow 0} K_-(\mathbf{1}, t).$$

Example 2.11. — Consider the kernels $k_d(x, y) = \sum_{u=1}^d x^u y^u$, with $x, y \in [-1, 1]$ and $1 \leq d \leq 4$. These corresponding RKHSs do not contain the constant functions. To test the kernel matrix based lower bound in a simple experiment we are calculating upper bounds on $\inf_{c \in \mathbb{R}} \inf_{h \in \mathcal{H}, \|h\|=1} \|h - c\mathbf{1}\|_\infty$ in the following way: the functions x^u and x^v are orthogonal in the corresponding RKHSs whenever $u \neq v$ and have norm 1. Therefore, functions of the form $(1/\sqrt{d}) \sum_{u=1}^d x^u$ have norm 1. To get a good approximation of constant functions we use such functions for $d = 3, 4$, with signs adjusted so that the different terms cancel each other as well as possible. In detail, for $d = 1$ we use the function $h_1(x) = x$ which has approximation error 1 when approximating the (constant) function 0; for $d = 2$ we use $h_2(x) = x^2$; for $d = 3$ we use $h_3(x) = (1/\sqrt{3})(x + x^2 - x^3)$; and for $d = 4$, $h_4(x) = (1/\sqrt{4})(-x + x^2 + x^3 - x^4)$. The functions for $d = 2, 3$ and 4 are shown in Figure 2.2 in the left three plots in blue. The constants that are best approximated by these functions are shown in orange. In the right plot the corresponding approximation error in $\|\cdot\|_\infty$ norm is plotted against d (top curve; orange). The blue curve in the right plot corresponds to the lower bound where we use $-1 = x_1 < \dots < x_d = 1$ with equidistant spacing to get full rank kernel matrices.




Figure 2.2. The three plots on the left show in blue polynomials of degree 2, 3 and 4 respectively. The orange lines correspond to the constant functions that are best approximated by these polynomials. The right most plot shows the corresponding approximation error in $\|\cdot\|_\infty$ (orange curve) and our lower bound on the approximation error (blue curve). Note that the approximation error is calculated for the three curves in the left plots and is only an upper bound for the best approximation error that can be attained.

2.3.6. Quantifying the width of the empirical convex set C_n


The above techniques can also be applied to the empirical convex set C_n . An easy way to do so is to identify the subspace spanned by C_n with a new RKHS. In particular, the subspace spanned by the empirical convex set C_n can be identified with an RKHS in a similar way to how we dealt with measures that attain values in a subspace in Section 4; see also Section 2.4.3 for a more detailed discussion. For an experiment $\omega \in \Omega$ let $S_\omega = \{X_1(\omega), \dots, X_n(\omega)\}$ be the support of the empirical measure for the realization ω . If k is our original kernel function then let k_ω be $k|_{S_\omega \times S_\omega}$ and let \mathcal{H}_ω be the corresponding RKHS. The empirical convex set C_n , as an element of \mathcal{H} , has then a corresponding convex set $C_\omega = \text{ch}\{k_\omega(X_i(\omega), \cdot) : i \leq n\}$ within \mathcal{H}_ω . For ease of notation fix an $\omega \in \Omega$ and let $x_1, \dots, x_n \in \mathcal{X}$ be $x_1 = X_1(\omega), \dots, x_n = X_n(\omega)$ for the rest of this section.

Importantly, there is a linear map $\psi : \mathcal{H} \rightarrow \mathcal{H}_\omega$ defined in the following way: if $h \in \mathcal{H}$ is of the form $\sum_{i=1}^n \alpha_i k(x_i, \cdot)$ for some $\alpha_i \in \mathbb{R}$ then let $\psi(h) = \sum_{i=1}^n \alpha_i k_S(x_i, \cdot)$. Also, let $U = \text{span}\{k(x_i, \cdot) : i \leq n\}$ be the subspace of \mathcal{H} corresponding to these functions h . For functions $g \in U^\perp$, define $\psi(g) = 0$, and extend ψ to all of \mathcal{H} by linearity. The function $\psi : \mathcal{H} \rightarrow \mathcal{H}_\omega$ defined in this way has the following properties: for all $g, h \in \text{span}\{k(X_1(\omega), \cdot), \dots, k(X_n(\omega), \cdot)\} \subset \mathcal{H}$ we have $\langle g, h \rangle = \langle \psi(g), \psi(h) \rangle_{\mathcal{H}_\omega}$ (this follows right away from the kernel expansion of g, h because k and k_S are equivalent on x_1, \dots, x_n) and $\psi(f) = 0$ if f is orthogonal to the subspace spanned by the data. In other words, ψ is a partial isometry between \mathcal{H} and \mathcal{H}_ω and an isometry between U , with the inherited inner product, and \mathcal{H}_ω .

Beside this natural link between \mathcal{H} and \mathcal{H}_ω there is also the linear map A that we considered in the proof of Lemma 2.3. We have to adapt the approach from Lemma 2.3 slightly to make use of it in this new context. First, observe that if \mathcal{H}_ω is d_ω -dimensional then we have the operators $A_\omega : \mathcal{H}_\omega \rightarrow \mathbb{R}^{d_\omega}$ defined by $A_\omega f = (f(x_{\iota(1)}), \dots, f(x_{\iota(d_\omega)}))^\top$ for a given injective function $\iota : \{1, \dots, d_\omega\} \rightarrow \{1, \dots, n\}$

such that the matrix $(k(x_{\iota(i)}, x_{\iota(j)}))_{i,j \leq d_\omega}$ has full rank. This dimension can obviously depend on ω and will always be upper bounded by the dimension $d_{\mathcal{H}}$ of \mathcal{H} . Consider now the kernel matrix $K_\omega = (k(x_{\iota(i)}, x_{\iota(j)}))_{i,j \leq d_\omega}$ and equip \mathbb{R}^{d_ω} with the inner product $\langle a, b \rangle_{K_\omega^{-1}} = a^\top K_\omega^{-1} b$, $a, b \in \mathbb{R}^{d_\omega}$. As in the proof of Lemma 2.3 it follows that $\langle g, h \rangle_{\mathcal{H}_\omega} = \langle Ag, Ah \rangle_{K_\omega^{-1}}$ for all $g, h \in \mathcal{H}_\omega$ and $A_\omega : \mathcal{H}_\omega \rightarrow \mathbb{R}^{d_\omega}$ is an isometry.

We have the following commutative diagram summarizing the relationship between the three spaces.



Furthermore, when $U = \text{span}\{k(x_1, \cdot), \dots, k(x_n, \cdot)\}$ is the subspace of \mathcal{H} induced by the data it follows that the following three spaces are isometric isomorphic

$$(U, \langle \cdot, \cdot \rangle) \cong (\mathcal{H}_\omega, \langle \cdot, \cdot \rangle_{\mathcal{H}_\omega}) \cong (\mathbb{R}^{d_\omega}, \langle \cdot, \cdot \rangle_{K_\omega^{-1}}).$$

In particular, $A_\omega \circ \psi$ is an isometry between U and \mathbb{R}^{d_ω} . This isometry has takes a simple form: let $h = \sum_{i=1}^n \alpha_i k(x_i, \cdot)$ then

$$(A_\omega \circ \psi)(h) = \sum_{i=1}^n \alpha_i (k(x_i, x_{\iota(1)}), \dots, k(x_i, x_{\iota(d_\omega)}))^\top = (h(x_{\iota(1)}), \dots, h(x_{\iota(d_\omega)}))^\top.$$

This relation allows us to apply the techniques we developed for measuring the size of C to the empirical convex set C_n . For example, if \mathcal{H}_ω does not contain constant functions then using the kernel k_S^+ and denoting the corresponding RKHS by \mathcal{H}_ω^+ , we can lower bound the width of C_ω . The RKHS \mathcal{H}_ω^+ has dimension $d_\omega + 1$ and there exists an injection $\iota : \{1, \dots, d_\omega + 1\} \rightarrow \{1, \dots, n\}$ such that $k_S^+(x_{\iota(1)}, \cdot), \dots, k_S^+(x_{\iota(d_\omega+1)}, \cdot)$ are linearly independent. Then, as above, $A_\omega^+ : \mathcal{H}_\omega^+ \rightarrow \mathbb{R}^{d_\omega+1}$ defined by $A_\omega^+(h) = (h(x_{\iota(1)}), \dots, h(x_{\iota(d_\omega+1)}))^\top$ is an isometry between \mathcal{H}_ω^+ and $\mathbb{R}^{d_\omega+1}$ when the latter is equipped with the inner product $\langle a, b \rangle_{(K_\omega^+)^{-1}}$, for all $a, b \in \mathbb{R}^{d_\omega+1}$ and K_ω^+ is the kernel matrix corresponding to the points $x_{\iota(1)}, \dots, x_{\iota(d_\omega+1)}$. From this we can infer a lower bound on the width of C_ω within \mathcal{H}_ω . Alternatively, we can apply directly Proposition 2.10 to \mathcal{H}_ω to get this lower bound. Since we have an isometry between U and \mathcal{H}_ω these lower bounds translate directly to lower bounds on the width of C_n within U .

There is another point worth noting. The lower bound on the width of C_n depends on the choice of ι . Finding the subset of points x_1, \dots, x_n that maximizes this lower bound seems like a hard problem. Therefore one might wonder if there is a simpler way to *optimize the lower bound*. In particular, there seems hope to get the largest smallest eigenvalue λ_{d_ω} when using the full kernel matrix. To that end, let $K_\omega^* = (k(x_i, x_j))_{i,j \leq n}$ be the kernel matrix corresponding to all the data. Since the subspace spanned by the data has dimension d_ω it follows that there are exactly d_ω non-zero eigenvalues $\lambda_1^*, \dots, \lambda_{d_\omega}^*$. There is a useful interplay between K_ω^* and the following linear operator $A_\omega^* : \mathcal{H}_\omega \rightarrow \mathbb{R}^n$ given by $A_\omega^*(h) = (h(x_1), \dots, h(x_n))^\top$. First note that for $h = \sum_{i=1}^n \alpha_i k(x_i, \cdot)$, with suitable $\alpha_i \in \mathbb{R}$, we have that $A_\omega^*(h) = K_\omega^* \alpha$. Also observe that A_ω^* is injective because if $A_\omega^*(f) = (f(x_1), \dots, f(x_n))^\top = (g(x_1), \dots, g(x_n))^\top = A_\omega^*(g)$ for two functions $f, g \in \mathcal{H}_\omega$, $f : S \rightarrow \mathbb{R}, g : S \rightarrow \mathbb{R}$, then

f and g are equal on S and are therefore the same function. While A_ω^\star is injective there are generally for a given $h \in \mathcal{H}_\omega$ many $\alpha \in \mathbb{R}^n$ such that $A_\omega^\star(h) = K_\omega^\star \alpha$ and K_ω^\star is not invertible. Therefore, consider the Moore–Penrose pseudo-inverse $(K_\omega^\star)^\dagger$, and observe that with $\alpha_h^\star = (K_\omega^\star)^\dagger A_\omega^\star(h)$ we get $K_\omega^\star \alpha_h^\star = K_\omega^\star (K_\omega^\star)^\dagger A_\omega^\star(h) = A_\omega^\star(h)$ since $A_\omega^\star(h)$ lies in the range of K_ω^\star [CM09, Definition 1.1.2(a)]. In particular, for $f, g \in \mathcal{H}_\omega$,

$$\langle f, g \rangle = (\alpha_f^\star)^\top K_\omega^\star \alpha_g^\star = (A_\omega^\star(f))^\top (K_\omega^\star)^\dagger K_\omega^\star (K_\omega^\star)^\dagger A_\omega^\star(g) = (A_\omega^\star(f))^\top (K_\omega^\star)^\dagger A_\omega^\star(g).$$

From this relation we get a lower bound on the supremum norm of a function $h \in \mathcal{H}_\omega$,

$$\begin{aligned} \|h\|_{\mathcal{H}_\omega}^2 &= (A_\omega^\star(f))^\top (K_\omega^\star)^\dagger A_\omega^\star(f) = \text{tr}((K_\omega^\star)^\dagger A_\omega^\star(f) (A_\omega^\star(f))^\top) \\ &\leq \|(K_\omega^\star)^\dagger\|_{\text{op}} (A_\omega^\star(f))^\top A_\omega^\star(f) \leq n \|(K_\omega^\star)^\dagger\|_{\text{op}} \|h\|_\infty^2. \end{aligned}$$

The term $\|(K_\omega^\star)^\dagger\|_{\text{op}}$ is equal to $1/\lambda_{d_\omega}^\star$ but, unfortunately, instead of the constant d_ω we have now the constant n .

In Proposition 2.10 seemingly no price had to be paid for the dimension of \mathcal{H} when using a *Mercer decomposition*. Since intuitively K_ω^\star is closely related to the integral operator that appears in Mercer’s theorem when the underlying measure is the empirical measure P_n one might wonder if the constant d_ω , or n , can be removed by following that route. Unfortunately, this approach does not remove the constant: consider the integral operator

$$(T_\omega f)(y) = \int f(x) k(x, y) dP_n(x),$$

for $f \in \mathcal{L}^2(S, P_n)$ where $S = \{x_1, \dots, x_n\}$. Observe that $\mathcal{L}^2(S, P_n)$ is the same set of functions as \mathcal{H}_ω but the \mathcal{L}^2 -inner product does not have to be equal to the inner product of \mathcal{H}_ω . Then for $f \in \mathcal{L}^2(S, P_n) = \mathcal{H}_\omega$ and $j \leq n$,

$$(T_\omega f)(x_j) = \frac{1}{n} \sum_{i=1}^n f(x_i) k(x_i, x_j) = n^{-1} ((A_\omega^\star(f))^\top K_\omega^\star)_j$$

and for $f, g \in \mathcal{L}^2(S, P_n)$,

$$\langle T_\omega f, g \rangle_2 = n^{-2} (A_\omega^\star(f))^\top K_\omega^\star A_\omega^\star(g).$$

The eigenfunctions of T_ω are closely related to the eigenvectors of K_ω^\star . Let $q_1, \dots, q_n \in \mathbb{R}^n$ be the eigenvectors of K_ω^\star and let $\lambda_1^\star, \dots, \lambda_n^\star$ be the corresponding eigenvalues. Observe that q_1, \dots, q_{d_ω} lie in the range of A_ω^\star since for any $i \leq d_\omega$, $q_i = \lambda_i^\star K_\omega^\star q_i = A_\omega^\star(e_i)$ where $e_i = \sum_{j=1}^n \lambda_i^\star (q_i)_j k(x_j, \cdot)$. Also, it follows directly that $n^{1/2} e_1, \dots, n^{1/2} e_{d_\omega}$ are an orthonormal basis in $\mathcal{L}^2(S, P_n)$ as $\langle n^{1/2} e_i, n^{1/2} e_j \rangle_2 = A_\omega^\star(e_i)^\top A_\omega^\star(e_j) = q_i^\top q_j$ and $\mathcal{L}^2(S, P_n)$ is d_ω -dimensional. Furthermore, $n^{1/2} e_1, \dots, n^{1/2} e_{d_\omega}$ are the eigenfunctions of T_ω ,

$$n \langle T_\omega e_i, e_j \rangle_2 = n^{-1} q_i^\top K_\omega^\star q_j = (\lambda_j^\star/n) q_i^\top q_j$$

and the corresponding eigenvalues of T_ω are $\lambda_1^\star/n, \dots, \lambda_{d_\omega}^\star/n$.

To summarize, we discussed two approaches in this section to get a lower bound on the width of the empirical convex set. The first approach uses a selection of d_ω sample points and the eigenvalues of the kernel matrix corresponding to these points. It is unclear if there is an efficient way to optimize over this subset selection.

The second approach uses instead the full kernel matrix, which sidesteps the problem of selecting sample points, and leads to a larger eigenvalue but then the constant degrades significantly if $n \gg d_\omega$. There is a third way which “interpolates” between the two approaches. For instance, it might be reasonable to use $2d_\omega$ many points to help with the subset selection problem while keeping the constant small.

There are multiple hurdles to using these approaches in practice. First off, it is not just the width that needs to be controlled but also how centered \mathbf{m}_n lies within C_n . Furthermore, the current approach is only applicable in the small sample regime since we need the smallest eigenvalue of the kernel matrix to control the width. This eigenvalue can be computed by applying the power iteration method. The power iteration returns the largest absolute value of a matrix. A standard way to find λ_d is the following: apply the power iteration to K to find λ_1 ; then move to matrix $B = K - \lambda_1 I$, which is negative definite, and apply the power iteration to get $\lambda_d - \lambda_1$. Each iteration of the power iteration relies on a multiplication of an $n \times n$ matrix with a vector. This makes this method prohibitively costly to apply in the large sample regime. We come back to these issues in Section 4 where we study, among other things, algorithms which split the data into small batches. In such settings it becomes possible to control the width of the empirical convex sets that correspond to the small batches of data.

2.4. Locating \mathbf{m}

For various convex approximation methods the distance from \mathbf{m} to the boundary of the convex set characterizes the rate of convergence: the larger the distance the faster the rate of convergence. A crude way to measure the distance is to consider the largest ball that fits within the convex set around \mathbf{m} . Having a closed ball of size $\delta > 0$ around \mathbf{m} in C is equivalent to

$$(2.12) \quad \inf_{\|h\|=1} \sup_{x \in \mathcal{X}} \langle h, k(x, \cdot) - \mathbf{m} \rangle \geq \delta,$$

and similarly for affine subspaces. This can be seen in the following way: clearly when there exists a closed ball around \mathbf{m} with the stated properties then for any $h, \|h\| = 1$, some extreme of the convex set must fulfill (2.12). On the other hand, $C - \mathbf{m}$ is equal to the intersection of the closed half-spaces tangent to it [Roc97, Theorem 18.8]. To each of these half-spaces there exists a normal $h \in \mathcal{H}, h \neq 0$, and an $\alpha_h \in \mathbb{R}$ such that $\langle g, h \rangle \leq \alpha_h$ whenever g lies in the half-space. In particular, for any such normal $\langle g, h \rangle \leq \alpha_h$ whenever $g \in C - \mathbf{m}$. Without loss of generality we can assume that the normals have norm one and by assuming that (2.12) holds we know that for any such normal $h, \alpha_h \geq \delta$. If there would not exist a ball of size δ around \mathbf{m} in C then there would be an element $g \notin C - \mathbf{m}, \|g\| \leq \delta$. But then $\langle h, g \rangle \leq \delta$ for all $h \in \mathcal{H}, \|h\| = 1$, and g would lie in the intersection of the half-spaces and then also in $C - \mathbf{m}$ due to [Roc97, Theorem 18.8] which cannot be.

In the previous section we quantified the width of the set C in direction h , i.e. $\text{width}_h(C) = \sup_{x \in \mathcal{X}} h(x) - \inf_{x \in \mathcal{X}} h(x)$. The width tells us how large a ball around \mathbf{m} can be in the ideal case where \mathbf{m} lies centered within C , however, we do not know

how centered \mathbf{m} lies within C . Obviously, \mathbf{m} can lie in the boundary for instance when $\mathbf{m} = k(x, \cdot)$ and $k(x, \cdot)$ is an extreme of C , and assumptions on the distribution of the data are needed to guarantee the existence of a ball around \mathbf{m} . Our aim in this section is to show how natural assumptions on the probability distribution translate to statements of how centered \mathbf{m} lies. In the following, we are studying two such conditions: (1) a lower bound on the density of the law of X_1 together with a Lipschitz condition on $\phi : \mathcal{X} \rightarrow \mathcal{H}$; (2) an assumption on the covariance operator $\tilde{\mathcal{C}}_c$. We finish this section with a look at the case where the law of X_1 does not have full support in \mathcal{X} .

Before looking at these conditions let us add a short comment about the relation between the extremes of C and \mathbf{m} . For the convex set $C = \text{cch}\{\phi(x) : x \in \mathcal{X}\}$ the extremes of C which are close to \mathbf{m} are images under ϕ of points x which lie close to each other. In detail, consider a kernel function with $k(x, x) = 1$ for all $x \in \mathcal{X}$. Whenever $\|\phi(x) - \mathbf{m}\| < \epsilon$ and $\|\phi(y) - \mathbf{m}\| < \epsilon$ for some $\epsilon > 0$ then $4\epsilon^2 \geq \|\phi(x) - \phi(y)\|^2 = 2(1 - k(x, y))$. In other words, if there exists an extreme $\phi(x_0)$ of C that lies ϵ close to \mathbf{m} then all the extremes of C that are ϵ close to \mathbf{m} are contained in

$$\phi[\{y : k(x_0, y) \geq 1 - 2\epsilon^2\}].$$

Obviously, the case that we have an extreme $\phi(x_0)$ close to \mathbf{m} is rare since this means that for all functions $h \in \mathcal{H}$, $\|h\| \leq 1$, the expected value $E(h(X)) \approx h(x_0)$ and \mathcal{H} cannot distinguish between P and a probability measure that puts mass one on x_0 .

2.4.1. Assumptions on the density

In [BLJO12] it was observed that when the probability measure corresponding to \mathbf{m} has a density on X which is bounded away from 0 and \mathcal{H} is finite dimensional then it is at least guaranteed that some open ball exists around \mathbf{m} in C . This result can be strengthened and turned into a quantitative statement by using a simple observation.

Consider first the Lebesgue integral on \mathbb{R} . If we have a (non-atomic) probability measure on \mathbb{R} which has a mean value of 0 and there exists some measurable set B with $\inf B \geq \epsilon$ and $P(B) > 0$, then there will be probability mass on the negative axis to counter the “pull” from B since otherwise

$$0 = \int_{\mathbb{R}} x dP = \int_{[0, \infty)} x dP \geq \int_B x dP \geq \epsilon P(B) > 0.$$

This argument can also be applied to \mathbf{m} . Consider the set $X = [0, 1]$, an RKHS \mathcal{H} with continuous kernel function $k(x, y)$ and assume that $k(x, \cdot) \in \mathcal{L}^1(P; \mathcal{H})$ with Bochner-integral \mathbf{m} and the probability measure P has a density function that is bounded away from 0. For every $y \in X$ with $k(y, \cdot) - \mathbf{m} \neq 0$ there exists an $x \in X$ such that $\langle k(y, \cdot) - \mathbf{m}, k(x, \cdot) - \mathbf{m} \rangle < 0$. Otherwise, let $\epsilon = \|k(y, \cdot) - \mathbf{m}\|^2/2$ then $B = \{x : \langle k(y, \cdot) - \mathbf{m}, k(x, \cdot) - \mathbf{m} \rangle > \epsilon\}$ is non-empty as $y \in B$ and contains

an open interval I of X , with $P(I) > 0$. Hence, $P(B) > 0$ and because $\|\mathbf{m}\|^2 = \int_X \langle \mathbf{m}, k(x, \cdot) \rangle dP(x)$,

$$0 = \int_X \langle k(y, \cdot) - \mathbf{m}, k(x, \cdot) - \mathbf{m} \rangle dP(x) \geq \int_B \langle k(y, \cdot) - \mathbf{m}, k(x, \cdot) - \mathbf{m} \rangle dP(x) \geq \epsilon P(B) > 0.$$

This implies that we have on both sides of \mathbf{m} (with respect to the direction $k(y, \cdot) - \mathbf{m}$) elements of $\text{cch}\{k(x, \cdot) : x \in X\} = C$.

To provide lower bounds on the radius of a ball around \mathbf{m} in C we need more. Ideally, we like to have assumptions on the kernel function and the measure which guarantee the existence of some strictly positive function $\psi : (0, \infty) \rightarrow (0, \infty)$ such that for any $h \in \mathcal{H}$, $\|h\| \leq 1$, $x \in \mathcal{X}$, if $\langle h, k(x, \cdot) - \mathbf{m} \rangle > 0$ then

$$\inf_{y \in \mathcal{X}} \langle h, k(y, \cdot) - \mathbf{m} \rangle \leq -\psi(\langle h, k(x, \cdot) - \mathbf{m} \rangle).$$

Under a Lipschitz assumption on the functions in \mathcal{H} we can provide such a function ψ . The Lipschitz assumption we are using is that any $h \in \mathcal{H}$ fulfills

$$(2.13) \quad \sup_{x \neq x'} \frac{|h(x) - h(x')|}{\|x - x'\|} \leq \|h\|L,$$

where $L > 0$ is the Lipschitz-constant. When the space \mathcal{X} is a compact subset of \mathbb{R} this Lipschitz assumption is often fulfilled. For instance, when a polynomial or Gaussian kernel is used. In fact, whenever we have a well behaved domain \mathcal{X} like $[0, 1]^d$, $h \upharpoonright \text{int } \mathcal{X} \in C^1(\text{int } \mathcal{X})$, $h \in C(\mathcal{X})$ and $\|D_x h\|_{\text{op}} \leq \|h\|L$ (compare to [SC08, Corollary 4.36]) then the condition is fulfilled.

In the following, β_d denotes d -dimensional Lebesgue measure of the unit ball in \mathbb{R}^d and μ_d denotes Lebesgue measure.

PROPOSITION 2.12. — *Let $\mathcal{X} = [0, 1]^d$ and let \mathcal{H} be an RKHS such that for all $h \in \mathcal{H}$ and $x, x' \in \mathcal{X}$, $|h(x) - h(x')| \leq L\|h\|\|x - x'\|$. Furthermore, let P be a probability measure on \mathcal{X} and assume that P has a density p with $\inf_{x \in \mathcal{X}} p(x) \geq c > 0$. Then for any $h \in \mathcal{H}$, $\|h\| \leq 1$,*

$$\max_{y \in \mathcal{X}} \langle -h, k(y, \cdot) - \mathbf{m} \rangle \geq \frac{c\gamma^{d+1}}{(d+1)(2L)^d} \beta_d.$$

whenever there exists an $x \in \mathcal{X}$ such that $\langle h, k(x, \cdot) - \mathbf{m} \rangle \geq \gamma > 0$ and $\gamma/L \leq 1$.

Proof. — Fix any h in the unit ball of \mathcal{H} and let $f(x) = \langle h, k(x, \cdot) - \mathbf{m} \rangle$. Let $x^* \in \mathcal{X}$ be a point at which $f(x^*) = \langle h, k(x^*, \cdot) - \mathbf{m} \rangle \geq \gamma$. The function f is also Lipschitz continuous with Lipschitz-constant $\|h\|L \leq L$ and f is therefore non-negative on the set $A = \{y : \|y - x^*\| \leq \gamma/L, y \in \mathcal{X}\}$. Let $B = \{y : \|y\| \leq \gamma/L, y \in \mathcal{X}\}$ then $P(A) \geq c\mu_d(A) \geq c\mu_d(B)$ because B , being centered at 0, minimizes the volume of the intersection of $\mathcal{X} = [0, 1]^d$ with a ball of radius γ/L with center in \mathcal{X} . Furthermore, $\mu_d(B) = (\gamma/2L)^d \beta_d$; this is the volume of a d -dimensional ball

of radius γ/L scaled by 2^{-d} . Now, integrating over A and using [Fre01, 265G, 265H] again

$$\begin{aligned} \int_A f(x) dP(x) &\geq \int_{A-x^*} c(f(x^*) - L\|x\|) dx \geq c \int_B f(x^*) - L\|x\| dx \\ &\geq \frac{c\gamma^{d+1}}{(2L)^d} \beta_d - cL2^{-d} \frac{d}{d+1} \frac{\gamma^{d+1}}{L^{d+1}} \beta_d = \frac{c\gamma^{d+1}}{(2L)^d} \beta_d \left(1 - \frac{d}{d+1}\right). \end{aligned}$$

Since $\int_{\mathcal{X}} f(x) dP(x) = 0$ there must be a point $y \in \mathcal{X}$ such that

$$f(y) \leq -\frac{c\gamma^{d+1}}{(2L)^d} \beta_d \left(\frac{1}{d+1}\right). \quad \square$$

Under the conditions of the proposition we can state a lower bound on the size of a ball included in C around \mathbf{m} . Let $h \in \mathcal{H}$, $\|h\| = 1$, and assume w.l.o.g. that $s := \sup_{x \in \mathcal{X}} \langle h, k(x, \cdot) - \mathbf{m} \rangle \geq -\inf_{x \in \mathcal{X}} \langle h, k(x, \cdot) - \mathbf{m} \rangle =: i$. Then $s \geq (1/2) \text{width}_h(C)$ and

$$(2.14) \quad i \geq \frac{cs^{d+1}}{(d+1)(2L)^d} \beta_d \geq \frac{c((1/2) \text{width}_h(C))^{d+1}}{(d+1)(2L)^d} \beta_d.$$

If we have a lower bound b on $\text{width}_h(C)$ for all such h then we can conclude that there exists a ball of radius

$$\delta = \min \left\{ (b/2), \frac{c(b/2)^{d+1} \beta_d}{(d+1)(2L)^d} \right\}$$

around \mathbf{m} in C .

2.4.2. Assumptions on the covariance operator

Let us start with a useful relationship between $\mathbf{1}$ and \mathbf{m} whenever $\mathbf{1}$ lies in the RKHS. For any measure P , with expectation \mathbf{m} we have that $\langle \mathbf{1}, \mathbf{m} \rangle = 1$. Also, for any $x \in \mathcal{X}$, $\langle k(x, \cdot), \mathbf{1} \rangle = 1$ and C lies within the closed affine subspace $\{h \in \mathcal{H} : \langle h, \mathbf{1} \rangle = 1\}$, where closure follows from $h \mapsto \langle h, \mathbf{1} \rangle$ being a continuous function and $\{1\}$ being closed. Also, since $1 = \langle \mathbf{m}, \mathbf{1} \rangle \leq \|\mathbf{1}\| \|\mathbf{m}\|$ it has to hold that $\|\mathbf{m}\| \geq 1/\|\mathbf{1}\|$. An upper bound on $\|\mathbf{1}\|$ is therefore giving us a lower bound on $\|\mathbf{m}\|$. For instance, we can get a lower bound on $\|\mathbf{m}\|$ by inspecting the kernel function k of the RKHS in the sense that

$$(2.15) \quad \|\mathbf{m}\| \geq 1/\inf\{c : c^2 k \succeq \mathbf{1} \otimes \mathbf{1}\}.$$

One might be tempted to move to \mathcal{H}^+ whenever $\mathbf{1}$ does not lie in the RKHS \mathcal{H} ; recall that \mathcal{H}^+ has the kernel $k + \mathbf{1} \otimes \mathbf{1}$. Since \mathcal{H} can be regarded as a subspace of \mathcal{H}^+ we have that $\|\mathbf{m}\| = \|\mathbf{m}\|_+$; however, only for $h \in \mathcal{H}$ do we have that $\int h dP = \langle \mathbf{m}, h \rangle_+$ since $\mathbf{1} \in \mathcal{H}^\perp$ by construction. But there is then an element $\mathbf{m}^+ \in \mathcal{H}_+$ for which $P_{\mathcal{H}} \mathbf{m}^+ = \mathbf{m}$ and $(I - P_{\mathcal{H}}) \mathbf{m}^+ = \langle \mathbf{1}, \mathbf{m}^+ \rangle = 1$, where $P_{\mathcal{H}}$ is the orthogonal projection onto the subspace \mathcal{H} and I the identity operator. Now, $1 = \langle \mathbf{1}, \mathbf{m}^+ \rangle_+ \leq \|\mathbf{m}^+\|_+ = 1 + \|\mathbf{m}\|$ and we only learn from this the trivial fact that $0 \leq \|\mathbf{m}\|$.

An alternative approach gives us more insight. Whenever there is a function $\tilde{\mathbf{1}} \in \mathcal{H}$ such that $\|\mathbf{1} - \tilde{\mathbf{1}}\|_\infty \leq \beta < 1$ then $1 - \beta \leq \langle \mathbf{m}, \tilde{\mathbf{1}} \rangle \leq 1 + \beta$ and

$$(2.16) \quad \|\mathbf{m}\| \geq (1 - \beta)/\|\tilde{\mathbf{1}}\| = (1 - \beta)/\inf\{c : c^2 k \succeq \tilde{\mathbf{1}} \otimes \tilde{\mathbf{1}}\}.$$

In the following we will make use of the covariance operator $\tilde{\mathfrak{C}}$ (see Section 1.9) to determine the location of \mathfrak{m} . Before exploring the relation between the covariance operator and the location of \mathfrak{m} we note the following adaptation of the above discussion. If $\mathbf{1} \in \mathcal{H}$ then $\langle \tilde{\mathfrak{C}}\mathbf{1}, \mathbf{1} \rangle = 1$ and whenever $\tilde{\mathbf{1}} \in \mathcal{H}$ fulfills $\|\mathbf{1} - \tilde{\mathbf{1}}\|_\infty \leq \beta < 1$ then

$$(2.17) \quad \langle \tilde{\mathfrak{C}}\mathbf{1}, \mathbf{1} \rangle \geq (1 - \beta)^2 \|\tilde{\mathbf{1}}\| = (1 - \beta)^2 / \inf \{c : c^2 k \succeq \tilde{\mathbf{1}} \otimes \tilde{\mathbf{1}}\}.$$

Coming now to the problem of locating \mathfrak{m} within C we can take note of the following fundamental relationship. Whenever \mathcal{H} is d -dimensional, $\|k\|_\infty < \infty$ and $\tilde{\mathfrak{C}}$ has an eigen-decomposition with smallest eigenvalue $\bar{\lambda}_d$ then for any $h \in \mathcal{H}$, $\|h\| = 1$, with $\langle h, \mathfrak{m} \rangle = 0$, it follows that $\int h(x) dP = 0$ and

$$\int (\langle h, k(x, \cdot) - \mathfrak{m} \rangle)^2 dP = \int h^2(x) dP \geq \bar{\lambda}_d.$$

Since $\sup_{x \in \mathcal{X}} |h(x)| \leq \|k\|_\infty^{1/2}$ and $\int h(x) dP = 0$, a short calculation shows that

$$r(\mathfrak{m}, h) \geq \frac{\bar{\lambda}_d}{\|k\|_\infty^{1/2}},$$

where with $h \in \mathcal{H}$,

$$r(\mathfrak{m}, h) := \left(\sup_{x \in \mathcal{X}} \langle h, k(x, \cdot) - \mathfrak{m} \rangle \right) \wedge \left(- \inf_{x \in \mathcal{X}} \langle h, k(x, \cdot) - \mathfrak{m} \rangle \right).$$

In detail, it is sufficient to consider the case of a discrete measure where with probability p the function $\langle h, k(x, \cdot) - \mathfrak{m} \rangle$ attains value a and with probability $1 - p$ attains value b , with $a < 0 < b$, and $a^2 p + b^2(1 - p) = \bar{\lambda}_d$. The condition $\langle h, \mathfrak{m} \rangle = 0$ then implies that $ap + b(1 - p) = 0$. For a particular value of a we get that $p = b/(b - a)$ and $b = \bar{\lambda}_d/(-a)$. The value b is minimized by maximizing $-a$ but $-a = \langle h, k(x, \cdot) \rangle$ for some $x \in \mathcal{X}$ and $-a = |a| \leq \|k\|_\infty^{1/2}$. By symmetry we get that $(-a) \wedge b \geq \bar{\lambda}_d/\|k\|_\infty^{1/2}$.

The remaining direction we have to take care of is $h^* = \mathfrak{m}/\|\mathfrak{m}\|$. The distance of \mathfrak{m} to the boundary in direction h^* can be lower bounded away from zero when the smallest eigenvalue of the covariance operator is sufficiently large since

$$\int (\langle \mathfrak{m}, k(x, \cdot) - \mathfrak{m} \rangle)^2 dP = \int (\langle \mathfrak{m}, k(x, \cdot) \rangle)^2 dP - \|\mathfrak{m}\|^4 = \langle \tilde{\mathfrak{C}}\mathfrak{m}, \mathfrak{m} \rangle - \|\mathfrak{m}\|^4$$

and

$$\int (\langle h^*, k(x, \cdot) - \mathfrak{m} \rangle)^2 dP = \langle \tilde{\mathfrak{C}}h^*, h^* \rangle - \|\mathfrak{m}\|^2 \geq \bar{\lambda}_d - \|\mathfrak{m}\|^2.$$

Also, $|\langle h^*, k(x, \cdot) - \mathfrak{m} \rangle| \leq \|k(x, \cdot)\| + \|\mathfrak{m}\| \leq 2\|k\|_\infty^{1/2}$ and

$$r(\mathfrak{m}, h^*) \geq \frac{\bar{\lambda}_d - \|\mathfrak{m}\|^2}{2\|k\|_\infty^{1/2}}.$$

For this approach to yield a useful bound $\bar{\lambda}_d$ has to be strictly greater than $\|\mathfrak{m}\|^2$. However, $\bar{\lambda}_d$ can even be smaller than $\|\mathfrak{m}\|^2$. A better bound can be gained by using an eigen-decomposition of $\tilde{\mathfrak{C}}_c = \tilde{\mathfrak{C}} - \mathfrak{m} \otimes \mathfrak{m}$. In the following let $\bar{\lambda}_1, \bar{\lambda}_2, \dots$ be the eigenvalues of $\tilde{\mathfrak{C}}_c$ then by the same argument as for $\tilde{\mathfrak{C}}$, whenever $h \in \mathcal{H}$, $\|h\| = 1$, is such that $\langle h, \mathfrak{m} \rangle = 0$, it follows that $\int (\langle h, k(x, \cdot) - \mathfrak{m} \rangle)^2 dP = \langle \tilde{\mathfrak{C}}_c h, h \rangle \geq \bar{\lambda}_d$, where

$\bar{\lambda}_d$ is the smallest eigenvalue of $\tilde{\mathfrak{C}}_c$, and $r(\mathbf{m}, h) \geq \bar{\lambda}_d / \|k\|_\infty^{1/2}$. For $h^* = \mathbf{m} / \|\mathbf{m}\|$ we get now that

$$\int (\langle h^*, k(x, \cdot) - \mathbf{m} \rangle)^2 dP = \langle \tilde{\mathfrak{C}}_c h^*, h^* \rangle \geq \bar{\lambda}_d$$

and $r(\mathbf{m}, h^*) \geq \bar{\lambda}_d / 2 \|k\|_\infty^{1/2}$. Also notice that when $\mathbf{1} \in \mathcal{H}$ then $\langle \tilde{\mathfrak{C}}_c \mathbf{1}, \mathbf{1} \rangle = 0$ and $\bar{\lambda}_d = 0$. We summarize these findings in the following proposition.

PROPOSITION 2.13. — *Let $(\mathcal{X}, \mathcal{A}, P)$ be some probability space with measurable kernel function k defined on it and such that the corresponding RKHS \mathcal{H} has dimension $d < \infty$. Furthermore, assume that $\|k\|_\infty < \infty$ and that the centered covariance operator $\tilde{\mathfrak{C}}_c$ has an eigen-decomposition with smallest eigenvalue $\bar{\lambda}_d > 0$. Then $\mathbf{1} \notin \mathcal{H}$ and for any $h \in \mathcal{H}$, $\|h\| = 1$,*

$$r(\mathbf{m}, h) \geq \frac{\bar{\lambda}_d}{2 \|k\|_\infty^{1/2}}.$$

The *dimension dependence* is in this result not as obvious as in Proposition 2.12, but observe that when \mathcal{X} is the d -dimensional unit sphere, $d = 2l + 1$ for some $l \in \mathbb{N}$, P is the Lebesgue measure restricted to \mathcal{X} and normalized, $k(x, y) = \langle x, y \rangle_{\mathbb{R}^d}$ for any $x, y \in \mathcal{X}$ and \mathcal{H} is the corresponding RKHS which is of dimension d , then any $h \in \mathcal{H}$, $\|h\| = 1$, is of the form $\langle x, \cdot \rangle_{\mathbb{R}^d}$ for some $x \in \mathcal{X}$, $\|x\|_{\mathbb{R}^d} = 1$ and for such an h

$$\begin{aligned} \langle \tilde{\mathfrak{C}}_c h, h \rangle &= \beta_d^{-1} \int h^2(y) dP(y) = \beta_d^{-1} \int \langle x, y \rangle^2 dP(y) \\ &= \beta_d^{-1} \int_{-1}^1 \langle x, \tilde{y}x \rangle^2 \mu_{d-1}(B_{d-1}(\sqrt{1 - \tilde{y}^2})) d\mu(\tilde{y}) \\ &= \beta_{d-1} \beta_d^{-1} \int_{-1}^1 \tilde{y}^2 (1 - \tilde{y}^2)^l d\mu(\tilde{y}) = \frac{2^l l! \beta_{d-1} \beta_d^{-1}}{\prod_{i=1}^l (2i + 1)} \int_{-1}^1 \tilde{y}^{2(l+1)} d\mu(\tilde{y}) \\ &= \frac{2^l l! \beta_{d-1} \beta_d^{-1}}{(\prod_{i=1}^l (2i + 1))(l + 3/2)} = \frac{(2l + 1)!}{l! 2^{l+1} (\prod_{i=1}^l (2i + 1))(l + 3/2)} = \frac{1}{d + 2}, \end{aligned}$$

where μ_{d-1} denotes here the $d - 1$ -dimensional Lebesgue measure and β_{d-1} the Lebesgue measure of the $d - 1$ -dimensional unit sphere. Hence, the eigenvalues of $\tilde{\mathfrak{C}}$ shrink to zero as the dimension d increases.

If $\mathbf{1} \in \mathcal{H}$ then we get a similar result with $\bar{\lambda}_d$ being replaced by $\bar{\lambda}_{d-1}$.

COROLLARY 2.14. — *Let $(\mathcal{X}, \mathcal{A}, P)$ be some probability space with measurable kernel function k defined on it and such that the corresponding RKHS \mathcal{H} has dimension $d < \infty$ and $\mathbf{1} \in \mathcal{H}$. Furthermore, assume that $\|k\|_\infty < \infty$ and that the centered covariance operator $\tilde{\mathfrak{C}}_c$ has an eigen-decomposition with eigenvalue $\bar{\lambda}_{d-1} > 0$. Then for any $h \in \mathcal{H}$, $\|h\| = 1$, $\langle h, \mathbf{1} \rangle = 0$,*

$$r(\mathbf{m}, h) \geq \frac{\bar{\lambda}_{d-1}}{2 \|k\|_\infty^{1/2}}.$$

Proof. — First note that $\langle \tilde{\mathfrak{C}}_c \mathbf{1}, \mathbf{1} \rangle = 0$ and for any $h \in \mathcal{H}$ s.t. $\langle h, \mathbf{1} \rangle = 0$, $\langle \tilde{\mathfrak{C}}_c \mathbf{1}, h \rangle = \langle \tilde{\mathfrak{C}}_c h, \mathbf{1} \rangle = E(h) - E(h) = 0$. In particular, $\tilde{\mathfrak{C}}_c \mathbf{1} = 0$ and for any $h \in \mathcal{H}$, $\|h\| = 1$, $\langle h, \mathbf{1} \rangle = 0$, $\langle \tilde{\mathfrak{C}}_c h, h \rangle \geq \bar{\lambda}_{d-1}$. Now, for $h \in \mathcal{H}$, $\|h\| = 1$, $\langle h, \mathbf{1} \rangle = 0$,

$$\int (\langle h, k(x, \cdot) - \mathbf{m} \rangle)^2 dP = \langle \tilde{\mathfrak{C}}_c h, h \rangle \geq \bar{\lambda}_{d-1}$$

and the lower bound follows by the same argument as in Proposition 2.13. □

2.4.3. Data attaining values in a subset

Unless $h \in \mathcal{H}$ is constant on the support of P it holds that $E((h - E(h))^2) > 0$ and by the above arguments it follows that $r(\mathbf{m}, h) > 0$. But \mathbf{m} can certainly lie in the boundary, for example, when $\mathbf{m} = k(x, \cdot)$ and $k(x, \cdot)$ is an extreme of C . Therefore, there must be some direction h in which $r(\mathbf{m}, h) = 0$. The point is that when \mathbf{m} is an extreme or lies in a face of the convex set C and this face has extremes $\{k(x, \cdot) : x \in S \subsetneq \mathcal{X}\}$ then $P(X \in \mathcal{X} \setminus S) = 0$ has to hold. Furthermore, there then exists a normal h^* to this face and h^* is constant on the support of P which implies that $\langle \tilde{\mathfrak{C}}_c h^*, h^* \rangle = 0$.

This observation suggests that \mathbf{m} will either be an extreme or there will be a ball in an affine subset of \mathcal{H} around \mathbf{m} within a face of the convex set. Furthermore, we can hope that this affine subspace is directly related to the non-zero eigenvalues of $\tilde{\mathfrak{C}}_c$ and that these eigenvalues characterize a lower bound on the width of this ball. Alternatively, it is natural to consider the space $\mathcal{H}_S = \{h|_S : h \in \mathcal{H}\}$ where S is the support of P [Fre01, Definition 411N]. The space \mathcal{H}_S is again an RKHS with kernel $k_S = k|_{S \times S}$ and norm $\|h\|_{k_S} = \inf\{\|u\| : u|_S = h, u \in \mathcal{H}\}$ [PR16, Corollary 5.8]. In the proposition below we show that the covariance operator $\tilde{\mathfrak{C}}_c^S$ corresponding to P and \mathcal{H}_S characterizes the ball around \mathbf{m} within the affine subspace spanned by C .

PROPOSITION 2.15. — *Let $(\mathcal{X}, \mathcal{A}, P)$ be a probability space with P being a topological τ -additive probability measure which has support S and let k be a continuous kernel function k defined on \mathcal{X} such that the corresponding RKHS \mathcal{H} has dimension $d < \infty$. Furthermore, assume that $\|k\|_\infty < \infty$, and that the centered covariance operator $\tilde{\mathfrak{C}}_c$ has an eigen-decomposition with eigenvalue such that:*

- (i) $\bar{\lambda}_l > 0 = \bar{\lambda}_{l+1}$ for some $l < d$ or
- (ii) $\bar{\lambda}_d > 0$.

It follows that \mathcal{H}_S has dimension $l + 1$ and $\tilde{\mathfrak{C}}_c^S$ has eigenvalues $\bar{\lambda}_i^S = \bar{\lambda}_i$ for all $i \leq l$ and $\bar{\lambda}_{l+1}^S = 0$ under (i), and \mathcal{H}_S has dimension d and the eigenvalues of $\tilde{\mathfrak{C}}_c^S$ are the same eigenvalues as of $\tilde{\mathfrak{C}}_c$ under (ii). Furthermore, under condition (ii) there exists a closed ball B centered at \mathbf{m} with radius $\bar{\lambda}_d/2\|k\|_\infty^{1/2}$ inside C . If $\lambda_1 = 0$ and there are no two points $x, y \in \mathcal{X}$ such that $k(x, x) = k(y, y) = k(x, y)$ then S consists of a single element and $P = k(x, \cdot)$ for some $x \in \mathcal{X}$. If $\lambda_1 > 0$ then \mathbf{m} lies in the relative interior of $F = \text{cch}\{k(x, \cdot) : x \in S\}$. In particular, under condition (i) there exists a closed ball B centered at \mathbf{m} such that $B \cap \text{aff } F \subset F$ and B has radius $\bar{\lambda}_l/2\|k\|_\infty^{1/2}$.

Proof of Proposition 2.15.

(a). — When $\bar{\lambda}_d > 0$ it follows that $E(h^2(X)) > E(h(X))^2$ for all $h \in \mathcal{H}, h \neq 0$. In particular, let $e_1, \dots, e_d \in \mathcal{H}$ be linearly independent and fix any $a_1, \dots, a_d \in \mathbb{R}$ such that a_1, \dots, a_d are not simultaneously equal to zero. Consider $e_1 \upharpoonright S, \dots, e_d \upharpoonright S$ then

$$E\left(\left(\sum_{i=1}^d a_i e_i \upharpoonright S\right)^2\right) = E\left(\left(\sum_{i=1}^d a_i e_i\right)^2\right) > 0.$$

Since this holds for all such a_1, \dots, a_d it follows that $e_1 \upharpoonright S, \dots, e_d \upharpoonright S$ are linearly independent and \mathcal{H}_S is d -dimensional. Similarly, when $\lambda_l > 0$ and $\lambda_{l+1} = 0$ it follows that there is an $h_1 \in \mathcal{H}_S$ which is almost surely equal to $\mathbf{1} \upharpoonright S$ and l linearly independent functions $e_1 \upharpoonright S, \dots, e_l \upharpoonright S \in \mathcal{H}$ which are also linearly independent of h_1 (for any non-trivial linear combination of $e_1 \upharpoonright S, \dots, e_l \upharpoonright S$ it follows that the second moment is strictly larger than the squared expected value and therefore these linear combinations are not equal to a constant function). In fact, h_1 is equal to $\mathbf{1} \upharpoonright S$ since by assumption the kernel function is continuous: assume that h_1 is not equal to $\mathbf{1}$, take a $x \in S$ such that $h_1(x) \neq 1$ and let $\epsilon = |h_1(x) - 1|$. Take a function $h \in \mathcal{H}$ such that $h \upharpoonright S = h_1$ and $\|h\| \leq \|h_1\|_{\mathcal{H}_S} + \epsilon/4$. The set $A = h^{-1}[\{y : |y - h_1(x)| < \epsilon/4\}]$ is open and has non-empty intersection with S . Also, h_1 is different from $\mathbf{1}$ on all of A . Due to [Fre01, 411N] it follows that $P(A) > 0$ and $\mathbf{1}$ is not almost surely equal to h_1 which is impossible. By the same argument it follows that the other eigenfunctions of $\tilde{\mathcal{C}}_c$ with zero eigenvalues are constant on S and therefore lie in the span of h_1 and the dimension of \mathcal{H}_S is $l + 1$.

(b). — In case (ii), let e_1, \dots, e_d be orthonormal in \mathcal{H} then, due to (a), the functions $e_1 \upharpoonright S, \dots, e_d \upharpoonright S$ are linearly independent. Also $\|e_i \upharpoonright S\|_{\mathcal{H}_S} = \|e_i\|$ for all $i \leq d$, because for any given $i \leq d$, $e_i \upharpoonright S$ does not lie in the linear subspace spanned by $\{e_j \upharpoonright S\}_{j \neq i}$. Similarly, for any $i, j \leq d$, $\|e_i \upharpoonright S + e_j \upharpoonright S\|_{\mathcal{H}_S} = \|e_i + e_j\|$ and $\|e_i \upharpoonright S + e_j \upharpoonright S\|_{\mathcal{H}_S}^2 = \|e_i \upharpoonright S\|_{\mathcal{H}_S}^2 + \|e_j \upharpoonright S\|_{\mathcal{H}_S}^2$. Hence, $\{e_i \upharpoonright S\}_{i \leq d}$ is an orthonormal basis of \mathcal{H}_S . In particular, when e_1, \dots, e_d are the eigenfunctions of $\tilde{\mathcal{C}}_c$ then $e_1 \upharpoonright S, \dots, e_d \upharpoonright S$ are the eigenfunctions of $\tilde{\mathcal{C}}_c^S$ since for $i \neq j$,

$$\langle \tilde{\mathcal{C}}_c^S e_i \upharpoonright S, e_j \upharpoonright S \rangle_{\mathcal{H}_S} = \langle \tilde{\mathcal{C}}_c e_i, e_j \rangle = 0$$

and for any $i \leq d$, $\langle \tilde{\mathcal{C}}_c^S e_i, e_i \rangle = \bar{\lambda}_i$.

By the same argument it follows that in case (i) that $e_1 \upharpoonright S, \dots, e_l \upharpoonright S, \mathbf{1} \upharpoonright S / \|\mathbf{1} \upharpoonright S\|_{\mathcal{H}_S}$ is an orthonormal basis of \mathcal{H}_S . Hence, if e_1, \dots, e_l are the first l eigenfunctions of $\tilde{\mathcal{C}}_c$ then $e_1 \upharpoonright S, \dots, e_l \upharpoonright S$ are the eigenfunctions of $\tilde{\mathcal{C}}_c^S$ and the eigenvalues match.

(c). — When $\bar{\lambda}_1 = 0$ it follows that $\mathcal{H}_S = \text{span}\{\mathbf{1} \upharpoonright S\}$ and $\langle \mathbf{m}_S, c\mathbf{1} \rangle_{\mathcal{H}_S} = c = \langle \mathbf{m}_{S,n}, c\mathbf{1} \rangle_{\mathcal{H}_S}$ for all $c \in \mathbb{R}$ where \mathbf{m}_S and $\mathbf{m}_{S,n}$ are just \mathbf{m} and \mathbf{m}_n when the kernel is restricted to S . Also, every function $h \in \mathcal{H}$ is constant on S and $\langle \mathbf{m}, h \rangle = \langle \mathbf{m}_n, h \rangle = \langle h, k(x, \cdot) \rangle$ for all $x \in S$. In particular, if S does not consist of a single element it follows that $k(x, y) = k(x, x) = k(y, y)$ for all $x, y \in S$. Reversing this statement leads to the claim made in the proposition.

(d). — Let $U = \text{aff } F$ be the affine subspace spanned by $k(x, \cdot), x \in S$. The element \mathbf{m} lies within U since for any $h \in \mathcal{H}$ for which $\langle k(x, \cdot), h \rangle = c$ for all $x \in S$

and for some $c \in \mathbb{R}$ it follows that $\langle \mathbf{m}, h \rangle = \int_S \langle k(x, \cdot), h \rangle dP = c$. In other words, if \mathbf{m} would not lie in U then there would be an $h, \|h\| > 0$, that stands perpendicular on U and such that $\mathbf{m} = h + \arg \min_{g \in U} \|g - \mathbf{m}\|$, and $\langle \mathbf{m}, h \rangle \neq \langle k(x, \cdot), h \rangle$ for all $x \in S$.

(e). — Under condition (ii) the constant functions are not in \mathcal{H} and the conditions of Proposition 2.13 are fulfilled and the existence of the ball with the specified radius follows directly.

Under condition (i) the constant functions are contained in \mathcal{H}_S and there exists a function $h_0 \in \mathcal{H}$ such that $h_0(x) = 1$ for all $x \in S$. Consider $\inf\{\|h\| : h(x) = 1 \text{ for all } x \in S\} \leq \|h_0\|$. Since $h \mapsto \|h\|$ is continuous and $\{h : \|h\| \leq \|h_0\|\}$ is compact it follows that the infimum is attained at some $h^* \in \mathcal{H}$. Also, $\|\mathbf{1} \upharpoonright S\|_{\mathcal{H}_S} = \inf\{\|h\| : h \upharpoonright S = \mathbf{1} \upharpoonright S\} = \|h^*\|$ and $\mathbf{1} \upharpoonright S = h^* \upharpoonright S$.

Let $U = \text{aff}\{k(x, \cdot) : x \in S\} \subset \mathcal{H}$ and $V = U - U$ be the subspace parallel to U . For any $h \in U$ there exists $m \in \mathbb{N}, \lambda_1, \dots, \lambda_m \in \mathbb{R}, \lambda_1 + \dots + \lambda_m = 1$ and $x_1, \dots, x_m \in S$ such that $h = \sum_{i=1}^m \lambda_i k(x_i, \cdot)$. In particular, if $g \in \mathcal{H}$ is constant on S and attains value c , then $\langle g, h \rangle = c \sum_{i=1}^m \lambda_i = c$. This implies that g is orthogonal to V since $\langle g, h_1 - h_2 \rangle = 0$ for all $h_1, h_2 \in U$. Due to assumption (i) there are $d - l$ eigenfunctions e_{l+1}, \dots, e_d of $\tilde{\mathcal{C}}_c$ which are constant on S , that is e_{l+1}, \dots, e_d are orthonormal and each of them is orthogonal to V . Also, any function h that is orthogonal to V has to be constant on S since $h(x) = \langle k(x, \cdot), h \rangle = \langle k(y, \cdot), h \rangle = h(y)$ for all $x, y \in S$. Since the eigenfunctions e_1, \dots, e_l have corresponding eigenvalues which are strictly greater than zero it follows that e_1, \dots, e_l cannot be constant and V has dimension l . Also, note that U cannot be equal to V , or better, U cannot be a subspace but only an affine subspace: assume otherwise then $0 \in U$ and if h is constant on S , attaining value $c \neq 0$, then $0 = \langle h, k(x, \cdot) \rangle = h(x) = c$ for all $x \in S$. In other words there cannot be functions that are constants on S in \mathcal{H} , but we know already that there are functions which are constant on S in \mathcal{H} .

(f). — We claim that $W = \text{span}(V \cup \{h^*\})$, when equipped with the inner product of \mathcal{H} , is isometric isomorphic to $\mathcal{H}_S, U \subset W$ and $(\text{span}\{h^*\}) \cap U \neq \emptyset$. We start with the latter claim. Since U is not a subspace it follows that the orthogonal projection of 0 onto U is not 0 itself. In detail, take any $x \in S$ and let P_V the projection onto the subspace V then the orthogonal projection onto U is the operator defined by $P_U h = k(x, \cdot) + P_V(h - k(x, \cdot))$. Now $P_U 0 = k(x, \cdot) - P_V k(x, \cdot) \neq 0$ and, $P_U 0$ is orthogonal to V . But that means the $P_U 0$ is constant on S and lies in the span of e_{l+1}, \dots, e_d . In particular, there is function that is constant on S which lies in U . Notice that for any function h which is constant on S there exists an element in U that lies in the span of h if $\|h\| = \sup_{g \in U} |\langle h, g \rangle| / \|g\|$. Recall that a function $g \in U$ can be written as $g = \sum_{i=1}^m \lambda_i k(x_i, \cdot)$ for some $m \in \mathbb{N}, x_1, \dots, x_m \in S$ and such that $\lambda_1 + \dots + \lambda_d = 1$. Hence, $\langle h, g \rangle = c$ if h attains value c and $h \in U$ if $\|h\| = |c| \sup_{g \in U} 1 / \|g\|$. Consider now a function h that is constant on S and attains value 1 then $\|h\| = \sup_{g \in U} 1 / \|g\|$. For $0 < \epsilon < \|h^*\|$ take $g \in U$ such that $\|h\| \leq \epsilon + 1 / \|g\|$, then

$$\frac{|\langle h^*, g \rangle|}{\|h^*\|} \leq \|g\| \leq \frac{|\langle h, g \rangle|}{\|h\| - \epsilon} = \frac{|\langle h^*, g \rangle|}{\|h\| - \epsilon} \leq \frac{|\langle h^*, g \rangle|}{\|h^*\| - \epsilon}.$$

In other words,

$$\frac{|\langle h^*, g \rangle|}{\|g\|} \leq \|h^*\| \leq \frac{|\langle h^*, g \rangle|}{\|g\|} + \epsilon$$

and $\sup_{g \in U} |\langle g, h^* \rangle| / \|h^*\| = \|h^*\|$ which implies that $(\text{span}\{h^*\}) \cap U \neq \emptyset$ and $U \subset W$

Let $\psi : W \rightarrow \mathcal{H}_S$ be the function that associates with $h \in W$ the function $h|_S \in \mathcal{H}_S$. The function ψ is linear since $(af + g)|_S = a(f|_S) + g|_S$ for all $f, g \in W$, $a \in \mathbb{R}$. We have seen already that $\psi(h^*) = \mathbf{1}|_S$ and $\|\psi(h^*)\|_{\mathcal{H}_S} = \|h^*\|$. Also any $h \in U$ lies in the span of $\{k(x, \cdot) : x \in S\}$ and, since \mathcal{H} is finite dimensional, h can be written as $\sum_{i=1}^d a_i k(x_i, \cdot)$ for some $a_1, \dots, a_d \in \mathbb{R}$ and $x_1, \dots, x_d \in S$. Hence, any $g, h \in U \subset W$ can be written as $h = \sum_{i=1}^d a_i k(x_i, \cdot)$, $g = \sum_{i=1}^d b_i k(y_i, \cdot)$, with some $a_i, b_i \in \mathbb{R}, x_i, y_i \in S$ for all $i \leq d$, and

$$\langle g, h \rangle = \sum_{i,j} a_i b_j k(x_i, y_j) = \sum_{i,j} a_i b_j k_S(x_i, y_j) = \langle \psi(g), \psi(h) \rangle_{\mathcal{H}_S}.$$

But this implies that for any $g, h \in V$, that is in the subspace parallel to U , there exists $\tilde{g}, \tilde{h} \in U$ and $x \in S$ such that $g = \tilde{g} - k(x, \cdot)$, $h = \tilde{h} - k(x, \cdot)$ and

$$\begin{aligned} \langle g, h \rangle &= \langle \tilde{g} - k(x, \cdot), \tilde{h} - k(x, \cdot) \rangle \\ &= \langle \psi(\tilde{g}), \psi(\tilde{h}) \rangle_{\mathcal{H}_S} - \langle \psi(\tilde{g}), k_S(x, \cdot) \rangle_{\mathcal{H}_S} - \langle \psi(\tilde{h}), k_S(x, \cdot) \rangle_{\mathcal{H}_S} + k_S(x, x) \\ &= \langle \psi(\tilde{g} - k(x, \cdot)), \psi(\tilde{h} - k(x, \cdot)) \rangle_{\mathcal{H}_S} \\ &= \langle \psi(g), \psi(h) \rangle_{\mathcal{H}_S}, \end{aligned}$$

since $\psi(k(x, \cdot)) = k_S(x, \cdot)$ for any $x \in S$. Finally, for any $g \in V$, write $g = \tilde{g} - k(x, \cdot)$ for some $x \in S, \tilde{g} \in U$ then

$$\langle g, h^* \rangle = \langle \tilde{g}, h^* \rangle - \langle k(x, \cdot), h^* \rangle = 0 = \langle \psi(\tilde{g}) - k_S(x, \cdot), \mathbf{1}|_S \rangle_{\mathcal{H}_S} = \langle \psi(g), \psi(h^*) \rangle_{\mathcal{H}_S},$$

since h^* is constant on S , $\psi(\tilde{g}) = \sum_{i=1}^m a_i \psi(k(x_i, \cdot))$ for some $m \in \mathbb{N}, x_1, \dots, x_m \in S$ and a_1, \dots, a_m such that $\sum_{i=1}^m a_i = 1$, and $\langle \psi(\tilde{g}), \mathbf{1}|_S \rangle_{\mathcal{H}_S} = 1$. Hence, ψ is an isometry and since \mathcal{H}_S and $\text{span}(V \cup \{h^*\})$ have the same dimension it follows that they are isometric isomorphic. The existence of the ball around \mathbf{m} of the specified radius follows now directly from Corollary 2.14. \square

Notice that \mathcal{H}_S does not have to be equal to \mathcal{H} when $\bar{\lambda}_d > 0$. For instance, when \mathcal{H} consists of the quadratic functions on $[-1, 1]$ and has therefore dimension 1. If the measure P is discrete with $P(\{-1\}) = P(\{1\}) = 1/2$ then $\bar{\lambda}_1 > 0$ but $S = \{-1, 1\}$ and $\mathcal{H}_S \neq \mathcal{H}$.

2.5. Locating \mathbf{m}_n within the empirical convex set

We are now combining the various results that we have derived. Section 2.4 allows us to refer the size of a ball within C around \mathbf{m} back to the question of the width of C . In Section 2.3 we derived various ways to lower bound the width of C . We also know that C_n converges to C . Section 2.2 contains results in this direction: Lemma 2.1 and 2.2 relate empirical quantities through VC and Rademacher bounds

to population limits while the Examples 1–4 provide lower bounds on the population limits. Of particular importance are the lower bounds on $P\psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c)$ in Section 3–4. These lower bounds are closely related to the bounds in Section 2.4 since in both settings we need to measure how much probability mass lies in various directions behind some threshold. To get high probability bounds for the existence of a ball of a certain size around \mathbf{m}_n within C_n we also need to control the convergence of \mathbf{m}_n to \mathbf{m} . But that is easy to do with another VC or Rademacher argument. The following two theorems combine these results under some natural conditions. The first result applies when $\mathcal{X} = [0, 1]^l$, \mathcal{H} is finite dimensional, functions in \mathcal{H} are Lipschitz continuous and when there is a lower bound on the density of the law of X_1, \dots, X_n on \mathcal{X} . We also assume that $\mathbf{1} \in \mathcal{H}$ but the result can easily be adapted to the case of $\mathbf{1} \notin \mathcal{H}$.

THEOREM 2.16. — *Let $\mathcal{X} = [0, 1]^l, l \geq 1$, and k a continuous kernel function on \mathcal{X} such that the corresponding RKHS \mathcal{H} is d -dimensional, $1 \leq d < \infty$, functions $h \in \mathcal{H}$ are Lipschitz continuous in the sense of (2.13) with Lipschitz constant $L > 0$, and $\mathbf{1} \in \mathcal{H}$. Furthermore, let X_1, \dots, X_n be i.i.d. random variables defined on some probability space and such that the law P of X_1 has a density p on \mathcal{X} and $\inf_{x \in \mathcal{X}} p(x) \geq c > 0$ for some constant c . Mercer’s theorem applies to k . Let $\tilde{\lambda}_l$ be the smallest eigenvalue of the Mercer decomposition. There exists a ball of radius*

$$\delta = \frac{c(\tilde{\lambda}_l^{1/2} \wedge L)^{l+1} \beta_l}{(l + 1)(2L)^l}$$

around \mathbf{m} in C in the affine subspace spanned by C . Furthermore, for any $q \in (0, 1)$ and whenever

$$n > \left(\frac{\sqrt{2 \log(4/q)} + 96 \|k\|_\infty^{1/2} / \delta}{c \beta_l (\delta / 8L)^l} \right)^2 \vee \left(\frac{\sqrt{2} \|k\|_\infty^{1/2} (1 + \sqrt{\log(2/q)})}{\delta / 4} \right)^2$$

then with probability $1 - q$ there exists a ball of radius $\delta/4$ around \mathbf{m}_n in C_n within the affine subspace spanned by C .

Proof.

(a). — Proposition 2.12 tells us that there is a ball of radius

$$\frac{c \xi^{l+1} \beta_l}{(l + 1)(2L)^l}$$

around \mathbf{m} in C whenever, there exists point $x \in \mathcal{X}$ such that $\langle h, k(x, \cdot) - \mathbf{m} \rangle \geq \xi > 0$ and $\xi/L \leq 1$. Furthermore, Proposition 2.10(9) tells us that the width of C in any direction h is at least $2\tilde{\lambda}_l^{1/2}$. Hence, there is a point $x \in \mathcal{X}$ for which $\langle h, k(x, \cdot) - \mathbf{m} \rangle \geq \tilde{\lambda}_l^{1/2}$ and we can chose $\xi = \tilde{\lambda}_l^{1/2} \wedge L$. This choice of ξ gives us the δ which is stated in the theorem text.

(b). — We continue with a high probability bound on $\|\mathbf{m}_n - \mathbf{m}\|$. There are various ways to attain such bounds. The following bound is taken from [BBDG19, Lemma 1] which has a better constant than the corresponding Rademacher complexity bound. A minor technical detail that is worth paying attention to, for

a countable dense subset $\widetilde{\mathcal{F}}$ of the unit ball of \mathcal{H}_- it holds that $\|\mathbf{m}_n - \mathbf{m}\| = \sup_{h \in \widetilde{\mathcal{F}}} |(1/n) \sum_{i=1}^n h(X_i) - Ph|$ since the kernel function is bounded. Hence, $\|\mathbf{m}_n - \mathbf{m}\|$ is measurable. Applying [BBDG19, Lemma 1] gives us

$$(2.18) \quad \Pr(\|\mathbf{m}_n - \mathbf{m}\| \geq \sqrt{2}\|k\|_\infty^{1/2}(1 + \sqrt{\log(1/q)})n^{-1/2}) \leq q.$$

(c). — Next, we control the difference between C_n and C . Let $c' = -\delta/2$ then due to Lemma 2.2 for any $\gamma > 0$ and with probability $1 - q$ simultaneously for all $u \in \widetilde{\mathcal{F}}$,

$$(2.19) \quad \begin{aligned} P_n\psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c') \\ \geq P\psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c') - \left(\sqrt{2\log(2/q)} + 24\|k\|_\infty^{1/2}/\gamma\right)n^{-1/2}. \end{aligned}$$

Chose $\gamma = \delta/4$ and let $x_0 \in \mathcal{X}$ be a point such that $\langle u, \phi(x_0) - \mathbf{m} \rangle \leq -\delta$. According to Eq. (2.3),

$$P\psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c') \geq \Pr(\langle u, \phi(X_1) - \mathbf{m} \rangle \leq -\gamma + c').$$

As in the proof of Proposition 2.12, let $A = \{y : \|y - x_0\| \leq \delta/4L, y \in \mathcal{X}\}$ and B the translation of A to the origin. In other words, $B = \{y : \|y\| \leq \delta/4L, y \in \mathcal{X}\}$. Recall that β_l denotes the Lebesgue measure of the l -dimensional unit ball. The Lebesgue measure of a ball of radius $\delta/4L$ is $(\delta/4L)^l\beta_l$ [Fre01, 252Qc]. The set B is the intersection of a ball centered at the origin with the hypercube $[0, 1]^l$ which corresponds to 1 out of 2^{-l} quadrants of the ball whenever the radius $\delta/4L$ of the ball does not exceed one. Since,

$$\frac{\delta}{4L} \leq \frac{cL^{l+1}\beta_l}{4L(l+1)(2L)^l} = \frac{c(2^{-l}\beta_l)}{4(l+1)}$$

and $2^{-l}\beta_l$ is the volume of a ball of radius $1/2$ it follows that $c2^{-l}\beta_l \leq c \text{vol}(\mathcal{X}) \leq 1$ and $\delta/4L \leq 1$. Hence, the volume of B is $\mu_l(B) = (\delta/8L)^l\beta_l$. In the proof of Proposition 2.12 we show that $\Pr(X_1 \in A) \geq \Pr(X_1 \in B)$ and, hence, $\Pr(X_1 \in A) \geq c\mu_l(B)$. The same line of reasoning as in Section 3 can now be applied. In detail, since $-\gamma + c' = -(3/4)\delta$, $\langle u, \phi(x_0) - \mathbf{m} \rangle \leq -\delta$ and for any $x \in A$,

$$|\langle u, \phi(x) - \mathbf{m} \rangle - \langle u, \phi(x_0) - \mathbf{m} \rangle| \leq \|\phi(x) - \phi(x_0)\| \leq L\|x - x_0\| \leq \delta/4,$$

it follows that for all $x \in A$, $\langle u, \phi(x) - \mathbf{m} \rangle \leq -\gamma + c'$. Therefore,

$$P\psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c') \geq \Pr(\langle u, \phi(X_1) - \mathbf{m} \rangle \leq -\gamma + c') \geq \Pr(X_1 \in B) \geq c\beta_l(\delta/8L)^l.$$

For a given q let

$$N_q = \left(\frac{\sqrt{2\log(2/q)} + 96\|k\|_\infty^{1/2}/\delta}{c\beta_l(\delta/8L)^l}\right)^2$$

then whenever $n > N_q$ with probability $1 - q$ there is a ball of radius $\delta/2$ around \mathbf{m} in C_n .

(d). — Finally, we transfer the lower bound that we have for a ball within C_n around \mathbf{m} to \mathbf{m}_n . For $q \in (0, 1)$ let

$$\widetilde{N}_q = \left(\frac{\sqrt{2}\|k\|_\infty^{1/2}(1 + \sqrt{\log(1/q)})}{\delta/4}\right)^2.$$

Then for any $n > \tilde{N}_q$ with probability at least $1 - q$, $\|\mathbf{m}_n - \mathbf{m}\| \leq \delta/4$.

Combining the bounds it follows that with probability $1 - q$ there is ball of size $\delta/4$ in C_n around \mathbf{m}_n whenever $n > N_{q/2} \vee \tilde{N}_{q/2}$. \square

The second result uses an assumption on the centered covariance operator instead of an assumption on the density. For this result we actually do not need to use the results on the minimal width of C . We formulate this result directly for the case where P is allowed to attain values in a strict subset of \mathcal{X} . A fortunate circumstance in that setting is that the empirical convex set converges to the intersection of C with the minimal face that contains \mathbf{m} and algorithms that work with the empirical convex set adapt automatically to the structure of the covariance operator. The result relies on the existence of the support of the measure P . A weak assumption to guarantee this existence is that P is a τ -additive topological measure [Fre01, 411N].

THEOREM 2.17. — *Let $(\mathcal{X}, \mathcal{A}, P)$ be some probability space with P being a topological measure that is τ -additive, and with measurable kernel function k defined on \mathcal{X} such that the corresponding RKHS \mathcal{H} is finite dimensional. Furthermore, let X_1, \dots, X_n be i.i.d. random variables attaining values in \mathcal{X} and with law P and assume that $\|k\|_\infty < \infty$, and that the centered covariance operator $\tilde{\mathcal{C}}_c$ has an eigen-decomposition with smallest non-zero eigenvalue being $\bar{\lambda}_d$. There exists a ball of radius $\delta = \bar{\lambda}_d/2\|k\|_\infty^{1/2}$ around \mathbf{m} in C within the affine subspace spanned by C . Furthermore, for any $q \in (0, 1)$ and whenever n is (strictly) greater than*

$$\left(\frac{8\|k\|_\infty \left(\sqrt{2 \log(4/q)} + 192\|k\|_\infty/\bar{\lambda}_d \right)}{\bar{\lambda}_d^2} \right)^2 \vee \left(\frac{\sqrt{32}\|k\|_\infty^{1/2} \left(1 + \sqrt{\log(2/q)} \right)}{\delta} \right)^2$$

then with probability $1 - q$ there exists a ball of radius $\delta/4$ around \mathbf{m}_n in C_n within the affine subspace spanned by C .

Proof.

(a). — The existence of the ball with radius δ follows directly from Proposition 2.15. Furthermore, the same high probability bound for $\|\mathbf{m}_n - \mathbf{m}\|$ as in the proof of Theorem 2.16 applies. The bound for C_n also runs along the same line as in Theorem 2.16. Consider in the following the RKHS \mathcal{H}_S . Let $\gamma = \delta/4$ and $c = -\delta/2$ then Eq. (2.5) tells us that $P\psi_\gamma(\langle u, \phi_S(\cdot) - \mathbf{m}_S \rangle_{\mathcal{H}_S} - c) \geq \bar{\lambda}^2/8\|k\|_\infty$ whenever $u \in \mathcal{H}_S$ has unit norm. Hence, with probability $1 - q$ and simultaneously for all $u \in \tilde{\mathcal{F}}$, where $\tilde{\mathcal{F}}$ is countable dense subset of the unit ball of \mathcal{H}_S ,

$$P_n\psi_\gamma(\langle u, \phi(\cdot) - \mathbf{m} \rangle - c') \geq \bar{\lambda}_d^2/8\|k\|_\infty - \left(\sqrt{2 \log(2/q)} + 192\|k\|_\infty/\bar{\lambda}_d \right) n^{-1/2}.$$

(b). — For a given $q \in (0, 1)$ let $\tilde{N}_q = (\sqrt{32}\|k\|_\infty^{1/2}(1 + \sqrt{\log(1/q)})/\delta)^2$ then for any $n > \tilde{N}_q$, $\|\mathbf{m}_n - \mathbf{m}\| \leq \delta/4$ with probability $1 - q$. Similarly, with probability $1 - q$ for $n > N_q$ there is a ball of radius $\delta/2$ around \mathbf{m} in C_n (as a subset of the affine subspace spanned by C), where

$$N_q = \frac{64\|k\|_\infty^2 \left(\sqrt{2 \log(2/q)} + 192\|k\|_\infty/\bar{\lambda}_d \right)^2}{\bar{\lambda}_d^4}. \quad \square$$

The convergence of the empirical mean embedding and empirical convex set are both unproblematic in the large sample case in both theorems. The bottleneck of the approach is rather the size of the convex set C itself.

2.6. Rescaling the kernel function does not affect compression

We finish this section by studying the effect of modifying the kernel function, or the involved convex sets, on the approximation of \mathbf{m} . Given that the smallest eigenvalues of the covariance operator and the width of C control the approximation of \mathbf{m} it is natural to try to increase these. One way to do so is to scale the kernel function by a constant factor $\alpha > 0$, i.e. replace the kernel function k on \mathcal{X} by αk . However, due to [PR16, Proposition 5.20] the inner products corresponding to the two spaces are scaled versions of each other (for all $x, y \in \mathcal{X}$, $\langle k(x, \cdot), k(y, \cdot) \rangle_k = (1/\alpha) \langle \alpha k(x, \cdot), \alpha k(y, \cdot) \rangle_{\alpha k}$) and the algorithms that we discuss in the next section are unaffected by this change. One might also wonder if the error bounds are affected and if we can, at least, optimize these by choosing an appropriate scaling. It turns out that the error bounds are also invariant to the scaling of the kernel function. Let us start by analyzing the width of the convex set C . The change in width is easy to quantify: given a kernel function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ the width increases by $\alpha^{1/2}$ if we replace the kernel function by αk , where $\alpha > 0$. In detail, a function h , $\|h\|_k = 1$, lies in \mathcal{H}_k if, and only if, the function $\alpha^{1/2}h$ lies in $\mathcal{H}_{\alpha k}$ and has norm $\|\alpha^{1/2}h\|_{\alpha k} = 1$ [PR16, Proposition 5.20] and the width is

$$\text{width}_{\alpha^{1/2}h, \mathcal{H}_{\alpha k}}(C) = \alpha^{1/2} \left(\sup_{x \in \mathcal{X}} h(x) - \inf_{x \in \mathcal{X}} h(x) \right) = \alpha^{1/2} \text{width}_{h, \mathcal{H}_k}(C).$$

Similarly, when the smallest non-zero eigenvalue of the centered covariance operator $\tilde{\mathfrak{C}}_{c,k}$ is $\bar{\lambda}_l$, $l \geq 1$, it follows that the smallest non-zero eigenvalue of $\tilde{\mathfrak{C}}_{c,\alpha k}$ is $\alpha \bar{\lambda}_l$: let $h \in \mathcal{H}_k$, $\|h\|_k = 1$, be the eigenfunction corresponding to $\bar{\lambda}_l$ then $\alpha^{1/2}h$ has unit norm in $\mathcal{H}_{\alpha k}$ and

$$\langle \tilde{\mathfrak{C}}_{c,\alpha k} \alpha^{1/2}h, \alpha^{1/2}h \rangle_{\alpha k} = \alpha (E(h^2(X)) - (E(h(X)))^2) = \alpha \bar{\lambda}_l.$$

Recall that the relation between the width and the radius δ of the largest ball around \mathbf{m} is approximately $\delta \approx (\text{width}_h(C))^{d+1}$ when \mathcal{X} is a subset of \mathbb{R}^d and under suitable conditions on the density (end of Section 2.4.1). In other words, a scaling of the kernel function by α increases δ approximately by a factor of $\alpha^{(d+1)/2}$. This increase has to be compared to the increase in the Lipschitz constant which results from this scaling. In Eq. (2.14) the Lipschitz constant L enters into the lower bound on δ through L^{-d} . For a function $\alpha^{1/2}h \in \mathcal{H}_{\alpha k}$ we have that

$$\frac{|\alpha^{1/2}h(x) - \alpha^{1/2}h(y)|}{\|x - y\|} \leq \alpha^{1/2}L\|h\| = \alpha^{1/2}L\|\alpha^{1/2}h\|_{\alpha k},$$

for any $x, y \in \mathcal{X}$, $x \neq y$, if the Lipschitz assumption (2.13) holds with constant L for \mathcal{H} . Combining these we see that δ changes approximately by a factor of $\alpha^{(d+1)/2}L^{-d} = \alpha^{1/2}$. Hence, we can increase δ by increasing α . In the error bounds in the next section we will see that the key quantity for controlling the approximation

error is the ratio $\delta/\|k\|_\infty^{1/2}$, which has to be large to guarantee a good compression. Since $\|\alpha k\|_\infty = \alpha\|k\|_\infty$ this term is independent of α and rescaling does not change the rate of compression that is promised by the bounds. Also note that we are changing the norm by which we measure the error. For some approximation $\widehat{\mathbf{m}}$ of \mathbf{m} in \mathcal{H}_k we can observe that $\|\widehat{\mathbf{m}} - \mathbf{m}\|_k = \alpha^{1/2}\|\widehat{\mathbf{m}} - \mathbf{m}\|_{\alpha k}$. This $\alpha^{1/2}$ factor is canceled by the leading constant in the error bound of the algorithms, which is of the order $\|\alpha k\|_\infty^{1/2} = \alpha^{1/2}\|k\|_\infty^{1/2}$.

If we consider instead the bound on δ that is based on the covariance operator (Theorem 2.17) then δ is of order $\bar{\lambda}_l/\|k\|_\infty^{1/2}$ and a scaling of α leads again to $\alpha\bar{\lambda}_l/\alpha^{1/2}\|k\|_\infty^{1/2}$ and the scaling does not affect the error bound.

3. Related approximation problems

When confronted with a concrete statistical problem, it is typically insufficient to only approximate \mathbf{m} and it is useful to have techniques available to reduce machine learning problems to the problem of approximating \mathbf{m} by, for example, modifying the kernel function. We will focus in this section on the problem of kernel regression to demonstrate how the approach can be modified. In particular, we will discuss how changing the kernel function can allow us to approximate covariance operators and compress data (x, y) that we usually face in supervised learning problems. Moreover, we show how one can deal with unbounded response variables and we study the problem of compressing data simultaneously for multiple objectives. These are auxiliary results and a reader who is eager to get to the compression algorithms can safely jump ahead to Section 4. While kernel regression is used in this section as an example on which we can demonstrate different approaches, there is also the possibility that coresets methods can, in fact, be useful in practice for kernel regression (see Section 1.8 for a comparison to Nyström methods).

The least-squares error, when a regressor f from an RKHS \mathcal{H} with kernel k is used, is

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n (f(X_i) - Y_i)^2 \\ &= \frac{1}{n} \sum_{i=1}^n \langle f \otimes f, k(X_i, \cdot) \otimes k(X_i, \cdot) \rangle_{\otimes} - \frac{2}{n} \sum_{i=1}^n \langle f, Y_i k(X_i, \cdot) \rangle + \frac{1}{n} \sum_{i=1}^n Y_i^2. \end{aligned}$$

Looking at the right hand side we can note that we have to deal with *multiple approximation problems*. There are two high level approaches to addressing multiple approximation problems. We can either solve each approximation problem individually or we can solve them simultaneously. In terms of finding core-sets this means that we will get three different core-sets when solving the approximation problems individually and a single core-set when we solve the approximation problems simultaneously. Before getting back to this discussion, let us have a look at the individual terms in the above least-squares problem.

The *third term* on the right hand side is rather unproblematic since it does not depend on f and can be summarized by a single real number. In particular, if we

approximate each term individually then we can compress this term down to a single real number. The *first term* on the right hand side corresponds to an empirical covariance and can be treated in a similar way to the empirical measure, i.e.

$$\frac{1}{n} \sum_{i=1}^n k(X_i, \cdot) \otimes k(X_i, \cdot)$$

attains values in $\mathcal{H} \otimes \mathcal{H}$. In fact, since we are only interested in the terms $f^2(X_i)$, there is an RKHS that is better suited for our purposes than $\mathcal{H} \otimes \mathcal{H}$. Due to [PR16, Theorem 5.16] there exists a function $g \in \mathcal{H} \odot \mathcal{H}$ such that $f^2(X_i) = g(X_i)$, where $\mathcal{H} \odot \mathcal{H}$ is the RKHS that corresponds to the kernel function $\kappa(x, y) = k^2(x, y)$. The empirical covariance, when restricted to $\{(h, h) : h \in \mathcal{H}\}$, can be identified with

$$\mathfrak{C}_n = \frac{1}{n} \sum_{i=1}^n \kappa(X_i, \cdot) \in \mathcal{H} \odot \mathcal{H},$$

i.e. for any $h \in \mathcal{H}$,

$$\langle \mathfrak{C}_n, h^2 \rangle_{\mathcal{H} \odot \mathcal{H}} = \frac{1}{n} \sum_{i=1}^n h^2(X_i) = \frac{1}{n} \sum_{i=1}^n \langle k(X_i, \cdot) \otimes k(X_i, \cdot), h \otimes h \rangle_{\otimes}.$$

The random element \mathfrak{C}_n attains values within the empirical convex set

$$C_{\odot, n} = \text{cch}\{\kappa(X_i, \cdot) : i \leq n\}.$$

The corresponding population covariance element is given by

$$(3.1) \quad \mathfrak{C} = \int_{\mathcal{X}} \kappa(x, \cdot) dP(x)$$

which is contained in the convex set

$$C_{\odot} = \text{cch}\{\kappa(x, \cdot) : x \in \mathcal{X}\}.$$

In Section 2.4.2 we used the covariance operator $\tilde{\mathfrak{C}} : \mathcal{H} \rightarrow \mathcal{H}$. Notice that $\tilde{\mathfrak{C}}$ and \mathfrak{C} are closely related since for any $h \in \mathcal{H}$, $\langle \tilde{\mathfrak{C}}h, h \rangle = E(h^2(X)) = \langle \mathfrak{C}, h^2 \rangle_{\mathcal{H} \odot \mathcal{H}}$.

The *second term* in the above sum is more difficult to deal with than the other two due to the elements Y_i . We are looking at two approaches in Section 3.3: in the first approach, we consider $\mathfrak{m}_{y, n} = (1/n) \sum_{i=1}^n Y_i k(X_i, \cdot)$ as a subset of $\text{cch}\{Y_i k(X_i, \cdot) : i \leq n\}$. That approach works well when we consider the approximation problem in isolation, but it does lead to complications when trying to approximate $\mathfrak{m}_{y, n}$ simultaneously to \mathfrak{C}_n . In the second approach, we incorporate the Y_i 's into the kernel by using $\langle Y_i, \cdot \rangle_{\mathbb{R}} \otimes k(X_i, \cdot)$ as a kernel function and by mapping $f \in \mathcal{H}$ to $\langle 1, \cdot \rangle_{\mathbb{R}} \otimes f(\cdot) \in \mathbb{R}' \otimes \mathcal{H}$, i.e. for $i \leq n$,

$$\langle \langle 1, \cdot \rangle_{\mathbb{R}} \otimes f(\cdot), \langle Y_i, \cdot \rangle_{\mathbb{R}} \otimes k(X_i, \cdot) \rangle_{\mathbb{R}' \otimes \mathcal{H}} = Y_i f(X_i) = \langle f, Y_i k(X_i, \cdot) \rangle.$$

In this approach, we are aiming to approximate

$$\mathfrak{m}_{y, n}^{\otimes} = \frac{1}{n} \sum_{i=1}^n \langle Y_i, \cdot \rangle_{\mathbb{R}} \otimes k(X_i, \cdot) = \langle 1, \cdot \rangle_{\mathbb{R}} \otimes \mathfrak{m}_{y, n}.$$

In the least-squares problem we might like to use the same points X_i with the same weights w_i to *approximate \mathfrak{C}_n and $\mathfrak{m}_{y, n}$ simultaneously*. As we mentioned above, this approach is facilitated by incorporating the Y_i 's and by moving to $\mathfrak{m}_{y, n}^{\otimes}$. Similarly,

it is useful to extend the functions in $\mathcal{H} \odot \mathcal{H}$ to $\mathbb{R} \times \mathcal{X}$ by setting $\widehat{h}(y, x) = h(x)$ for $h \in \mathcal{H} \odot \mathcal{H}$. We denote the resulting space by $\widehat{\mathcal{H} \odot \mathcal{H}}$ which is again a Hilbert space when using the inner product $\langle \widehat{h}, \widehat{g} \rangle_{\widehat{\mathcal{H} \odot \mathcal{H}}} = \langle h, g \rangle_{\mathcal{H} \odot \mathcal{H}}$, for any $g, h \in \mathcal{H} \odot \mathcal{H}$. In fact, it is an RKHS with kernel function $\kappa_y((y_1, x_1), (y_2, x_2)) = \kappa(x_1, x_2)$, since

$$\langle \widehat{h}, \kappa_y((y, x), \cdot) \rangle_{\widehat{\mathcal{H} \odot \mathcal{H}}} = \langle h, \kappa(x, \cdot) \rangle_{\mathcal{H} \odot \mathcal{H}} = h(x) = \widehat{h}(y, x).$$

The empirical covariance operator now becomes $\mathfrak{C}_{y,n} = (1/n) \sum_{i=1}^n \kappa_y((Y_i, X_i), \cdot)$.

One way to achieve a simultaneous approximation of $\mathfrak{C}_{y,n}$ and $\mathfrak{m}_{y,n}^{\otimes}$ is to use a direct sum $\mathcal{G} = (\widehat{\mathcal{H} \odot \mathcal{H}}) \oplus (\mathbb{R}' \otimes \mathcal{H})$ and consider the convex set

$$C_{\oplus,n} = \text{cch}\{(\kappa((Y_i, X_i), \cdot), \langle Y_i, \cdot \rangle_{\mathbb{R}} \otimes k(X_i, \cdot)) : i \leq n\} \subset \mathcal{G}.$$

The element that we like to approximate is in this context $(\mathfrak{C}_{y,n}, \mathfrak{m}_{y,n}^{\otimes})$ which lies in $C_{\oplus,n}$. Let $(\bar{\mathfrak{C}}_{y,n}, \bar{\mathfrak{m}}_{y,n}^{\otimes})$ be some element in \mathcal{G} , then

$$\|(\bar{\mathfrak{C}}_{y,n}, \bar{\mathfrak{m}}_{y,n}^{\otimes}) - (\mathfrak{C}_{y,n}, \mathfrak{m}_{y,n}^{\otimes})\|_{\mathcal{G}}^2 = \|\bar{\mathfrak{C}}_{y,n} - \mathfrak{C}_{y,n}\|_{\widehat{\mathcal{H} \odot \mathcal{H}}}^2 + \|\bar{\mathfrak{m}}_{y,n}^{\otimes} - \mathfrak{m}_{y,n}^{\otimes}\|_{\mathbb{R}' \otimes \mathcal{H}}^2$$

and a good approximation in \mathcal{G} guarantees good approximations of $\mathfrak{C}_{y,n}$ and $\mathfrak{m}_{y,n}^{\otimes}$ simultaneously.

3.1. Assumptions

There are some *minimal assumptions* that we need to impose on \mathfrak{C} , \mathfrak{m}_y and variations thereof to be well defined. Generally, we assume that we have independent pairs of random variables $(X, Y), (X_1, Y_1), \dots$ defined on some probability space $(\Omega, \mathcal{A}, \mu)$. For \mathfrak{C} to be well-defined it suffices to assume that $\kappa(Y, \cdot) \in \mathcal{L}^1(\mu; \mathcal{H} \odot \mathcal{H})$ and, similarly, for \mathfrak{m}_y it suffices to assume that $Y \in \mathcal{L}^2(\mu)$ and $k(X, \cdot) \in \mathcal{L}^2(\mu; \mathcal{H})$ since then $\int \|Yk(X, \cdot)\| d\mu \leq \|Y\|_2 \|k(X, \cdot)\|_2 < \infty$ and $\mathfrak{m}_y = \int Yk(X, \cdot) d\mu \in \mathcal{H}$.

Some *further assumptions* are useful to facilitate the following analyses. In particular, in the least-squares setting it is natural to assume that $Y = f_0(X) + \epsilon$, where f_0 is a suitable function, ϵ is a zero mean real-valued random variable representing measurement noise, and X and ϵ are independent. When making this assumption we are assuming that ϵ is a random variable that is defined on the probability space $(\Omega, \mathcal{A}, \mu)$. To guarantee that $Y \in \mathcal{L}^2(\mu)$ it is enough to assume that $f_0(X) \in \mathcal{L}^2(\mu)$ and $\epsilon \in \mathcal{L}^2(\mu)$.

3.2. Covariance operators

Since Eq. (3.1) tells us that the covariance operator \mathfrak{C} can be treated like a mean element after changing the kernel, it follows immediately that the techniques we developed for approximating \mathfrak{m}_n and \mathfrak{m} can be applied to the covariance operator \mathfrak{C} and its empirical version \mathfrak{C}_n . One might also wonder if the approximation problems for \mathfrak{m}_n and \mathfrak{C}_n are related and if any information that we might deduce about C_n and \mathfrak{m}_n can give us insights into the approximation problem for \mathfrak{C}_n . For instance, can we say anything about the width of the convex set in \mathcal{H}_κ based on the width of the

convex set in \mathcal{H} ? Or, does an assumption on the variance of functions in \mathcal{H} translate to statements about the variance of certain functions in \mathcal{H}_κ ? Unsurprisingly, this seems to be impossible to do in general. However, for certain functions in \mathcal{H}_κ we can infer statements about the width and the variance. Similarly, under very stringent assumptions on certain eigenvalues we can say something about the width of the convex set with respect to any function in \mathcal{H}_κ . Before looking into these questions we start by taking a closer look at the RKHS $\mathcal{H} \odot \mathcal{H}$.

Note that the space $\mathcal{H} \odot \mathcal{H}$ is not just the RKHS corresponding to the kernel $\kappa(x, y) = k^2(x, y)$ but it is also closely related to the tensor product $\mathcal{H} \otimes \mathcal{H}$. In particular, $h \in \mathcal{H} \odot \mathcal{H}$ if, and only if, there exists $u \in \mathcal{H} \otimes \mathcal{H}$ such that $h(x) = u(x, x)$ for all $x \in \mathcal{X}$, and then $\|h\|_{\mathcal{H} \odot \mathcal{H}} = \inf\{\|u\|_{\mathcal{H} \otimes \mathcal{H}} : h(x) = u(x, x) \text{ for all } x \in \mathcal{X}\}$ [PR16, Theorem 5.16]. Observe that for any $h \in \mathcal{H} \odot \mathcal{H}$ the set $A_h := \{u \in \mathcal{H} \otimes \mathcal{H} : h(x) = u(x, x) \text{ for all } x \in \mathcal{X}\}$ is a convex subset of $\mathcal{H} \otimes \mathcal{H}$. Define the linear operator $T : \mathcal{H} \otimes \mathcal{H} \rightarrow \mathcal{H} \odot \mathcal{H}$ by $Tu = u \circ \psi$ where $\psi : \mathcal{X} \rightarrow \mathcal{X} \times \mathcal{X}, \psi(x) = (x, x)$, and observe that T is bounded since $\|Tu\|_{\mathcal{H} \odot \mathcal{H}} \leq \|u\|_{\mathcal{H} \otimes \mathcal{H}}$. Hence, $A_h = T^{-1}[\{h\}]$ is closed. Also, observe that when $h = (f \otimes g) \circ \psi$ for some $f, g \in \mathcal{H}$ then there exists functions $f_1, f_2, f_3 \in \mathcal{H}$ such that $h = (1/2)(f_1 \otimes f_1 - f_2 \otimes f_2 - f_3 \otimes f_3) \circ \psi$ and $\|f \otimes g\|_{\otimes} \geq \|f_1 \otimes f_1 + f_2 \otimes f_2 + f_3 \otimes f_3\|_{\otimes}$: choose $f_1 = (f - g), f_2 = f, f_3 = g$ then

$$f \otimes g + g \otimes f = (f + g) \otimes (f + g) - (f \otimes f) - (g \otimes g) = f_1 \otimes f_1 - f_2 \otimes f_2 - f_3 \otimes f_3$$

and, since $(f \otimes g) \circ \psi = (g \otimes f) \circ \psi$, the first statement follows. In terms of the norm observe that $\|f \otimes g + g \otimes f\|_{\otimes} = \|f_1 \otimes f_1 - f_2 \otimes f_2 - f_3 \otimes f_3\|_{\otimes} \leq 2\|f\| \|g\| = 2\|f \otimes g\|_{\otimes}$. Also, note that

$$(1/2)\|f \otimes g + g \otimes f\|_{\otimes}^2 = \|f\|^2 \|g\|^2 + |\langle f, g \rangle|^2 = \|f \otimes g\|_{\otimes}^2 + |\langle f, g \rangle|^2 (\leq 2\|f \otimes g\|_{\otimes}^2),$$

where the expression in the bracket follows from the Cauchy–Schwarz inequality. Hence, when f and g are orthogonal, we have that

$$\|f_1 \otimes f_1 - f_2 \otimes f_2 - f_3 \otimes f_3\|_{\otimes} = \sqrt{2}\|f \otimes g\|_{\otimes}.$$

The point is that for any tensor $f \otimes g$ we can express $(f \otimes g) \circ \psi$ as a linear combination of “symmetric” tensor elements applied to ψ without increasing the tensor norm.

Let us next consider the closed subspace $U = \overline{\text{span}}\{k(x, \cdot) \otimes k(x, \cdot) : x \in \mathcal{X}\} \subset \mathcal{H} \otimes \mathcal{H}$ and the orthogonal projection P_U onto it. For any $u \in \mathcal{H} \otimes \mathcal{H}$ we have that $P_U u$ lies in the subspace spanned by $k(x, \cdot) \otimes k(x, \cdot)$ and $\|P_U u\|_{\otimes} \leq \|u\|_{\otimes}$. In fact, for any $x \in \mathcal{X}$,

$$\begin{aligned} u \circ \psi(x) &= \langle u, k(x, \cdot) \otimes k(x, \cdot) \rangle_{\otimes} = \langle u, P_U(k(x, \cdot) \otimes k(x, \cdot)) \rangle_{\otimes} \\ &= \langle P_U u, k(x, \cdot) \otimes k(x, \cdot) \rangle_{\otimes} = (P_U u) \circ \psi(x). \end{aligned}$$

In other words, for $h \in \mathcal{H} \odot \mathcal{H}$, if we can show that the infimum will be attained over A_h , then there exists an element u in U such that $h = u \circ \psi, \|u\|_{\otimes} = \|h\|_{\odot}$, and for any $v \notin U$ that fulfills $v \circ \psi = h$ it follows that $\|v\|_{\otimes} > \|u\|_{\otimes}$.

When \mathcal{H} is finite dimensional and the kernel function is bounded it follows that $A_h \cap U$ is compact: the tensor space $\mathcal{H} \otimes \mathcal{H}$ is finite dimensional since \mathcal{H} is finite dimensional. Also, h is bounded since the kernel function is bounded. For $\epsilon > 0$

consider the centered closed ball B of radius $\|h\|_{\odot} + \epsilon$ within $\mathcal{H} \otimes \mathcal{H}$. The intersection $B \cap A_h \cap U$ is non-empty and compact and the infimum is attained within this compact set.

3.2.1. Lower bound on the width of C_{\odot}

How can this *tensor product characterization* of $\mathcal{H} \odot \mathcal{H}$ be used to characterize the width of C_{\odot} ? Let us first consider tensors of the form $f \otimes f$, $f \in \mathcal{H}$, $\|f \otimes f\| = 1 = \|f\|$. The width of C_{\odot} in direction h , where $h = (f \otimes f) \circ \psi$, is

$$\begin{aligned} \text{width}_h C_{\odot} &= \sup_{x \in \mathcal{X}} h(x) - \inf_{x \in \mathcal{X}} h(x) \\ &= \sup_{x \in \mathcal{X}} \langle f \otimes f, k(x, \cdot) \otimes k(x, \cdot) \rangle_{\otimes} - \inf_{x \in \mathcal{X}} \langle f \otimes f, k(x, \cdot) \otimes k(x, \cdot) \rangle_{\otimes} \\ &= \sup_{x \in \mathcal{X}} f^2(x) - \inf_{x \in \mathcal{X}} f^2(x). \end{aligned}$$

If \mathcal{X} is *path connected*, k is *continuous* and, hence, f is continuous, then we can relate $\text{width}_h C_{\odot}$ to $\text{width}_f(C)$: whenever there exist $x, x' \in \mathcal{X}$ such that $f(x) > 0 > f(x')$, then there also exists an $\tilde{x} \in \mathcal{X}$ with $f(\tilde{x}) = 0$ due to the mean value theorem. In this case,

$$\text{width}_h C_{\odot} \geq \left(\sup_{x \in \mathcal{X}} f(x) \vee - \inf_{x \in \mathcal{X}} f(x) \right)^2 \geq ((1/2) \text{width}_f(C))^2.$$

If there is no x such that $f(x) = 0$, that is f attains only positive or only negative values, then we can argue in the following way. W.l.o.g. assume that f attains only positive values. Since for $a, b \geq 0$, $a \geq b$, $a^2 - b^2 \geq (a - b)^2$, it also follows in this case that

$$\text{width}_h C_{\odot} \geq ((1/2) \text{width}_f(C))^2.$$

This lower bound can fail to hold when the assumptions about k and \mathcal{X} are not fulfilled. Consider $\mathcal{X} = \{-1, 1\}$ with $k(x, y) = \delta_{x,y}$ and the function $f = k(1, \cdot) - k(-1, \cdot)$ which lies in the RKHS. For this function f ,

$$\sup_{x \in \mathcal{X}} f^2(x) - \inf_{x \in \mathcal{X}} f^2(x) = 0 < 1 = \left((1/2) \left(\sup_{x \in \mathcal{X}} f(x) - \inf_{x \in \mathcal{X}} f(x) \right) \right)^2.$$

The factor 1/2 in the lower bound is redundant when f attains only positive or negative values since $a^2 - b^2 \geq (a - b)^2$ whenever a and b have the same sign. More importantly, the bound becomes loose when $\inf_{x \in \mathcal{X}} f^2(x)$ is large since $a^2 - b^2 - (a - b)^2 = 2b(a - b)$ whenever $a \geq b > 0$.

Moving on to other directions $h \in \mathcal{H} \otimes \mathcal{H}$, $\|h\|_{\otimes} = 1$, we can first observe that the arguments are not as straightforward as elements in U are of the form $\sum_{i=1}^d \alpha_i k(x_i, \cdot) \otimes k(x_i, \cdot)$ for some $d \in \mathbb{N}$, $\alpha_i \in \mathbb{R}$ and $x_i \in \mathcal{X}$ for $i \leq d$, and, when considering the supremum over $x \in \mathcal{X}$, the different terms $\alpha_i k^2(x_i, x)$ can potentially cancel each other; in particular, there is no reason why the α_i should all be positive.

Alternatively, we can consider the set $S = \{k(x, \cdot) \otimes k(x, \cdot) - f k(z, \cdot) \otimes k(z, \cdot) : x \in \mathcal{X}\}$ and the convex hull of S . This convex hull is closely related to C_{\odot} . Note that $\text{ch } S$ is a subset of $\mathcal{H} \otimes \mathcal{H}$ and contains the origin. Fix now any $u \in \text{span } S$, $\|u\|_{\otimes} = 1$, then $\text{span}\{u\}$ intersects $\text{ch } S$. In fact, $\text{span}\{u\} \cap \text{ch } S$ consists of more than

a single point since otherwise $\langle k(x, \cdot) \otimes k(x, \cdot), u \rangle_{\otimes} = \langle k(y, \cdot) \otimes k(y, \cdot), u \rangle_{\otimes}$ for all $x, y \in \mathcal{X}$ and $\langle u, s \rangle_{\otimes} = 0$ for any $s \in S$ which would imply that $u = 0$. Consider the two points at which $\text{span}\{u\}$ intersects with the boundary of $\text{ch } S$ and let d be the dimension of $\text{span } S$. Each of these points can be expressed as a convex combination of $l \leq d + 1$ points in S due to Carathéodory's theorem. Hence, $u = \tilde{u}/\|\tilde{u}\|$ where $\tilde{u} = \sum_{i=1}^l \alpha_i k(x_i, \cdot) \otimes k(x_i, \cdot)$ for some strictly positive α_i 's that sum to one and suitable points x_1, \dots, x_l . We can note right away that

$$(3.2) \quad \inf_{x \in \mathcal{X}} k(x, x)/(d + 1) \leq \left(\sum_{i,j=1}^l \alpha_i \alpha_j k^2(x_i, x_j) \right)^{1/2} = \|\tilde{u}\|_{\otimes} \leq \sup_{x \in \mathcal{X}} k(x, x).$$

Also, when $k(x, y)$ is a non-negative function an application of Jensen's inequality yields further results. In detail,

$$\begin{aligned} \sup_{x,y \in \mathcal{X}} \tilde{u}(x, x) - \tilde{u}(y, y) &= \sup_{x,y \in \mathcal{X}} \sum_{i=1}^l \alpha_i (k^2(x_i, x) - k^2(x_i, y)) \\ &\geq \sup_{x \in \mathcal{X}} \left(\sum_{i=1}^l \alpha_i k(x_i, x) \right)^2 - \inf_{y \in \mathcal{X}} \sum_{i=1}^l \alpha_i k^2(x_i, y). \end{aligned}$$

In the following, let $c = \inf_g \text{width}_g(C)$, where the infimum is taken over $\{g : g \in \mathcal{H}, \|g\| = 1\}$. When c is large, k is non-negative and $\|k\|_{\infty} \leq 1$, then this simple argument might be of use: since $k(x_i, y) \leq 1$ it follows that $k(x_i, y) \geq k^2(x_i, y)$ and

$$\sup_{x \in \mathcal{X}} \left(\sum_{i=1}^l \alpha_i k(x_i, x) \right)^2 - \inf_{y \in \mathcal{X}} \sum_{i=1}^l \alpha_i k^2(x_i, y) \geq \sup_{x \in \mathcal{X}} \left(\sum_{i=1}^l \alpha_i k(x_i, x) \right)^2 - \inf_{y \in \mathcal{X}} \sum_{i=1}^l \alpha_i k(x_i, y).$$

Furthermore, $(\sum_{i=1}^l \alpha_i k(x_i, y))^2 \geq \sum_{i=1}^l \alpha_i k(x_i, y) - 1/4$ and

$$\sup_{x,y \in \mathcal{X}} \tilde{u}(x, x) - \tilde{u}(y, y) \geq c - 1/4.$$

This is only useful for large c . If, in fact, $c > 1/4$, we can proceed and

$$\sup_{x,y \in \mathcal{X}} u(x, x) - u(y, y) \geq \frac{c - 1/4}{\sup_{x \in \mathcal{X}} k(x, x)} \geq c - 1/4.$$

3.2.2. Lower bounds on the fourth moments

Instead of controlling the width of C_{\odot} we can also aim to control the covariance operator corresponding to the kernel κ . Effectively, this corresponds to bounds on the fourth moments. A bound on the non-centered fourth moment is, in fact, easy to derive: let $\tilde{\mathcal{C}}$ be the covariance operator (now interpreted as a linear operator) corresponding to the kernel k . We need to control $E((h(X) - E(h(X)))^2)$ for a function $h \in \mathcal{H}_{\kappa}$. Choosing again a suitable $u \in \mathcal{H} \otimes \mathcal{H}$ such that $h = u \circ \psi$, we find that

$$E(h^2(X)) = E((u \circ \psi(X))^2).$$

If U is d -dimensional then, as above, we can write $u = \tilde{u}/\|\tilde{u}\|$ where $\tilde{u} = \sum_{i=1}^l \alpha_i k(x_i, \cdot) \otimes k(x_i, \cdot)$ for some positive α_i , $\sum_{i=1}^{d+1} \alpha_i = 1$, and suitable points

$x_i \in \mathcal{X}$ with $l \leq d + 1$. Given this representation of u and assuming that the smallest eigenvalue of $\tilde{\mathfrak{C}}$ is $\bar{\lambda}$,

$$\begin{aligned} \|\tilde{u}\|^2 E(h^2(X)) &= \sum_{i,j=1}^l \alpha_i \alpha_j E((k(x_i, X)k(x_j, X))^2) \\ &\geq \sum_{i,j=1}^l \alpha_i \alpha_j \left| \langle \tilde{\mathfrak{C}}k(x_i, X), k(x_j, X) \rangle \right|^2 \\ &\geq \sum_{i=1}^l \alpha_i^2 \frac{\bar{\lambda}}{k(x_i, x_i)}. \end{aligned}$$

In particular, when $k(x, x) = 1$ for all $x \in \mathcal{X}$ then

$$E(h^2(X)) \geq \frac{\bar{\lambda}}{(d+1)\|\tilde{u}\|^2} \geq \frac{\bar{\lambda}}{d+1}$$

follows from Eq. (3.2).

However, to say something about the largest ball that lies around \mathfrak{C} within C_\odot we need a lower bound on the variance of $h \in \mathcal{H}_\kappa$. This is not straight-forward and will need, in all likelihood, some stringent assumptions: consider the variance of an arbitrary functions $h \in \mathcal{H}_\kappa$, when \mathcal{H}_κ has dimension $d < \infty$. By the above argument, there exists a $u = \tilde{u}/\|\tilde{u}\|$, where $\tilde{u} = \sum_{i=1}^{d+1} \alpha_i k(x_i, \cdot) \otimes k(x_i, \cdot)$ for some non-negative α_i that sum to one and points $x_i \in \mathcal{X}$. Hence,

$$\begin{aligned} E((h(X) - E(h(X)))^2) &= E((u \circ \psi(X) - E(u \circ \psi(X)))^2) \\ &= \|\tilde{u}\|^{-2} \sum_{i,j=1}^{d+1} \alpha_i \alpha_j E\left((k^2(x_i, X) - E(k^2(x_i, X)))(k^2(x_j, X) - E(k^2(x_j, X)))\right). \end{aligned}$$

There is no reason why the sum over the off-diagonal elements

$$\sum_{i \neq j} \alpha_i \alpha_j E\left((k^2(x_i, X) - E(k^2(x_i, X)))(k^2(x_j, X) - E(k^2(x_j, X)))\right)$$

should be positive or should be of considerably smaller magnitude than the sum over the diagonal elements.

3.3. Weighted mean embedding

There are different ways to address the term $Y_i f(X_i)$, $i \leq n$, that occurs in the least squares problem and there are a variety of natural assumptions under which one can study the corresponding compression problem. Let us first have a look at how $Y_i f(X_i)$ can be lifted into the RKHS so that we can apply the compression techniques. A first approach to do so was introduced at the beginning of Section 3, where we wrote $Y_i f(X_i) = \langle f, Y_i k(X_i, \cdot) \rangle$. Using this representation we can try to approximate $(1/n) \sum_{i=1}^n Y_i k(X_i, \cdot) \in \mathcal{H}$. A second approach is to map the Y_i 's to linear functionals, that is to elements in the dual space \mathbb{R}' , and to consider the tensor products

$$\langle Y_i, \cdot \rangle_{\mathbb{R}} \otimes k(X_i, \cdot).$$

We still would like to work with an inner product similarly to $\langle f, k(X_i, \cdot) \rangle$ and we can do so if we work with $\langle 1, \cdot \rangle_{\mathbb{R}} \otimes \langle f, \cdot \rangle$ instead of f . In particular,

$$(3.3) \quad Y_i f(X_i) = \langle \langle Y_i, \cdot \rangle_{\mathbb{R}} \otimes k(X_i, \cdot), \langle 1, \cdot \rangle_{\mathbb{R}} \otimes \langle f, \cdot \rangle \rangle_{\otimes}.$$

Both approaches are natural when the Y_i 's are bounded but various issues arise when they are not. In particular, for the latter approach the elements $\langle Y_i, \cdot \rangle_{\mathbb{R}} \otimes k(X_i, \cdot)$ are not contained with probability one in a ball in the corresponding tensor product space. We therefore discuss the case where the Y_i 's are bounded first before moving on to the unbounded case. Finally, it is often natural to impose an assumption on the relation between X_i and Y_i , like

$$(3.4) \quad Y_i = f_0(X_i) + \epsilon_i$$

with $f_0 \in \mathcal{L}^2(P)$, ϵ_i independent of X_i , $E(\epsilon_i) = 0$, and the X_i and ϵ_i are i.i.d.

This leaves us with a total of eight different settings. But not all of these settings are useful for deepening our understanding. In particular, little can be said without the assumption (3.4) and we assume in the following, up to short discussions, that Eq. (3.4) holds. Beyond that we focus on three settings: the first setting uses the assumption that Y_i is bounded and we use the form $Y_i k(X_i, \cdot)$. We then move on to translate the results to the approach where we model $Y_i f(X_i)$ through the tensor product as in Eq. (3.3). Finally, we lift the assumption that Y_i is bounded and study the problem in the context of the assumptions (3.3) and (3.4). We assume throughout that the X_i 's are i.i.d. and that $k(X_i, \cdot) \in \mathcal{L}^2(P; \mathcal{H})$.

3.3.1. First setting: bounded Y_i 's & $Y_i f(X_i) = \langle f, Y_i k(X_i, \cdot) \rangle$

There are few natural question when working with the empirical estimate $\mathbf{m}_{y,n} = (1/n) \sum_{i=1}^n Y_i k(X_i, \cdot)$: what is a natural convex set which contains $\mathbf{m}_{y,n}$ and over which we can optimize efficiently? Do we have suitable population limits of the empirical quantities? What can be said about the diameter of the empirical convex set, about how centered $\mathbf{m}_{y,n}$ lies within the set and are assumptions on the covariance operator of use? In terms of an empirical convex set which contains $\mathbf{m}_{y,n}$ it is natural to consider the set

$$C_{y,n} = \text{cch}\{Y_i k(X_i, \cdot) : i \leq n\}$$

and optimization over this set is possible since we have control over the extremes of it.

Under the assumption that the Y_i 's are of the form $f_0(X_i) + \epsilon_i$, there are natural expressions for the population limits \mathbf{m}_y and the convex set C_y . For concreteness, we assume in the following that \mathcal{X} is a *Borel space*, \mathcal{H} is *separable*, f_0 and the feature map $\phi : \mathcal{X} \rightarrow \mathcal{H}$ are measurable, f_0 , the kernel function k and the ϵ_i 's are *bounded*, ϵ_i is independent of X_i , and $E(\epsilon_i) = 0$. We can define the population limit of $\mathbf{m}_{y,n}$ through

$$\mathbf{m}_y = \int f_0(x) k(x, \cdot) dP(x),$$

where P is the law of X_1 . The element \mathbf{m}_y lies in \mathcal{H} : the function $f_0 \times \phi : \mathcal{X} \rightarrow \mathcal{H}$ is weakly measurable since when $h \in \mathcal{H}$, then $\langle f_0(x) \phi(x), h \rangle = f_0(x) h(x)$ is the product of two Borel measurable functions and is therefore Borel measurable. Because \mathcal{H} is

separable it follows that $f_0 \times \phi$ is Bochner measurable. Furthermore, $\|f_0(x)\phi(x)\| \leq |f_0(x)|k^{1/2}(x, x)$ is a bounded function of x and $\mathbf{m}_y = \int f_0(x)\phi(x) dP$ is well defined and lies in \mathcal{H} .

Controlling $\|\mathbf{m}_{y,n} - \mathbf{m}_y\|$. We can quantify the deviation of $\mathbf{m}_{y,n}$ from \mathbf{m}_y in the following way. Let \mathcal{F} be a countable dense subset of the unit ball of \mathcal{H} and using that dual elements can be moved through the Bochner integral, we get for $\alpha > 0$,

$$\begin{aligned} (3.5) \quad & \Pr(\|\mathbf{m}_y - \mathbf{m}_{y,n}\| \geq \alpha) \\ &= \Pr\left(\sup_{h \in \mathcal{F}} \left\langle h, \int f_0(X_1)k(X_1, \cdot) d\mu - \frac{1}{n} \sum_{i=1}^n Y_i k(X_i, \cdot) \right\rangle \geq \alpha\right) \\ &\leq \Pr\left(\sup_{h \in \mathcal{F}} \int f_0(X_1)h(X_1) d\mu - \frac{1}{n} \sum_{i=1}^n f_0(X_i)h(X_i) \geq \alpha/2\right) \\ &\quad + \Pr\left(\sup_{h \in \mathcal{F}} -\frac{1}{n} \sum_{i=1}^n \epsilon_i h(X_i) \geq \alpha/2\right). \end{aligned}$$

The latter term can be bounded by means of [Pin94, Theorem 3.3] which is a Bernstein type theorem for Banach spaced valued random variables. Note that the theorem statement in [Pin94] contains an error which is corrected in [Pin99]. The bound is the following,

$$\sup_{h \in \mathcal{F}} -\frac{1}{n} \sum_{i=1}^n \epsilon_i h(X_i) = \left\| \frac{1}{n} \sum_{i=1}^n \epsilon_i \phi(X_i) \right\|$$

and $v_n = \sum_{i=1}^n \epsilon_i \phi(X_i)$ attains values in \mathcal{H} . The sequence v_1, \dots, v_n is a martingale sequence in \mathcal{H} with regard to the filtration $F_t = \sigma(X_1, \dots, X_t, \epsilon_1, \dots, \epsilon_t)$, $t \leq n$, since $E(v_t | F_{t-1}) = E(\epsilon_t \phi(X_t)) + v_{t-1} = v_{t-1}$ (a.s.) for $2 \leq t \leq n$ due to the independence between ϵ_t and X_t . Furthermore, let $v_0 = 0$ and $F_0 = \{\emptyset, \Omega\}$ so that $E(v_1 | F_0) = E(v_1) = v_0$ a.s. We can continue the sequence by letting $v_t = v_n$ and $F_t = F_n$ for all $t > n$, which preserves the martingale property. To apply [Pin94, Pin99, Theorem 3.3] we need the following moment bounds; for all $m \geq 2$,

$$\sum_{t=1}^n E(\|\epsilon_t \phi(X_i)\|^m | F_{t-1}) \leq n(c\|k\|_\infty^{1/2})^m \quad (\text{a.s.}),$$

where c is an upper bound on $|\epsilon_t|$. This implies that we can set $\Gamma = c\|k\|_\infty^{1/2}$, $B = n^{1/2}c\|k\|_\infty^{1/2}$ in [Pin94, Pin99, Theorem 3.3] and

$$\begin{aligned} (3.6) \quad & \Pr\left(\left\| \frac{1}{n} \sum_{i=1}^n \epsilon_i \phi(X_i) \right\| \geq \frac{\alpha}{2}\right) \\ &\leq 2 \exp\left(-\frac{n\alpha^2}{4c\|k\|_\infty^{1/2} \left(c\|k\|_\infty^{1/2} + \alpha + \sqrt{c^2\|k\|_\infty + \alpha c\|k\|_\infty^{1/2}/n}\right)}\right). \end{aligned}$$

The first term can be controlled with a standard Rademacher argument after changing the kernel. Define the kernel $l = f_0 \otimes f_0$ and consider the product kernel $l \times k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ with RKHS $\mathcal{H}_{l \times k}$. For $h \in \mathcal{H}$ it follows that $f_0 \times h \in \mathcal{H}_{l \times k}$;

for an h of the form $\sum_{i=1}^m \alpha_i k(x_i, \cdot)$ one can write down the representation explicitly as $f_0 \times h = \sum_{i=1}^m \alpha_i f_0(x) k(x_i, x) = \sum_{i=1}^m (\alpha_i / f_0(x_i))(l \times k)(x_i, x)$, whenever $f_0(x_i) \neq 0$ for all $i \leq m$. We can thus write

$$(3.7) \quad \Pr\left(\sup_{h \in \tilde{\mathcal{F}}} \int f_0(X_1) h(X_1) d\mu - \frac{1}{n} \sum_{i=1}^n f_0(X_i) h(X_i) \geq \alpha/2\right) \\ = \Pr\left(\sup_{g \in \tilde{\mathcal{F}}_{l \times k}} \int g(X_1) d\mu - \frac{1}{n} \sum_{i=1}^n g(X_i) \geq \alpha/2\right),$$

where $\tilde{\mathcal{F}}_{l \times k}$ is dense subset of the unit ball of $\mathcal{H}_{l \times k}$. The Rademacher argument that we are using in (2.18) can now be applied.

Population limit of $C_{y,n}$. The next question to address is how to *define the population limit of $C_{y,n}$* . Under the boundedness assumption of Y_1 we can characterize the limit of $C_{y,n}$ in the following way. Let $\bar{b} = \inf\{b : b \in \mathbb{R}, \epsilon_1 \leq b \text{ a.s.}\}$, $\underline{b} = -\sup\{b : b \in \mathbb{R}, \epsilon_1 \geq b \text{ a.s.}\}$ and let

$$C_y = \text{cch}\left(\{(f_0(x) + \bar{b})k(x, \cdot) : x \in \mathcal{X}\} \cup \{(f_0(x) - \underline{b})k(x, \cdot) : x \in \mathcal{X}\}\right).$$

Note that for any $x \in \mathcal{X}$, $(f_0(x) + \bar{b})k(x, \cdot)$ and $(f_0(x) - \underline{b})k(x, \cdot)$ lie in \mathcal{H} and $C_y \subset \mathcal{H}$. Furthermore, $\epsilon_i \in [-\underline{b}, \bar{b}]$ (a.s.) and for $i \leq n$, $(f_0(X_i) + \epsilon_i)k(X_i, \cdot)$ is almost surely a convex combination of $(f_0(X_i) + \bar{b})k(X_i, \cdot)$ and $(f_0(X_i) - \underline{b})k(X_i, \cdot)$. Therefore, $C_{y,n}$ is almost surely contained within C_y .

While C_y is a natural limit of $C_{y,n}$, we face the problem that the convergence towards C_y can be arbitrarily slow since ϵ_1 can have a low probability of attaining values close to \bar{b} or \underline{b} . Assumptions on the distribution of ϵ_1 are one way to address this problem. Alternatively, we can work directly with $C_{y,n}$ and study how deep $\mathbf{m}_{y,n}$ lies in $C_{y,n}$ by controlling events of the form $\langle h, \tilde{Y}k(\tilde{X}, \cdot) - \mathbf{m}_y \rangle \leq c$, where (\tilde{X}, \tilde{Y}) has law P_n , and by comparing the random variables $\langle h, \tilde{Y}k(\tilde{X}, \cdot) - \mathbf{m}_y \rangle$ and $\langle h, \tilde{Y}k(\tilde{X}, \cdot) - \mathbf{m}_{y,n} \rangle$. We follow this latter approach and we use Rademacher complexities to control these events uniformly over the unit ball of \mathcal{H} . In the Rademacher approach, we control such events by lower bounding terms of the form

$$P_n \psi_\gamma(\langle h, \tilde{y}k(\tilde{x}, \cdot) - \mathbf{m}_{y,n} \rangle - c) \left(= \int \psi_\gamma(\langle h, \tilde{y}k(\tilde{x}, \cdot) - \mathbf{m}_{y,n} \rangle - c) dP_n(\tilde{x}, \tilde{y}) \right)$$

for suitable $c, \gamma \in \mathbb{R}$ and all unit norm elements $h \in \mathcal{H}$. The element $\mathbf{m}_{y,n}$ converges to \mathbf{m}_y and, because ψ_γ is $1/\gamma$ -Lipschitz continuous, it follows that

$$(3.8) \quad |P_n \psi_\gamma(\langle h, \tilde{y}k(\tilde{x}, \cdot) - \mathbf{m}_{y,n} \rangle - c) - P_n \psi_\gamma(\langle h, \tilde{y}k(\tilde{x}, \cdot) - \mathbf{m}_y \rangle - c)| \\ \leq P_n(1/\gamma) |\langle h, \mathbf{m}_{y,n} - \mathbf{m}_y \rangle| \leq (1/\gamma) \|\mathbf{m}_{y,n} - \mathbf{m}_y\|, \quad (\text{a.s.})$$

where the P_n term becomes redundant since no variables \tilde{y} and \tilde{x} are present in the last line.

Next, we consider the convergence of $P_n \psi_\gamma(\langle h, \tilde{y}k(\tilde{x}, \cdot) - \mathbf{m}_y \rangle - c)$ to its population limit $P(\psi_\gamma(\langle h, yk(x, \cdot) - \mathbf{m}_y \rangle - c))$ uniformly over the unit ball of \mathcal{H} . The convergence can be controlled by using Rademacher complexities. Because ϵ is used in this section

to denote the noise terms we will use ζ to denote Rademacher variables. Since ψ_γ is continuous and is applied to a subset of \mathbb{R} we can note that

$$\sum_{i=1}^n \zeta_i \psi_\gamma(Y_i h(X_i) - E(Yh(X)) - c)$$

is well defined. Also, $\gamma\psi_\gamma$ is a contraction vanishing at zero [GN16, Section 5.2.1] and for any finite subset F of the unit ball of \mathcal{H} it follow that

$$\begin{aligned} E_\zeta \left(\sup_{h \in F} \sum_{i=1}^n \zeta_i (\gamma/2) \psi_\gamma(Y_i h(X_i) - E(Yh(X)) - c) \right) \\ \leq E_\zeta \left(\sup_{h \in F} \sum_{i=1}^n \zeta_i Y_i h(X_i) \right) \leq E_\zeta \left(\sup_{h \in F} \left| \sum_{i=1}^n \zeta_i Y_i h(X_i) \right| \right) \quad (\text{a.s.}), \end{aligned}$$

from [GN16, Theorem 5.2.1, Eq. 5.50]. Note that, conditional on Y_i , the probability laws of $\zeta_i Y_i$ and $\zeta_i |Y_i|$ are the same. In particular,

$$E_\zeta \left(\sup_{h \in F} \left| \sum_{i=1}^n \zeta_i Y_i h(X_i) \right| \right) = E_\zeta \left(\sup_{h \in F} \left| \sum_{i=1}^n \zeta_i |Y_i| h(X_i) \right| \right) \quad (\text{a.s.}).$$

Recall that Rademacher complexities are stable under taking absolute convex hulls. Let

$$C_F = \{(|Y_1|h(X_1), \dots, |Y_n|h(X_n)) : h \in F\}$$

and

$$\widehat{C}_F = \{(\|Y\|_\infty h(X_1), \dots, \|Y\|_\infty h(X_n)) : h \in F\},$$

then $\text{abs conv } C_F \subset \text{abs conv } \widehat{C}_F$. Furthermore,

$$\begin{aligned} E_\zeta \left(\sup_{h \in F} \left| \sum_{i=1}^n \zeta_i |Y_i| h(X_i) \right| \right) &= E_\zeta \left(\sup_{t \in C_F} \left| \sum_{i=1}^n \zeta_i t_i \right| \right) = E_\zeta \left(\sup_{t \in \text{abs conv } C_F} \left| \sum_{i=1}^n \zeta_i t_i \right| \right) \\ &\leq E_\zeta \left(\sup_{t \in \text{abs conv } \widehat{C}_F} \left| \sum_{i=1}^n \zeta_i t_i \right| \right) = \|Y\|_\infty E_\zeta \left(\sup_{h \in F} \left| \sum_{i=1}^n \zeta_i h(X_i) \right| \right). \quad (\text{a.s.}) \end{aligned}$$

In summary, we have shown that

$$\begin{aligned} E_\zeta \left(\sup_{h \in F} \sum_{i=1}^n \zeta_i \psi_\gamma(Y_i h(X_i) - E(Yh(X)) - c) \right) \\ \leq (2/\gamma) \|Y\|_\infty E_\zeta \left(\sup_{h \in F} \left| \sum_{i=1}^n \zeta_i h(X_i) \right| \right). \quad (\text{a.s.}) \end{aligned}$$

A simple variation of the above argument gives us a bound on the absolute value. In detail,

$$\begin{aligned}
 (3.9) \quad & E_\zeta \left(\sup_{h \in F} \left| \sum_{i=1}^n \zeta_i (\gamma/2) \psi_\gamma(Y_i h(X_i) - E(Yh(X)) - c) \right| \right) \\
 & \leq E_\zeta \left(\sup_{h \in F} \left| \sum_{i=1}^n \zeta_i Y_i h(X_i) \right| \right) + \sup_{h \in F} |E(Yh(X)) - c| E \left(\left| \sum_{i=1}^n \zeta_i \right| \right) \\
 & \leq \|Y\|_\infty E_\zeta \left(\sup_{h \in F} \left| \sum_{i=1}^n \zeta_i h(X_i) \right| \right) + (\|Y\|_\infty \|k\|_\infty^{1/2} + |c|) \sqrt{2\pi n}, \quad (\text{a.s.})
 \end{aligned}$$

where the last inequality follows from integrating a Hoeffding bound on $\Pr(|\sum_{i=1}^n \zeta_i| \geq t)$. Since this holds for all finite F we can take the supremum over finite sets F on both sides and move to $\widetilde{\mathcal{F}}$ (see (1.1)).

Lower bounds. With the Rademacher argument we control the difference between the empirical and population value. To make use of this bound we need a lower bound on the population value. This can be attained in the following way. Let $p = \Pr(\epsilon_1 \geq 0) \wedge \Pr(\epsilon_1 \leq 0)$. Because $E(\epsilon_1) = 0$ it holds that $p > 0$. Using the towering rule for conditional expectations and that ψ_γ is monotonically decreasing, we can now argue in the following way for the population limit and any $h \in \mathcal{H}$,

$$\begin{aligned}
 (3.10) \quad & E(\psi_\gamma(\langle h, Yk(X, \cdot) - \mathbf{m}_y \rangle - c)) \\
 & \geq E(\psi_\gamma(\langle h, Yk(X, \cdot) - \mathbf{m}_y \rangle - c) \times \chi\{\epsilon h(X) \leq 0\}) \\
 & \geq E(\psi_\gamma(\langle h, f_0(X)k(X, \cdot) - \mathbf{m}_y \rangle - c) \\
 & \quad \times (E(\chi\{\epsilon \leq 0\} | X) \times \chi\{h(X) \geq 0\} \\
 & \quad \quad \quad + E(\chi\{\epsilon \geq 0\} | X) \times \chi\{h(X) \leq 0\})) \\
 & \geq pE(\psi_\gamma(\langle h, f_0(X)k(X, \cdot) - \mathbf{m}_y \rangle - c)).
 \end{aligned}$$

The final expectation term can be dealt with in the usual way after moving to the kernel function $l \times k$, where $l = f_0 \otimes f_0$. We demonstrate this for the case that we work with a covariance operator assumption and we derive high probability lower bounds on the radius of a ball centered on $\mathbf{m}_{y,n}$ which lies within $C_{y,n}$.

Assumptions on the covariance operator. We denote the centered covariance operator for the kernel $l \times k$ by $\widetilde{\mathfrak{C}}_{c,l \times k}$. As usual, we need an assumption on the smallest non-zero eigenvalue of this operator. Recall that $\|f_0\|_l = 1$ and

$$\|f_0 \times h\|_{l \times k} = \min \{ \|u\|_{\mathcal{H}_l \otimes \mathcal{H}_k} : f_0(x)h(x) = u(x, x) \text{ for all } x \in \mathcal{X} \}.$$

In particular, $\|f_0 \times h\|_{l \times k} \leq \|f_0 \otimes h\|_{\mathcal{H}_l \otimes \mathcal{H}_k} = \|h\|_k$. In fact, this can be tightened by using [PR16, Proposition 5.20]: the RKHS $\mathcal{H}_{l \times k}$ is the set $\{f_0 \times h : h \in \mathcal{H}\}$ and for any $g, h \in \mathcal{H}$ we have that $\langle f_0 \times g, f_0 \times h \rangle_{l \times k} = \langle g, h \rangle$. In particular, all eigenfunctions

of $\tilde{\mathfrak{C}}_{c,l \times k}$ are of the form $f_0 \times h$ for some $h \in \mathcal{H}, \|h\| = 1$. Therefore, our bounds will depend on

$$(3.11) \quad \bar{\lambda}_* = \inf \left\{ \left\langle \tilde{\mathfrak{C}}_{c,l \times k} f_0 \times h, f_0 \times h \right\rangle_{l \times k} : \|h\| = 1, f_0 \times h \in \left(\ker \tilde{\mathfrak{C}}_{c,l \times k} \right)^\perp \right\}.$$

As before, it is beneficial to move to the RKHS $\mathcal{H}_{l \times k, S}$ corresponding to the kernel function $\kappa = (l \times k) \upharpoonright S \times S$, where S is the support of X . Observe that

$$(l \times k) \upharpoonright S \times S = (l \upharpoonright S \times S) \times (k \upharpoonright S \times S) = ((f_0 \upharpoonright S) \otimes (f_0 \upharpoonright S)) \times (k \upharpoonright S \times S).$$

Using [PR16, Proposition 5.20] again shows that $\mathcal{H}_{l \times k, S} = \{(f_0 \upharpoonright S) \times h : h \in \mathcal{H}_S\}$, where \mathcal{H}_S is defined as before. Furthermore, for $g, h \in \mathcal{H}_S$,

$$\langle (f_0 \upharpoonright S) \times g, (f_0 \upharpoonright S) \times h \rangle_{l \times k, S} = \langle g, h \rangle_S.$$

Let us also introduce $\mathbf{m}_{l \times k, S} = \int \kappa(x, \cdot) dP(x)$, where P is the law of X , and which is well defined whenever $(l \times k)(X, \cdot)$ is Bochner integrable. For $h \in \mathcal{H}$,

$$\begin{aligned} \langle h, \mathbf{m}_y \rangle &= \int_{\mathcal{X}} \langle h, f_0(x)k(x, \cdot) \rangle dP(x) \\ &= \int_S (f_0 \upharpoonright S)(x)(h \upharpoonright S)(x) dP(x) \\ &= \int_S \langle h \upharpoonright S, (f_0 \upharpoonright S)(x)(k \upharpoonright S \times S)(x, \cdot) \rangle_S dP(x) \\ &= \int_S \langle (f_0 \upharpoonright S) \times (h \upharpoonright S), (f_0 \upharpoonright S) \times (f_0 \upharpoonright S)(x)(k \upharpoonright S \times S)(x, \cdot) \rangle_{l \times k, S} dP(x) \\ &= \int_S \langle (f_0 \upharpoonright S) \times (h \upharpoonright S), ((l \times k) \upharpoonright S \times S)(x, \cdot) \rangle_{l \times k, S} dP(x) \\ &= \langle (f_0 \upharpoonright S) \times (h \upharpoonright S), \mathbf{m}_{l \times k, S} \rangle_{l \times k, S}. \end{aligned}$$

The key observation is now the following, for any $h \in \mathcal{H}$ and almost surely

$$\begin{aligned} \langle h, f_0(X)k(X, \cdot) - \mathbf{m}_y \rangle &= (h \upharpoonright S)(X)(f_0 \upharpoonright S)(X) - \langle (f_0 \upharpoonright S) \times (h \upharpoonright S), \mathbf{m}_{l \times k, S} \rangle_{l \times k, S} \\ &= \langle (f_0 \upharpoonright S) \times (h \upharpoonright S), \kappa(X, \cdot) - \mathbf{m}_{l \times k, S} \rangle_{l \times k, S}. \end{aligned}$$

This leads directly to a first result. Under suitable assumptions and with $\delta = \bar{\lambda}_*/2 \| (f_0 \otimes f_0) \times k \|_\infty^{1/2}$, $\gamma = \delta/4$ and $c = -\delta/2$, Eq. (2.5) shows that for any $h \in \mathcal{H}$, such that $\|h \upharpoonright S\|_S = 1$,

$$\begin{aligned} (3.12) \quad E(\psi_\gamma(\langle h, f_0(X)k(X, \cdot) - \mathbf{m}_y \rangle - c)) &= E(\psi_\gamma(\langle (f_0 \upharpoonright S) \times (h \upharpoonright S), \kappa(X, \cdot) - \mathbf{m}_{l \times k, S} \rangle_{l \times k, S} - c)) \\ &\geq \bar{\lambda}_*/8 \| (f_0 \otimes f_0) \times k \|_\infty. \end{aligned}$$

Combining the various steps above leads to the following proposition, which is an adaptation of Theorem 2.17.

PROPOSITION 3.1. — *Let $(\mathcal{X} \times \mathbb{R}, \mathcal{A}, P)$ be some probability space, let P be a topological measure that is τ -additive, and let k be a measurable kernel function defined on \mathcal{X} s.t. the corresponding RKHS \mathcal{H} is finite dimensional. Furthermore, let $(X_1, Y_1), \dots, (X_n, Y_n)$ be i.i.d. random variables attaining values in $\mathcal{X} \times \mathbb{R}$, with law P , and of the form $Y_i = f_0(X_i) + \epsilon_i$ where $\epsilon_1, \dots, \epsilon_n$ are centered i.i.d. random variables which are independent of X_1, \dots, X_n and such that $|\epsilon_1| \leq c_\epsilon$ (a.s.), and*

$f_0 \in \mathcal{L}^2(P)$. Assume that $\|(f_0 \otimes f_0) \times k\|_\infty < \infty$, and that the centered covariance operator $\tilde{\mathfrak{C}}_{c,l \times k}$ has an eigen-decomposition with smallest non-zero eigenvalue being $\bar{\lambda}_*$. Let $\delta = \bar{\lambda}_*/2\|(f_0 \otimes f_0) \times k\|_\infty^{1/2}$ and $p = \Pr(\epsilon_1 \geq 0) \wedge \Pr(\epsilon_1 \leq 0)$. For any $q \in (0, 1)$ and whenever n is (strictly) greater than

$$\begin{aligned} & \frac{1024\|f_0 \otimes f_0 \times k\|_\infty^2}{p^2 \bar{\lambda}_*^4 \delta^2} \left(16 \left(8\|f_0 \otimes f_0 \times k\|_\infty + 6\sqrt{2 \log(3/q)} \right)^2 \right. \\ & \quad \vee 16 \left(4c_\epsilon \|k\|_\infty^{1/2} \log(6/q) + \left((4c_\epsilon^2 \|k\|_\infty + \sqrt{1 + c_\epsilon^2 \|k\|_\infty}) \log(6/q) \right)^{1/2} \right)^2 \\ & \quad \left. \vee \left(\delta \sqrt{2 \log(6/q)} + 16\sqrt{2\pi} \left((\|f_0\|_\infty + c_\epsilon) \|k\|_\infty^{1/2} + \delta/2 \right) + 32\|k\|_\infty^{1/2} (\|f_0\|_\infty + c_\epsilon) \right)^2 \right) \end{aligned}$$

then, with probability at least $1 - q$, there exists a ball of radius $\delta/4$ around $\mathbf{m}_{y,n}$ in $C_{y,n}$ within the affine subspace spanned by $C_{y,n}$.

Proof.

(a). — Since $f_0 \otimes f_0 \times k$ is a bounded kernel function we can apply (3.12) and conclude for $\gamma = \delta/4$ and $c = -\delta/2$ that $E(\psi_\gamma(\langle h, f_0(X)k(X, \cdot) - \mathbf{m}_y \rangle - c)) \geq \bar{\lambda}_*/8\|(f_0 \otimes f_0) \times k\|_\infty$ and from (3.10) it follows that

$$E(\psi_\gamma(\langle h, Yk(X, \cdot) - \mathbf{m}_y \rangle - c)) \geq p\bar{\lambda}_*/8\|(f_0 \otimes f_0) \times k\|_\infty.$$

(b). — Next, we have to incorporate a few triangle inequalities. First, we fill in the details in (3.5). From (3.7) it follows that for $p_1 \in (0, 1)$ and $\alpha_1 = 8n^{-1/2}\|f_0 \otimes f_0 \times k\|_\infty + 6\sqrt{2 \log(1/p_1)}n^{-1/2}$,

$$\Pr\left(\sup_{h \in \tilde{\mathcal{F}}} Pf_0 \times h - P_n f_0 \times h \geq \frac{\alpha_1}{2}\right) = \Pr\left(\sup_{g \in \tilde{\mathcal{F}}_{l \times k}} |Pg - P_n g| \geq \frac{\alpha_1}{2}\right) \leq p_1.$$

Also note that we can simplify (3.6) to

$$\Pr\left(\left\|\frac{1}{n} \sum_{i=1}^n \epsilon_i \phi(X_i)\right\| \geq \frac{\alpha_2}{2}\right) \leq 2 \exp\left(-\frac{n\alpha_2^2}{4c_\epsilon \|k\|_\infty^{1/2} (c_\epsilon \|k\|_\infty^{1/2} + \alpha_2 + \sqrt{1 + c_\epsilon^2 \|k\|_\infty})}\right),$$

whenever $n \geq \alpha_2 c_\epsilon \|k\|_\infty^{1/2}$. Hence, for such n , for $p_2 \in (0, 1)$ and

$$\alpha_2 = \frac{4c_\epsilon \|k\|_\infty^{1/2} \log(2/p_2)}{n} + \left(\frac{(4c_\epsilon^2 \|k\|_\infty + \sqrt{1 + c_\epsilon^2 \|k\|_\infty}) \log(2/p_2)}{n}\right)^{1/2}$$

it follows that

$$\Pr\left(\left\|\frac{1}{n} \sum_{i=1}^n \epsilon_i \phi(X_i)\right\| \geq \frac{\alpha_2}{2}\right) \leq p_2.$$

In particular, for $\alpha_{12} = \alpha_1 \vee \alpha_2$ and whenever $n \geq \alpha_2 c_\epsilon \|k\|_\infty^{1/2}$,

$$\Pr(\|\mathbf{m}_y - \mathbf{m}_{y,n}\| \geq \alpha_{12}) \leq p_1 + p_2.$$

(c). — From (3.8) we can infer that almost surely

$$\begin{aligned} & P\psi_\gamma(\langle h, yk(x, \cdot) - \mathbf{m}_y \rangle - c) - P_n\psi_\gamma(\langle h, \tilde{y}k(\tilde{x}, \cdot) - \mathbf{m}_{y,n} \rangle - c) \\ & \leq \gamma^{-1}\|\mathbf{m}_{y,n} - \mathbf{m}_y\| + P\psi_\gamma(\langle h, yk(x, \cdot) - \mathbf{m}_y \rangle - c) - P_n\psi_\gamma(\langle h, \tilde{y}k(\tilde{x}, \cdot) - \mathbf{m}_y \rangle - c). \end{aligned}$$

The same inequality holds almost surely if we consider the supremum over $\widetilde{\mathcal{F}}$. Therefore, with $\gamma = \delta/4$ and $\alpha > 0$,

$$\begin{aligned} & \Pr\left(\sup_{h \in \widetilde{\mathcal{F}}} P\psi_\gamma(\langle h, yk(x, \cdot) - \mathbf{m}_y \rangle - c) - P_n\psi_\gamma(\langle h, \widetilde{y}k(\widetilde{x}, \cdot) - \mathbf{m}_{y,n} \rangle - c) \geq 2\alpha\right) \\ & \leq \Pr(\|\mathbf{m}_{y,n} - \mathbf{m}_y\| \geq \delta\alpha/4) \\ & \quad + \Pr\left(\sup_{h \in \widetilde{\mathcal{F}}} P\psi_\gamma(\langle h, yk(x, \cdot) - \mathbf{m}_y \rangle - c) - P_n\psi_\gamma(\langle h, \widetilde{y}k(\widetilde{x}, \cdot) - \mathbf{m}_y \rangle - c) \geq \alpha\right). \end{aligned}$$

The latter term can be dealt with by a Rademacher argument when using (3.9). In detail, for any $p_3 \in (0, 1)$, with probability $1 - p_3$ simultaneously for all $h \in \widetilde{\mathcal{F}}$,

$$\begin{aligned} P_n\psi_\gamma(\langle h, \widetilde{y}k(\widetilde{x}, \cdot) - \mathbf{m}_y \rangle - c) & \geq P\psi_\gamma(\langle h, yk(x, \cdot) - \mathbf{m}_y \rangle - c) - \sqrt{2 \log(2/p_3)/n} \\ & \quad - 2E_\zeta\left(\sup_{h \in \widetilde{\mathcal{F}}}\left|\frac{1}{n} \sum_{i=1}^n \zeta_i \psi_\gamma(Y_i h(X_i) - E(Yh(X))) - c\right|\right) \end{aligned}$$

follows from [GN16, Theorem 3.4.5]; also see p. 496. Substituting (3.9) leads to the following lower bound on the P_n term,

$$\begin{aligned} & P\psi_\gamma(\langle h, yk(x, \cdot) - \mathbf{m}_y \rangle - c) - \sqrt{2 \log(2/p_3)/n} \\ & \quad - (4/\gamma)\left(\|Y\|_\infty E_\zeta\left(\sup_{h \in \widetilde{\mathcal{F}}}\left|\frac{1}{n} \sum_{i=1}^n \zeta_i h(X_i)\right|\right)\right) + (\|Y\|_\infty \|k\|_\infty^{1/2} + |c|)\sqrt{2\pi/n}. \end{aligned}$$

Filling in γ, c , the upper bound on $|Y|$ and the Rademacher complexity of $\widetilde{\mathcal{F}}$, reduces the lower bound to

$$\alpha_3 = \frac{\sqrt{2 \log(2/p_3)}}{\sqrt{n}} + \frac{16\sqrt{2\pi}(\|f_0\|_\infty + c_\epsilon)\|k\|_\infty^{1/2} + \delta/2}{\delta\sqrt{n}} + \frac{32\|k\|_\infty^{1/2}(\|f_0\|_\infty + c_\epsilon)}{\delta\sqrt{n}}$$

and

$$\Pr\left(\sup_{h \in \widetilde{\mathcal{F}}} P\psi_\gamma(\langle h, yk(x, \cdot) - \mathbf{m}_y \rangle - c) - P_n\psi_\gamma(\langle h, \widetilde{y}k(\widetilde{x}, \cdot) - \mathbf{m}_y \rangle - c) \geq \alpha_3\right) \leq p_3.$$

(d). — Combining these bounds we can derive a lower bound on $P_n\psi_\gamma(\langle h, \widetilde{y}k(\widetilde{x}, \cdot) - \mathbf{m}_{y,n} \rangle - c)$. In detail, let $p_1 = p_2 = p_3 = q/3$ and set

$$\alpha^* = \frac{4\alpha_{12}}{\delta} \vee \alpha_3$$

then with probability $1 - q$ simultaneously for all $h \in \widetilde{\mathcal{F}}$,

$$P_n\psi_\gamma(\langle h, \widetilde{y}k(\widetilde{x}, \cdot) - \mathbf{m}_{y,n} \rangle - c) \geq p\bar{\lambda}_*^2/8\|(f_0 \otimes f_0) \times k\|_\infty - 2\alpha^*.$$

To guarantee that the right hand side is strictly positive we can choose the n which is provided in the statement of the Proposition 3.1. \square

3.3.2. Second setting: bounded Y_i 's & Eq. (3.3)

In this section, we map $h \in \mathcal{H}$ to $\check{h} = \langle 1, \cdot \rangle_{\mathbb{R}} \otimes h(\cdot)$ and we work with the kernel function

$$\rho((y_1, x_1), (y_2, x_2)) = \langle y_1, y_2 \rangle_{\mathbb{R}} k(x_1, x_2).$$

In the introduction to this section we denoted the RKHS \mathcal{H}_ρ by $\mathbb{R}' \otimes \mathcal{H}$. We will use in the following the more compact notation $\mathcal{H}_\rho, \|\cdot\|_\rho$ etc.

Compared to the approach in the previous section, using \mathcal{H}_ρ offers a dramatic simplification of the analysis and leads to improved bounds. As usual, we are interested in approximating a mean element. In the current context, this is $\mathbf{m}_y^\otimes = \int \rho((y, x), \cdot) dP(x, y)$ and there is a straight forward relation to the element \mathbf{m}_y that was used in the previous section,

$$\mathbf{m}_y^\otimes = \int \langle y, \cdot \rangle_{\mathbb{R}} \otimes k(x, \cdot) dP(x, y) = \langle 1, \cdot \rangle_{\mathbb{R}} \otimes \int yk(x, \cdot) dP(x, y) = \langle 1, \cdot \rangle_{\mathbb{R}} \otimes \mathbf{m}_y,$$

using (1.2). The element \mathbf{m}_y^\otimes lies in $\mathbb{R}' \otimes \mathcal{H}$ whenever $Y^2k(X, X) \in \mathcal{L}^1(\mu)$. Under our assumption that $Y = f_0(X) + \epsilon$, X and ϵ independent random variables, the representation of \mathbf{m}_y^\otimes simplifies to

$$\begin{aligned} \mathbf{m}_y^\otimes &= E(\langle f_0(X) + \epsilon, \cdot \rangle_{\mathbb{R}} \otimes k(X, \cdot)) \\ &= E(f_0(X)\langle 1, \cdot \rangle_{\mathbb{R}} \otimes k(X, \cdot)) + \langle E(\epsilon), \cdot \rangle \otimes E(k(X, \cdot)) \\ &= \langle 1, \cdot \rangle_{\mathbb{R}} \otimes E(f_0(X)k(X, \cdot)), \end{aligned}$$

where we used (1.2) in the second and in the last equality. This is again just $\langle 1, \cdot \rangle_{\mathbb{R}} \otimes \mathbf{m}_y$.

If P is τ -additive as a measure on $\mathcal{X} \times \mathbb{R}$ then the support S of P is well defined and we have a natural population limit

$$C_y^\otimes = \text{cch}\{\langle y, \cdot \rangle_{\mathbb{R}} \otimes k(x, \cdot) : (x, y) \in S\}$$

of the empirical convex set

$$C_{y,n}^\otimes = \text{cch}\{\langle Y_i, \cdot \rangle_{\mathbb{R}} \otimes k(X_i, \cdot) : i \leq n\}.$$

These are just the convex sets associated with the kernel function ρ acting on $\mathbb{R} \times \mathcal{X}$ and $\mathbf{m}_y^\otimes, \mathbf{m}_{y,n}^\otimes$ are the corresponding mean and empirical mean elements. In fact, we can apply right away Theorem 2.17. Our sample space is then $[-\|f_0\|_\infty - c_\epsilon, \|f_0\|_\infty + c_\epsilon] \times \mathcal{X}$, where c_ϵ is a constant such that $|Y| \leq c_\epsilon$ a.s. The kernel function is ρ restricted to the sample space and

$$\|\rho|_S \times S\|_\infty \leq (\|f_0\|_\infty + c_\epsilon)\|k\|_\infty.$$

In this formulation it might not be directly obvious how assumptions on the distribution of Y enter. Using a Rademacher argument we can control the difference between P and P_n when acting on indicator functions. To do so we do not need any assumption on the distribution of Y beside boundedness. But, if you recall our earlier arguments, you will notice that we used lower bounds on P when applied to indicator functions to control the size of $C_{y,n}^\otimes$. This lower bound on P depends on the distribution of Y . In particular, with the covariance operator approach, it depends on the variance of Y .

Let $\tilde{\mathfrak{C}}_c^\otimes$ be the centered covariance operator corresponding to kernel ρ then for $g, h \in \mathcal{H}$,

$$\begin{aligned} \langle \tilde{\mathfrak{C}}_c^\otimes \check{g}, \check{h} \rangle_\rho &= E(Y^2 g(X)h(X)) - E(Yg(X))E(Yh(X)) \\ &= E(f_0^2(X)g(X)h(X)) + \sigma^2 E(g(X)h(X)) - E(f_0(X)g(X))E(f_0(X)h(X)), \end{aligned}$$

where σ^2 is the variance of ϵ . We can also relate this expression back to the covariance operator discussed in the earlier approach, for $h \in \mathcal{H}$,

$$\langle \tilde{\mathfrak{C}}_c^\otimes \check{g}, \check{h} \rangle_\rho = \langle \tilde{\mathfrak{C}}_{c,l \times k} f_0 \otimes g, f_0 \otimes h \rangle_{l \times k} + \sigma^2 E(g(X))E(h(X)).$$

Recall that for $g, h \in \mathcal{H}$, $\langle g, h \rangle = \langle f_0 \times g, f_0 \times h \rangle_{l \times k}$ and note that $\langle \check{g}, \check{h} \rangle_\rho = \langle g, h \rangle$. In particular,

$$(3.13) \quad \|h\| = \|\check{h}\|_\rho = \|f_0 \times h\|_{l \times k}$$

for all $h \in \mathcal{H}$. If $h \in \mathcal{H}$, $\|h\| = 1$, is such that $f_0 \times h$ is an eigenfunction of $\tilde{\mathfrak{C}}_{c,l \times k}$ with eigenvalue λ and σ^2 is the variance of ϵ then

$$\langle \tilde{\mathfrak{C}}_c^\otimes \check{h}, \check{h} \rangle_\rho = \lambda + \sigma^2 E(h^2(X)).$$

Furthermore, if $f_0 \times h$ is an eigenfunction of $\tilde{\mathfrak{C}}_{c,l \times k}$ and g is such that $\langle g, h \rangle = 0$ then $\langle \tilde{\mathfrak{C}}_c^\otimes \check{h}, \check{g} \rangle_\rho = \langle \tilde{\mathfrak{C}}_{c,l \times k} f_0 \times h, f_0 \times g \rangle_{l \times k} + \sigma^2 E(g(X))E(h(X)) = \sigma^2 E(g(X))E(h(X))$.

There is no reason why the latter term should be zero and the two operators will generally not have the same eigenfunctions (in the sense that \check{h} is an eigenfunction of $\tilde{\mathfrak{C}}_c^\otimes$ iff $f_0 \otimes h$ is an eigenfunction of $\tilde{\mathfrak{C}}_{c,l \times k}$).

Remark 3.2. — If $v = \inf\{E(h^2(X)) : h \in \mathcal{H}, \|h\| = 1\} > 0$ then $\tilde{\mathfrak{C}}_c^\otimes$ has no eigenvalue below $\sigma^2 v$. This can help with the compression, but notice that larger values of σ^2 are related to larger values $\|\rho\|_\infty$ which hinders the compression.

Lower bounds on the width. We could also look at the width of the convex set C_y^\otimes by means of the kernel function ρ . While this is a useful exercise we only want to highlight here a simple relation between the width of C_y^\otimes and the width of usual convex set C_S (as a subset of \mathcal{H}_S). For $h \in \mathcal{H}$, $\|h\| = 1$,

$$\begin{aligned} \text{width}_{\check{h}}(C_y^\otimes) &= \sup_{(x,y) \in S} (y - f_0(x))\langle h, k(x, \cdot) \rangle + h(x)f_0(x) \\ &\quad - \inf_{(x',y') \in S} ((y' - f_0(x'))\langle h, k(x', \cdot) \rangle + h(x')f_0(x')) \\ &\geq \sup_{x \in \mathcal{X}_S} \sup_{y,y' \in S_x} (y - y')h(x) = \sup_{x \in \mathcal{X}_S} \sup_{y,y' \in S_x} (y - y')|h(x)|, \end{aligned}$$

where $\mathcal{X}_S = \{x : (x, y) \in S\}$ and $S_x = \{y : (x, y) \in S\}$. Also, note that

$$\sup_{x \in \mathcal{X}_S} |h(x)| \geq (1/2) \text{width}_h(C_S)$$

and we have a lower bound on $\text{width}_{\check{h}}(C_y^\otimes)$ which is a product of the width of C_S and the spread of ϵ .

3.3.3. Third setting: unbounded Y_i 's & Eq. (3.3)

If Y is unbounded then one way to approach the approximation problem is to cap the observations Y_i and to control the cap as a function of n . We demonstrate this for the case that we have the model stated in (3.4), f_0 is *bounded* and *measurable*, and ϵ is *sub-Gaussian* with variance factor ν (see [BLM13, Section 2.3]) but is not necessarily bounded. The natural sample space is now $\mathbb{R} \times \mathcal{X}$. The kernel ρ is well defined on $\mathbb{R} \times \mathcal{X}$ but is unbounded. When f_0 is bounded, k is measurable and bounded, and ϵ is sub-Gaussian, we have that $\langle Y, \cdot \rangle_{\mathbb{R}} k(X, \cdot) \in \mathcal{L}^2(\mu, \mathcal{H}_\rho)$ since

$$E(\|\langle Y, \cdot \rangle_{\mathbb{R}} k(X, \cdot)\|_\rho^2) = E(Y^2 k(X, X)) \leq \|k\|_\infty E(Y^2)$$

and the latter term is finite since Y is sub-Gaussian. Let $\{r_n\}_{n \geq 1}$ be a non-negative and non-decreasing sequence, and let $\widehat{Y}^{(n)} = (Y \wedge (r_n + \|f_0\|_\infty)) \vee -(r_n + \|f_0\|_\infty)$. The $\|f_0\|_\infty$ can obviously be replaced by an upper bound on the norm, but as the argument is developed such a bound is needed to control the error introduced by capping the observations Y_i . Define $\widehat{\mathbf{m}}_{y,n}^\otimes = (1/n) \sum_{i=1}^n \langle \widehat{Y}_i^{(n)}, \cdot \rangle_{\mathbb{R}} k(X_i, \cdot)$. In this section our aim is to derive a suitable adapted version of Theorem 2.17 for this setting where ϵ is sub-Gaussian. We start by investigating the effect of the capping of Y .

Bounding $\|\mathbf{m}_{y,n}^\otimes - \widehat{\mathbf{m}}_{y,n}^\otimes\|$. A simple expansion yields

$$\begin{aligned} \|\mathbf{m}_{y,n}^\otimes - \widehat{\mathbf{m}}_{y,n}^\otimes\|_\rho^2 &= \frac{1}{n^2} \left\| \sum_{i=1}^n \langle Y_i - \widehat{Y}_i^{(n)}, \cdot \rangle_{\mathbb{R}} k(X_i, \cdot) \right\|_\rho^2 \\ &= \frac{1}{n^2} \sum_{i,j=1}^n (Y_i - \widehat{Y}_i^{(n)}) (Y_j - \widehat{Y}_j^{(n)}) k(X_i, X_j). \end{aligned}$$

Due to the independence of the observations and by using the Cauchy–Schwarz inequality,

$$\begin{aligned} E\left(\|\mathbf{m}_{y,n}^\otimes - \widehat{\mathbf{m}}_{y,n}^\otimes\|_\rho^2\right) &\leq \frac{1}{n^2} \sum_{i=1}^n \left| E(Y_i - \widehat{Y}_i^{(n)})^2 k(X_i, X_i) \right| \\ &\quad + \frac{1}{n^2} \sum_{i \neq j} \left| \langle E((Y_i - \widehat{Y}_i^{(n)}) k(X_i, \cdot)), E((Y_j - \widehat{Y}_j^{(n)}) k(X_j, \cdot)) \rangle \right| \\ &\leq \frac{\|k\|_\infty}{n^2} \sum_{i=1}^n E(Y_i - \widehat{Y}_i^{(n)})^2 + \frac{n-1}{n} E^2\left(\left|Y_1 - \widehat{Y}_1^{(n)}\right| \|k(X_1, \cdot)\|\right) \\ &\leq \frac{\|k\|_\infty}{n^2} \sum_{i=1}^n E(Y_i - \widehat{Y}_i^{(n)})^2 + \frac{(n-1)\|k\|_\infty}{n} E^2\left(\left|Y_1 - \widehat{Y}_1^{(n)}\right|\right). \end{aligned}$$

Let λ denote the law of ϵ . Noting that ϵ is sub-Gaussian and using [BLM13, Section 2.3] together with [Fre01, 252O] and [Dud14, Proposition 2.5 (a)], we obtain

$$\begin{aligned}
 (3.14) \quad E\left(|Y_i - \widehat{Y}_i^{(n)}|\right) &\leq \int_0^\infty (t - r_n) \times \chi\{t \geq r_n\} d\lambda(t) + \int_{-\infty}^0 (-t - r_n) \times \chi\{-t \geq r_n\} d\lambda(t) \\
 &\leq \int_0^\infty \lambda\{t : t \geq s + r_n\} ds + \int_0^\infty \lambda\{t : -t \geq s + r_n\} ds \\
 &\leq 2 \int_{r_n}^\infty e^{-s^2/2\nu} ds \leq \sqrt{2\pi\nu} e^{-r_n^2/2\nu}.
 \end{aligned}$$

Similarly,

$$\begin{aligned}
 (3.15) \quad E\left(Y_i - \widehat{Y}_i^{(n)}\right)^2 &\leq \int_0^\infty (t - r_n)^2 \times \chi\{t \geq r_n\} d\lambda(t) + \int_{-\infty}^0 (-t - r_n)^2 \times \chi\{-t \geq r_n\} d\lambda(t) \\
 &\leq \int_0^\infty \lambda\{t : t \geq \sqrt{s} + r_n\} ds + \int_0^\infty \lambda\{t : t \leq -\sqrt{s} - r_n\} ds \\
 &\leq 2e^{-r_n^2/2\nu} \int_0^\infty e^{-s/2\nu} ds = 4\nu e^{-r_n^2/2\nu}.
 \end{aligned}$$

Combining these yields

$$E\left(\|\mathbf{m}_{y,n}^\otimes - \widehat{\mathbf{m}}_{y,n}^\otimes\|_\rho^2\right) \leq 2\nu \|k\|_\infty e^{-r_n^2/2\nu} \left(2/n + \pi e^{-r_n^2/2\nu}\right)$$

and

$$\Pr\left(\|\mathbf{m}_{y,n}^\otimes - \widehat{\mathbf{m}}_{y,n}^\otimes\|_\rho \geq t\right) \leq \frac{2\nu \|k\|_\infty e^{-r_n^2/2\nu} \left(2/n + \pi e^{-r_n^2/2\nu}\right)}{t^2}.$$

In other words, if we have an upper bound on $\|f_0\|_\infty$ and cap the observations as described above then with probability $1 - \delta$ for any $\delta \in (0, 1)$,

$$(3.16) \quad \|\mathbf{m}_{y,n}^\otimes - \widehat{\mathbf{m}}_{y,n}^\otimes\|_\rho \leq \sqrt{2\nu} \|k\|_\infty^{1/2} e^{-r_n^2/4\nu} \left(2/n + \pi e^{-r_n^2/2\nu}\right)^{1/2} \delta^{-1/2}.$$

Spectrum of the covariance operator. In Remark 3.2 we observed that no eigenvalue of $\widetilde{\mathfrak{C}}_c^\otimes$ can be lower than $\sigma^2 v$, where σ^2 is the variance of ϵ and $v = \inf\{E(h^2(X)) : h \in \mathcal{H}, \|h\| = 1\}$. There is a peculiar detail that we have to be careful about: if $h_1, h_2 \in \mathcal{H}$ are linearly independent but $h_1|_{\mathcal{X}_S} = h_2|_{\mathcal{X}_S}$, where $\mathcal{X}_S = \{x : (x, y) \in S\}$, then there exists an $h \in \mathcal{H}$ of norm one for which $E(h^2(X)) = 0$. Furthermore, the corresponding functions \check{h}_1, \check{h}_2 in \mathcal{H}_ρ are not constant on S but some linear combination of \check{h}_1 and \check{h}_2 is zero on S (at least when the involved functions are continuous). To make use of the lower bound $\sigma^2 v$ it makes therefore sense to move right away to functions restricted to $S_f = \{(y, x) : (x, y) \in S\}$ or \mathcal{X}_S .

Before coming back to the lower bound we want to take a paragraph to understand better \mathcal{H}_{ρ, S_f} , which is the RKHS corresponding to the kernel function $\rho|_{S_f \times S_f}$. First, notice that there are no (non-zero) constant functions in \mathcal{H}_{ρ, S_f} if ϵ is not almost surely zero. In particular, the covariance operator $\widetilde{\mathfrak{C}}_c^{\otimes, S}$ has then only eigenvalues

that are strictly positive. If $\tilde{\mathfrak{C}}_c^{\otimes}$ has zero eigenvalues then there must be elements \check{h} which are constant on S but this means that these elements have to be equal to zero on S and correspond to the origin in \mathcal{H}_{ρ,S_f} . Also notice that \mathcal{H}_{ρ,S_f} is not the same RKHS as the RKHS with kernel function $(l|\mathbb{R}_S \times \mathbb{R}_S) \times (k|\mathcal{X}_S \times \mathcal{X}_S)$, where $\mathbb{R}_S = \{y : (x, y) \in S\}$ and $l(y, y') = \langle y, y' \rangle_{\mathbb{R}}$ for all $y, y' \in \mathbb{R}$. This follows directly since they have different domains. The latter kernel is defined for pairs (x, y) in $\mathcal{X}_S \times \mathbb{R}_S$ while the former is defined for pairs $(x, y) \in S$. This is inconvenient since we like to use $E(h^2(X))$ for h in some RKHS of functions acting on the support of some measure and it is not directly obvious what this support should be like. In the following, let $\mathfrak{T}_{\mathcal{X}}$ be a topology on \mathcal{X} , let \mathfrak{T} be the corresponding product topology on $\mathcal{X} \times \mathbb{R}$, and assume that the law P of (X, Y) is a *Radon measure* with σ -algebra \mathcal{A} ; in particular, it is a topological τ -additive measure and $\mathfrak{T} \subset \mathcal{A}$. Then S is well defined as a subset of $\mathcal{X} \times \mathbb{R}$. Let us also introduce our probability space (Ω, Σ, μ) , assume that μ is *complete* and (X, Y) is a well defined random variable in the sense that $(X, Y)^{-1}[A] \in \Sigma$ for all $A \in \mathcal{A}$. Furthermore, consider the σ -algebra $\mathcal{A}_{\mathcal{X}} = \{A : A \times \mathbb{R} \in \mathcal{A}\}$ and let $P_X = P \circ \pi_{\mathcal{X}}^{-1}$ where the function $\pi_{\mathcal{X}} : \mathcal{X} \times \mathbb{R} \rightarrow \mathcal{X}$ projects onto the first coordinate. The σ -algebra $\mathcal{A}_{\mathcal{X}}$ contains $\mathfrak{T}_{\mathcal{X}}$ since for $O \in \mathfrak{T}_{\mathcal{X}}$ it holds that $O \times \mathbb{R}$ is in the product topology \mathfrak{T} . Hence, P_X is a topological measure. If \mathcal{X} is a *Hausdorff space* then P_X is, in fact, a Radon measure (apply [Fre01, 418I] to $\pi_{\mathcal{X}}$ and note that $\pi_{\mathcal{X}}$ is continuous). This implies that the support $S' \subset \mathcal{X}$ of P_X is well defined. Observe that $\overline{\mathcal{X}_S} = S'$: the projection $\pi_{\mathcal{X}}$ is a continuous inverse-measure preserving function from $\mathcal{X} \times \mathbb{R}$ to \mathcal{X} and, due to [Fre01, 411N.e], the support of P_X is $\pi_{\mathcal{X}}[S] = \{x : (x, y) \in S\} = \overline{S_{\mathcal{X}}}$. We also have to check that X is actually a well defined random variable in the sense that $X^{-1}[A] \in \Sigma$ for all $A \in \mathcal{A}_{\mathcal{X}}$, and that $E(h^2(X)) = \int h^2 dP_X = \int (h|\overline{\mathcal{X}_S})^2 dP_X$. The former can be seen in the following way. Since $\pi_{\mathcal{X}}$ is a measurable function from $(\mathcal{X} \times \mathbb{R}, \mathcal{A})$ to $(\mathcal{X}, \mathcal{A}_{\mathcal{X}})$ it follows that $X = \pi_{\mathcal{X}} \circ (X, Y)$ is measurable as a mapping from (Ω, Σ) to $(\mathcal{X}, \mathcal{A}_{\mathcal{X}})$ and is a well defined random variable. For the latter, if h is in $\mathcal{L}^2(\mathcal{X}, P_X)$ then $\int h^2 dP_X$ is well defined and obviously equal to $\int (h|\overline{\mathcal{X}_S})^2 dP_X$. It remains to show the $E(h^2(X)) = \int h^2 dP_X$. One way to show this is to use [Fre01, 235E]. This can be applied since $\mu X^{-1}[A] = \Pr(X \in A, Y \in \mathbb{R}) = P\pi_{\mathcal{X}}^{-1}[A] = P_X(A)$ for any $A \in \mathcal{A}_{\mathcal{X}}$.

We can now define an RKHS of functions that act on the support $\overline{\mathcal{X}_S}$. For compactness of notation let $\tilde{S} = \overline{\mathcal{X}_S}$, let $k_{\tilde{S}} = k|\tilde{S} \times \tilde{S}$ and denote the corresponding RKHS by $\mathcal{H}_{\tilde{S}}$. This RKHS allows us to carry over Remark 3.2 to the case where we work with \mathcal{H}_{ρ,S_f} . Let $\tilde{\mathfrak{C}}_c^{\otimes,S}$ be the centered covariance operator that corresponds to \mathcal{H}_{ρ,S_f} . Since, $\mathcal{H}_{\rho,S_f} = \{\check{h}|S_f : h \in \mathcal{H}\}$, it follows that

$$\begin{aligned}
 (3.17) \quad & \left\langle \tilde{\mathfrak{C}}_c^{\otimes,S} \check{h}|S_f, \check{h}|S_f \right\rangle_{\rho,S_f} \\
 &= E(f_0^2(X)h^2(X)) + \sigma^2 E(h^2(X)) - E^2(f_0(X)h(X)) \\
 &= E(f_0^2(X)h^2(X)) + \sigma^2 \left\langle \tilde{\mathfrak{C}}^{\tilde{S}} h|\tilde{S}, h|\tilde{S} \right\rangle_{\tilde{S}} - E^2(f_0(X)h(X)).
 \end{aligned}$$

When the kernel function k is *continuous*, then the only function $h \in \mathcal{H}_{\tilde{S}}$ for which $E(h^2(X)) = 0$ is $h = 0$ which has norm zero (otherwise there is an open set

on which $h^2(X)$ is bounded away from zero and the intersection of this open set with the support has measure strictly larger than zero [Fre01, 411N]). In this case, all the eigenvalues of $\tilde{\mathfrak{C}}^{\tilde{S}}$ are strictly positive and also all eigenvalues of $\tilde{\mathfrak{C}}_c^{\otimes, S}$ are strictly positive, implying that there is no constant function in \mathcal{H}_{ρ, S_f} . Also, note that $\|h\upharpoonright\tilde{S}\|_{\tilde{S}} = \|\check{h}\upharpoonright S_f\|_{\rho, S_f}$ for all $h \in \mathcal{H}$: first observe that $\check{g}\upharpoonright S_f = \check{h}\upharpoonright S_f$ if, and only if, $g\upharpoonright\tilde{S} = h\upharpoonright\tilde{S}$. If $g\upharpoonright\tilde{S} = h\upharpoonright\tilde{S}$ then for $(x, y) \in S$, $\check{g}(y, x) = yg(x) = yh(x) = \check{h}(y, x)$ because $x \in \mathcal{X}_S \subset \tilde{S}$. On the other hand, if $\check{g}\upharpoonright S_f = \check{h}\upharpoonright S_f$ then for any $x \in \mathcal{X}_S$ there exist a point y such that $(x, y) \in S$ and $yg(x) = yh(x)$ which implies $g(x) = h(x)$ if $y \neq 0$. In fact, if $\sigma^2 > 0$, there exist at least two such points and, in particular, there exists a $y \neq 0$ such that $yg(x) = yh(x)$. Since $g = h$ on the dense subset \mathcal{X}_S of \tilde{S} and both g, h are continuous (assuming k is continuous) it follows that $g = h$ on \tilde{S} (e.g. [Eng89, Theorem 1.5.4] and using that \mathbb{R} is a Hausdorff space). Using (3.13),

$$\begin{aligned}
 (3.18) \quad \|\check{h}\upharpoonright S_f\|_{\rho, S_f} &= \inf \left\{ \|\check{g}\|_{\rho} : \check{g}\upharpoonright S_f = \check{h}\upharpoonright S_f, g \in \mathcal{H} \right\} \\
 &= \inf \left\{ \|g\| : g\upharpoonright S_f = \check{h}\upharpoonright S_f, g \in \mathcal{H} \right\} \\
 &= \inf \left\{ \|g\| : g\upharpoonright\tilde{S} = h\upharpoonright\tilde{S}, g \in \mathcal{H} \right\} = \|h\upharpoonright\tilde{S}\|_{\tilde{S}}.
 \end{aligned}$$

This implies that a strictly positive lower bound on the $\langle \tilde{\mathfrak{C}}^{\tilde{S}} h\upharpoonright\tilde{S}, h\upharpoonright\tilde{S} \rangle_{\tilde{S}}$ is given by the smallest eigenvalue of $\tilde{\mathfrak{C}}^{\tilde{S}}$ (when \mathcal{H} is finite dimensional). We can also express this bound in terms of \mathcal{H} since for $h \in \mathcal{H}$, $\langle \tilde{\mathfrak{C}}h, h \rangle = \langle \tilde{\mathfrak{C}}^{\tilde{S}} h\upharpoonright\tilde{S}, h\upharpoonright\tilde{S} \rangle_{\tilde{S}}$ and $\|h\| \geq \|h\upharpoonright\tilde{S}\|_{\tilde{S}}$ it follows that the smallest eigenvalue of $\tilde{\mathfrak{C}}$ provides a lower bound on the smallest eigenvalue of $\tilde{\mathfrak{C}}^{\tilde{S}}$ and this lower bound is strictly positive. We might be tempted to improve this lower bound by recalling that the eigenvalues of $\tilde{\mathfrak{C}}_c$ tell us the dimension of $\mathcal{H}_{\tilde{S}}$, but notice that there is no reason why the eigenfunctions of $\tilde{\mathfrak{C}}_c$ and $\tilde{\mathfrak{C}}$ should be related since one corresponds to the variance and the other corresponds to the second moment, and it is not directly obvious of how to benefit from the additional information that $\tilde{\mathfrak{C}}_c$ provides.

A family of covariance operators. The move from Y to $\widehat{Y}^{(n)}$ affects the covariance, and the covariance operator corresponding to $\langle Y, \cdot \rangle_{\mathbb{R}} k(X, \cdot)$ is not the same as the covariance operator corresponding to $\langle \widehat{Y}^{(n)}, \cdot \rangle_{\mathbb{R}} k(X, \cdot)$, $n \geq 1$. Let us denote the covariance operators corresponding to the \widehat{Y} 's by the somewhat unwieldy $\widehat{\mathfrak{C}}_c^{\otimes, (n)}$ and the covariance operator corresponding to the support S_n of the law of $(X, \widehat{Y}^{(n)})$ by $\widehat{\mathfrak{C}}_c^{\otimes, S_n}$; we assume that the laws $P^{(n)}$ of $(X, \widehat{Y}^{(n)})$ are *Radon measures*, which guarantees that *the support of $P^{(n)}$ is well defined*. It is easy to verify that $P^{(n)}$ is a Radon measure if P itself is a Radon measure and the topology corresponding to P is a Hausdorff topology. Consider the set $A = \mathcal{X} \times [-(r_n + \|f_0\|_{\infty}), r_n + \|f_0\|_{\infty}]$ equipped with the subspace topology which is also a Hausdorff topology [Eng89, Theorem 2.1.6], and the continuous function $f : \mathcal{X} \times \mathbb{R} \rightarrow A$ given by $f(x, y) = (x, \widehat{y}^{(n)})$, where we mean the same transformation as for the random variable Y . The push-forward $P^{(n)} = f_{\#}P$ is a Radon measure according to [Fre01, 418I].

We need lower bounds on the smallest non-zero eigenvalues of the different $\widehat{\mathfrak{C}}_c^{\otimes, (n)}$ operators to use our compression approach. It seems natural to work with an assumption on the smallest eigenvalue of the covariance operator $\widetilde{\mathfrak{C}}_c^{\otimes, S}$, which corresponds to the original Y , and to relate the eigenvalues of $\widehat{\mathfrak{C}}_c^{\otimes, S_n}$ back to the eigenvalues of $\widetilde{\mathfrak{C}}_c^{\otimes, S}$. As discussed on the previous page, *the covariance operator $\widetilde{\mathfrak{C}}_c^{\otimes, S}$ does not have an eigenvalue that is zero if k is continuous.* In this case, the smallest eigenvalue $\widehat{\lambda}_\star^{(n)}$ of $\widehat{\mathfrak{C}}_c^{\otimes, S_n}$ is at least of size $\bar{\lambda}_\star/2$, where $\bar{\lambda}_\star$ is the smallest eigenvalue of $\widetilde{\mathfrak{C}}_c^{\otimes, S}$, whenever

$$\left\| \widetilde{\mathfrak{C}}_c^{\otimes, S} - \widehat{\mathfrak{C}}_c^{\otimes, S_n} \right\|_{\text{op}} \leq \frac{\sigma^2 \nu}{2},$$

where (3.17) tells us that we can choose ν either as the smallest eigenvalue of $\widetilde{\mathfrak{C}}^{\mathfrak{S}}$ or the smallest non-zero eigenvalue of $\widetilde{\mathfrak{C}}$, and where $0 < \sigma^2 \leq \nu$ is the variance of ϵ . Alternatively, we can obviously also directly impose assumptions on the eigenvalues of $\widetilde{\mathfrak{C}}_c^{\otimes, S}$. We can bound the operator norm in the following way,

$$\begin{aligned} & \left\| \widetilde{\mathfrak{C}}_c^{\otimes, S} - \widehat{\mathfrak{C}}_c^{\otimes, S_n} \right\|_{\text{op}} \\ &= \sup_{\| \check{h} \|_{S_f} = 1} \sup_{\| \check{g} \|_{S_f} = 1} \left\langle \widetilde{\mathfrak{C}}_c^{\otimes, S} \check{h} \check{g} - \widehat{\mathfrak{C}}_c^{\otimes, S_n} \check{h} \check{g}, \check{h} \check{g} \right\rangle_{\rho, S_f} \\ &= \sup_{\| h \|_{\check{S}} = 1} \sup_{\| g \|_{\check{S}} = 1} E \left((Y^2 - (\widehat{Y}^{(n)})^2) h(X) g(X) \right) - E(Y g(X)) E(Y h(X)) \\ & \qquad \qquad \qquad + E(\widehat{Y}^{(n)} g(X)) E(\widehat{Y}^{(n)} h(X)). \end{aligned}$$

Let us first address the second moment term. For h, g such that $\| h \|_{\check{S}} = 1 = \| g \|_{\check{S}}$,

$$\begin{aligned} & \left| E \left((Y^2 - (\widehat{Y}^{(n)})^2) h(X) g(X) \right) \right| \leq \| k \|_\infty E \left(Y^2 - (\widehat{Y}^{(n)})^2 \right) \\ &= \| k \|_\infty \int_0^\infty (t^2 - r_n^2) \times \chi \{ t \geq r_n \} d\lambda(t) + \| k \|_\infty \int_{-\infty}^0 (t^2 - r_n^2) \times \chi \{ -t \geq r_n \} d\lambda(t) \\ &\leq \| k \|_\infty \int_0^\infty \lambda \{ t : t \geq \sqrt{s + r_n^2} \} ds + \| k \|_\infty \int_0^\infty \lambda \{ t : t \leq -\sqrt{s + r_n^2} \} ds \\ &\leq 2 \| k \|_\infty \int_{r_n^2}^\infty e^{-s/2\nu} ds = 4\nu \| k \|_\infty e^{-r_n^2/2\nu}. \end{aligned}$$

The other term can be controlled in the following way,

$$\begin{aligned} & \left| E(Y g(X)) E(Y h(X)) - E(\widehat{Y}^{(n)} g(X)) E(\widehat{Y}^{(n)} h(X)) \right| \\ &\leq \left| E \left((Y - \widehat{Y}^{(n)}) g(X) \right) E(f_0(X) h(X)) \right| + \left| E(\widehat{Y}^{(n)} g(X)) E \left((Y - \widehat{Y}^{(n)}) h(X) \right) \right| \\ &\leq \| k \|_\infty (2 \| f_0 \|_\infty + r_n) E \left(|Y - \widehat{Y}^{(n)}| \right) \leq \sqrt{8\pi\nu} \| k \|_\infty (2 \| f_0 \|_\infty + r_n) e^{-r_n^2/2\nu}. \end{aligned}$$

Combining these yields

$$(3.19) \quad \left\| \widetilde{\mathfrak{C}}_c^{\otimes, S} - \widehat{\mathfrak{C}}_c^{\otimes, S_n} \right\|_{\text{op}} \leq \left(\sqrt{8\pi\nu} (2 \| f_0 \|_\infty + r_n) + 4\nu \right) \| k \|_\infty e^{-r_n^2/2\nu}.$$

In particular, if we use $\nu = \sigma^2$,

$$(3.20) \quad r_1 = (1 \vee 2\sigma^2) \vee 2\sigma \log^{1/2} \left(\frac{\|k\|_\infty (12\sqrt{\pi}(\|f_0\|_\infty + 1) + 8\sigma)}{\sigma \bar{\lambda}_{\star, \tilde{S}}} \right),$$

where $\bar{\lambda}_{\star, \tilde{S}}$ is the smallest non-zero eigenvalue of $\tilde{\mathfrak{C}}^{\tilde{S}}$, and let $\{r_n\}_{n \geq 1}$ be a non-decreasing sequence then for all $n \geq 1$,

$$\widehat{\lambda}_\star^{(n)} \geq \bar{\lambda}_{\star, \tilde{S}}/2.$$

This follows from the argument on the last page and because this choice guarantees that the right hand side of Eq. (3.19) is upper bounded by $\sigma^2 \bar{\lambda}_{\star, \tilde{S}}$: first notice that $r_1^2/4\sigma^2 \geq \log r_1$ for any $r_1 \geq 1 \vee 2\sigma^2$. Hence,

$$\begin{aligned} & (\sqrt{32\pi}\sigma\|f_0\|_\infty + 4\sigma^2) \|k\|_\infty e^{-r_1^2/2\sigma^2} + \sqrt{8\pi}\sigma \|k\|_\infty e^{-r_1^2/2\sigma^2 + \log r_1} \\ & \leq (\sqrt{32\pi}\sigma\|f_0\|_\infty + 4\sigma^2) \|k\|_\infty e^{-r_1^2/2\sigma^2} + \sqrt{8\pi}\sigma \|k\|_\infty e^{-r_1^2/4\sigma^2} \\ & \leq (\sqrt{32\pi}\sigma(\|f_0\|_\infty + 1) + 4\sigma^2) \|k\|_\infty e^{-r_1^2/4\sigma^2}. \end{aligned}$$

The same arguments applies to any $r_n > r_1$ and the final display is non-increasing in the r_1 argument. Setting this final display equal to $\sigma^2 \bar{\lambda}_{\star, \tilde{S}}/2$ yields the expression in (3.20).

Also, notice that r_1 depends logarithmically on the unknown terms $\|f_0\|_\infty$ and $\bar{\lambda}_{\star, \tilde{S}}$.

Compression in the case of sub-Gaussian noise. We have now all the ingredients to state a proposition for the *sub-Gaussian noise case under the assumption that we have an upper bound on $\|f_0\|_\infty$, a lower bound on $\bar{\lambda}_{\star, \tilde{S}}$ and some control over the variance term σ^2* . In particular, we know that when r_1 is chosen as in (3.20) that $\widehat{\mathfrak{C}}_c^{\otimes, S_n}$ has eigenvalues that are closely related to the eigenvalues of $\tilde{\mathfrak{C}}_c^{\otimes, S}$ and that we can apply our results to compress $\widehat{\mathfrak{m}}_{y,n}^\otimes$. Furthermore, we know how to control the difference between $\widehat{\mathfrak{m}}_{y,n}^\otimes$ and $\mathfrak{m}_{y,n}^\otimes$. The following proposition ties these results together.

PROPOSITION 3.3. — *Let $(\mathcal{X}, \mathfrak{T}_\mathcal{X})$ be a Hausdorff space, $(\mathcal{X} \times \mathbb{R}, \mathfrak{T}, \mathcal{A}, P)$ be a topological measure space such that P is a Radon probability measure which has support S , \mathfrak{T} is the product topology corresponding to $\mathfrak{T}_\mathcal{X}$ and the standard topology on \mathbb{R} , and let k be a continuous and bounded kernel function defined on \mathcal{X} such that the corresponding RKHS \mathcal{H} is finite dimensional. Furthermore, let $(X_1, Y_1), \dots, (X_n, Y_n)$ be i.i.d. random variables attaining values in $\mathcal{X} \times \mathbb{R}$, with law P , and of the form $Y_i = f_0(X_i) + \epsilon_i$, where $\epsilon_1, \dots, \epsilon_n$ are centered i.i.d. random variables which are independent of X_1, \dots, X_n and such that ϵ_1 is sub-Gaussian with variance $0 < \sigma^2$, and f_0 is a measurable and bounded function. Let $\bar{\lambda}_{\star, \tilde{S}}$ be the smallest eigenvalue of the covariance operator $\tilde{\mathfrak{C}}^{\tilde{S}}$ that corresponds to the kernel function $k|_{\tilde{S} \times \tilde{S}}$, where \tilde{S} is the closure of $\{x : (x, y) \in S\}$ in \mathcal{X} . Furthermore, let $\bar{\lambda}_\star$ be the smallest eigenvalue of the covariance operator $\tilde{\mathfrak{C}}_c^{\otimes, S}$ corresponding to the kernel function $\rho|_{S \times S}$, $\rho((y_1, x_1), (y_2, x_2)) = y_1 y_2 k(x_1, x_2)$ for all $x_1, x_2 \in \mathcal{X}, y_1, y_2 \in \mathbb{R}$.*

Given $q \in (0, 1)$ define the sequence $\{r_n\}_{n \geq 1}$ in the following way. Define r_1 as in (3.20) and for $n \geq 2$ through

$$r_n = r_1 \vee \sqrt{2}\sigma \log(n/q).$$

Under these conditions, for any $n \geq 1$, the smallest eigenvalue $\widehat{\lambda}_*^{(n)}$ of $\widehat{\mathfrak{C}}_c^{\otimes, S_n}$ fulfills $\widehat{\lambda}_*^{(n)} \geq \bar{\lambda}_*/2 \geq \sigma^2 \bar{\lambda}_{*, \bar{S}}/2 > 0$ and there exists a ball of radius $\widehat{\delta}^{(n)} = \widehat{\lambda}_*^{(n)} / (2(\|f_0\|_\infty + r_n)^{1/2} \|k\|_\infty^{1/2})$ around $\widehat{\mathfrak{m}}_y^\otimes$ within the affine space spanned by $\widehat{C}_\rho = \{\rho((y, x), \cdot) : (x, y) \in S_n\}$ as a subset of \mathcal{H}_ρ , where S_n is the support of the law of $P^{(n)}$ corresponding to $(X, \widehat{Y}^{(n)})$. Whenever n is (strictly) greater than

$$\left(\frac{8(\|f_0\|_\infty + r_n) \|k\|_\infty \left(\sqrt{2 \log(12/q)} + 192(\|f_0\|_\infty + r_n) \|k\|_\infty / \widehat{\lambda}_*^{(n)} \right)^2}{(\widehat{\lambda}_*^{(n)})^2} \right)^2 \vee \left(\frac{16(\|f_0\|_\infty + r_n)^{1/2} \|k\|_\infty^{1/2} + \sqrt{288 \log(4/q)}}{\widehat{\delta}^{(n)}} \right)^2$$

with probability $1 - q$ there exists a ball of radius $\widehat{\delta}^{(n)}/4$ around $\widehat{\mathfrak{m}}_{y,n}^\otimes$ in $\widehat{C}_{\rho,n}^\otimes$ within the affine subspace spanned by \widehat{C}_ρ and

$$\left\| \widehat{\mathfrak{m}}_{y,n}^\otimes - \mathfrak{m}_{y,n}^\otimes \right\|_\rho \leq 6\sigma \|k\|_\infty^{1/2} n^{-3/4}.$$

Proof. — We derived the inequalities concerning the eigenvalues earlier in this section. Furthermore, the bound on $\widehat{\delta}^{(n)}$ follows directly when applying Theorem 2.17 to the random variables $(X, \widehat{Y}^{(n)})$ and the kernel function $\rho|_{S_n \times S_n}$, noting that $\|\rho|_{S_n \times S_n}\|_\infty \leq (\|f_0\|_\infty + r_n) \|k\|_\infty$. The bound on n is also taken from Theorem 2.17 with the only modification being that a union bound is used to guarantee simultaneously the existence of the ball around $\widehat{\mathfrak{m}}_{y,n}^\otimes$ within $\widehat{C}_{\rho,n}^\otimes$ and that $\|\widehat{\mathfrak{m}}_{y,n}^\otimes - \mathfrak{m}_{y,n}^\otimes\|_\rho$ is upper bounded by $n^{-1/2}$. In detail, for the stated q with probability $1 - q/2$ there exists a ball around $\widehat{\mathfrak{m}}_{y,n}^\otimes$ and with the given choice of r_n , with probability $1 - q/2$

$$\begin{aligned} \left\| \widehat{\mathfrak{m}}_{y,n}^\otimes - \mathfrak{m}_{y,n}^\otimes \right\|_\rho &\leq \sqrt{2\sigma^2} \|k\|_\infty^{1/2} e^{-r_n^2/4\sigma^2} \left(2/n + \pi e^{-r_n^2/2\sigma^2} \right)^{1/2} (q/2)^{-1/2} \\ &\leq 2\sigma \|k\|_\infty^{1/2} n^{-1/2} n^{-1/4} \left(2n^{-1/2} + \pi \right)^{1/2} \leq 3\sigma \|k\|_\infty^{1/2} \pi^{1/2} n^{-3/4}. \end{aligned}$$

follows from (3.16). □

3.4. Simultaneous compression

In this section we are interested in compressing different quantities like the covariance operator and the mean element simultaneously, meaning that we want to find a *single* convex combination of a subset of the data that allows us to approximate both quantities well. As mentioned in the introduction, we are utilizing a direct sum approach to approach the simultaneous compression problem. In this section, we start our exploration with \mathfrak{C}_n and $\mathfrak{m}_{y,n}$ for bounded Y , which is in some sense easy to deal with since the RKHSs corresponding to them have intersection $\{0\}$ (after some

minor adjustments of the kernel functions) which makes the direct sum approach easy to apply. We then explore how we can deal with RKHSs $\mathcal{H}_1, \mathcal{H}_2$ for which the intersection is a non-trivial subspace. This problem is more challenging and we combine the direct sum approach with a quotient space approach to deal with it. We conclude this section by applying this approach to approximate simultaneously $\mathfrak{C}_n, \mathfrak{m}_{y,n}$ and $\sum_{i=1}^n Y_i$, which allows us to calculate the least squares error for RKHS functions using only a core set of the data.

3.4.1. Compressing the covariance and weighted mean embedding simultaneously

One of the main challenges when trying to control the approximation error of \mathfrak{C}_n and $\mathfrak{m}_{y,n}$ simultaneously is to determine the size of the convex set that contains $(\mathfrak{C}_n, \mathfrak{m}_{y,n})$ within the direct sum of two RKHSs and to locate $(\mathfrak{C}_n, \mathfrak{m}_{y,n})$ within this convex set, or, alternatively, to analyze the covariance operator corresponding to this new space. These problems would be easier to handle if we could identify the direct sum space with an RKHS and apply the techniques that we have developed for RKHSs. When using the first approach for $\mathfrak{m}_{y,n}$, we face directly a problem in that we will gain some weighted sum of $(\kappa(X_i, \cdot), k(X_i, \cdot))$ as an approximation, but we need a weighted sum of $(\kappa(X_i, \cdot), Y_i k(X_i, \cdot))$. This problem can be circumvented by incorporating the Y_i 's into the kernel as we have done in the second approach, i.e. for a given kernel function k on \mathcal{X} let,

$$(3.21) \quad \begin{aligned} \rho((y_1, x_1), (y_2, x_2)) &= y_1 y_2 k(x_1, x_2) \\ &= \langle \langle y_1, \cdot \rangle_{\mathbb{R}} \otimes k(x_1, \cdot), \langle y_2, \cdot \rangle_{\mathbb{R}} \otimes k(x_2, \cdot) \rangle_{\otimes} \end{aligned}$$

then ρ is a kernel function on $\mathbb{R} \times \mathcal{X}$ and we move from $\mathfrak{m}_{y,n}$ to $\mathfrak{m}_{y,n}^{\otimes}$. It helps to also extend κ to $\mathbb{R} \times \mathcal{X}$ by setting $\kappa_y((y_1, x_1), (y_2, x_2)) = \kappa(x_1, x_2)$. Let \widehat{h} be the extension of $h \in \mathcal{H}_{\kappa}$ to $\mathbb{R} \times \mathcal{X}$, i.e. $\widehat{h}(y, x) = h(x)$ for all $x \in \mathcal{X}, y \in \mathbb{R}$, then $\|\widehat{h}\|_{\kappa_y} = \|h\|_{\kappa}$. For finite linear combinations this follows from $\|\sum_{i=1}^n \alpha_i \kappa_y((y_i, x_i), \cdot)\|_{\kappa_y}^2 = \sum_{i,j=1}^n \alpha_i \alpha_j \kappa_y((y_i, x_i), (y_j, x_j)) = \|\sum_{i=1}^n \alpha_i \kappa(x_i, \cdot)\|_{\kappa}^2$, where $n \in \mathbb{N}, \alpha_i \in \mathbb{R}, x_i \in \mathcal{X}, y_i \in \mathbb{R}$ for all $i \leq n$ and extends to all of \mathcal{H}_{κ} by a denseness argument. By a similar argument we can see that the extension map is surjective.

Observe that $\widehat{\mathcal{H} \odot \mathcal{H}}$, that is the RKHS corresponding to κ_y , and $\mathcal{H}_{\rho} = \mathbb{R}' \otimes \mathcal{H}$ are linearly independent, i.e. $(\widehat{\mathcal{H} \odot \mathcal{H}}) \cap (\mathbb{R}' \otimes \mathcal{H}) = \{0\}$, because $\kappa_y((y_1, x_1), (y_2, x_2))$ does not depend on the values y_1, y_2 while ρ does. Due to this linear independence we have that $\mathcal{K} = (\widehat{\mathcal{H} \odot \mathcal{H}}) \oplus (\mathbb{R}' \otimes \mathcal{H})$ is isometrically isomorphic to $\mathcal{H}_{\kappa_y + \rho}$. Let $\mathcal{G} = \{g + h : (g, h) \in \mathcal{K}\}$ with norm $\|f\|_{\mathcal{G}} = \inf\{\|(g, h)\|_{\mathcal{K}} : g + h = f, (g, h) \in \mathcal{K}\}$. There exists a surjective isometry between \mathcal{K} and \mathcal{G} . Because $\widehat{\mathcal{H} \odot \mathcal{H}}$ and $\mathbb{R}' \otimes \mathcal{H}$ are linearly independent there is for every $f \in \mathcal{G}$ exactly one pair $(g, h) \in \mathcal{K}$ such that $g + h = f$ and $\|f\|_{\mathcal{G}} = \|(g, h)\|_{\mathcal{K}}$. Furthermore, we have an inner product on \mathcal{G} which is given by $\langle f_1, f_2 \rangle_{\mathcal{G}} = \langle (g_1, h_1), (g_2, h_2) \rangle_{\mathcal{K}}$ whenever $f_1 = g_1 + h_1$ and $f_2 = g_2 + h_2$. For $(g, h) \in \mathcal{K}$ we have that $g \in \mathcal{H}_{\kappa_y}$ and $h \in \mathcal{H}_{\rho}$. By [Aro50, Theorem, p. 353] the kernel $\kappa_y + \rho$ is the kernel of \mathcal{G} and, therefore, $\mathcal{H}_{\kappa_y + \rho}$ is isometrically isomorphic to \mathcal{K} .

When P is a Radon measure with support $S \subset \mathcal{X} \times \mathbb{R}$ then we can look at $\mathcal{K}_S = \widehat{(\mathcal{H} \odot \mathcal{H})}_S \oplus (\mathbb{R}' \otimes \mathcal{H})_S$, where $\widehat{(\mathcal{H} \odot \mathcal{H})}_S = \{u \upharpoonright S_f : u \in \widehat{\mathcal{H} \odot \mathcal{H}}\} = \mathcal{H}_{\kappa_y \upharpoonright S_f \times S_f}$ with norm $\|u\|_{\kappa_y \upharpoonright S_f \times S_f} = \inf\{\|v\| : u = v \upharpoonright S_f, v \in \widehat{\mathcal{H} \odot \mathcal{H}}\}$, with $S_f = \{(y, x) : (x, y) \in S\}$, and similarly we define $(\mathbb{R}' \otimes \mathcal{H})_S$. If $\mathbb{R}_S = \{y : (x, y) \in S\}$ contains at least two elements then $\widehat{(\mathcal{H} \odot \mathcal{H})}_S \cap (\mathbb{R}' \otimes \mathcal{H})_S = \{0\}$ and the above argument shows that \mathcal{K}_S is isometrically isomorphic to $\mathcal{H}_{(\kappa_y + \rho) \upharpoonright S_f \times S_f}$. We summarize this in the following lemma.

LEMMA 3.4. — *Let \mathcal{X} be a measurable space and k a measurable kernel function on \mathcal{X} with corresponding RKHS \mathcal{H} then*

$$\widehat{\mathcal{H} \odot \mathcal{H}} \oplus (\mathbb{R}' \otimes \mathcal{H}) \cong \mathcal{H}_{\kappa_y + \rho}.$$

Furthermore, if P is a Radon measure on $\mathcal{X} \times \mathbb{R}$ with support S and \mathbb{R}_S contains at least two elements then

$$\widehat{(\mathcal{H} \odot \mathcal{H})}_S \oplus (\mathbb{R}' \otimes \mathcal{H})_S \cong \mathcal{H}_{(\kappa_y + \rho) \upharpoonright S_f \times S_f}.$$

In the following, we focus on the case where P is a Radon measure and study the RKHS $\mathcal{H}_{(\kappa_y + \rho) \upharpoonright S_f \times S_f}$. For ease of notation let $\kappa_{\oplus} = (\kappa_y + \rho) \upharpoonright S_f \times S_f$. Similar to before, there is a natural definition for the convex set that contains our mean element. This convex set is

$$C_{\kappa_{\oplus}} = \text{cch}\{\kappa_{\oplus}((y, x), \cdot) : (x, y) \in S\}$$

and the empirical analogue is

$$C_{\kappa_{\oplus}, n} = \text{cch}\{\tilde{\kappa}_{\oplus}((Y_i, X_i), \cdot) : i \leq n\}.$$

The mean element that we want to approximate is then $\mathbf{m}_{\kappa_{\oplus}} = \int \kappa_{\oplus}((y, x), \cdot) dP(x, y)$ when this is well defined, and the empirical analogue is $\mathbf{m}_{\kappa_{\oplus}, n} = (1/n) \sum_{i=1}^n \kappa_{\oplus}((Y_i, X_i), \cdot)$.

In the following, we will assume that $Y = f_0(X) + \epsilon$ with both f_0 and ϵ being bounded and ϵ independent of X .

Covariance operator. We denote the covariance operator corresponding to κ_{\oplus} by $\tilde{\mathfrak{C}}_{\kappa_{\oplus}}$. Because we are dealing with a direct sum one might suppose that it follows directly that the covariance operator factors into the individual covariance operators corresponding to $\kappa_y \upharpoonright S_f \times S_f$ and $\rho \upharpoonright S_f \times S_f$. Unfortunately that is not the case: for $h_1, h_2 \in \mathcal{H}_{\kappa_{\oplus}}$ there exists $f_1, f_2 \in \mathcal{H}_{\kappa_y \upharpoonright S_f \times S_f}$ and $g_1, g_2 \in \mathcal{H}_{\rho \upharpoonright S_f \times S_f}$ such that $h_i = f_i + g_i$ for $i \in \{1, 2\}$, and $\|h_i\|_{\kappa_{\oplus}}^2 = \|f_i\|_{\kappa_y \upharpoonright S_f \times S_f}^2 + \|g_i\|_{\rho \upharpoonright S_f \times S_f}^2$. Hence,

$$(3.22) \quad \begin{aligned} \left\langle \tilde{\mathfrak{C}}_{\kappa_{\oplus}} h_1, h_2 \right\rangle_{\kappa_{\oplus}} &= \left\langle \tilde{\mathfrak{C}}_{\kappa_y \upharpoonright S_f \times S_f} f_1, f_2 \right\rangle_{\kappa_y \upharpoonright S_f \times S_f} \\ &\quad + \left\langle \tilde{\mathfrak{C}}_{\rho \upharpoonright S_f \times S_f} g_1, g_2 \right\rangle_{\rho \upharpoonright S_f \times S_f} + E(f_1 \times g_2) + E(f_2 \times g_1). \end{aligned}$$

The cross-terms do not vanish even if we use the centered covariance operator.

Width of $C_{\kappa_{\oplus}}$. We can apply our standard approach directly to the kernel function κ_{\oplus} (recall that in our definition of this kernel the reduction to the support of P is already incorporated) to gain insights into the convex set $C_{\kappa_{\oplus}}$. Alternatively, we can aim to link the width of $C_{\kappa_{\oplus}}$ back to the width of the corresponding convex sets corresponding to the kernel k and κ . Due to Lemma 3.4 we have that

$$\text{width}_u C_{\oplus} = \text{width}_h C_{\kappa_{\oplus}},$$

where $u \in (\widehat{\mathcal{H} \odot \mathcal{H}})_S \oplus (\mathbb{R}' \otimes \mathcal{H})_S, \|u\| = 1, h$, with $\|h\| = 1$, is the corresponding element in $\mathcal{H}_{\kappa_{\oplus}}$, and

$$C_{\oplus} = \text{cch}\{((\kappa_y \upharpoonright S_f \times S_f)((y, x), \cdot), (\rho \upharpoonright S_f \times S_f)((y, x), \cdot))\} \subset (\widehat{\mathcal{H} \odot \mathcal{H}})_S \oplus (\mathbb{R}' \otimes \mathcal{H})_S.$$

Hence, we can bound the width of C_{\oplus} instead of bounding directly the width of $C_{\kappa_{\oplus}}$. We can write any $u \in (\widehat{\mathcal{H} \odot \mathcal{H}})_S \oplus (\mathbb{R}' \otimes \mathcal{H})_S$ as $(\widehat{g} \upharpoonright S_f, v \upharpoonright S_f)$, where $g \in \mathcal{H} \odot \mathcal{H}$, \widehat{g} is the extension of g to $\mathbb{R} \times \mathcal{X}$, and $v \in \mathbb{R}' \otimes \mathcal{H}$. Observe that if v is given by a finite linear combination of elements $\langle y_i, \cdot \rangle_{\mathbb{R}} \otimes h_i, y_i \in \mathbb{R}, h_i \in \mathcal{H}$, then

$$(3.23) \quad v = \sum_{i=1}^n \alpha_i (\langle y_i, \cdot \rangle_{\mathbb{R}} \otimes h_i) = \langle 1, \cdot \rangle_{\mathbb{R}} \otimes \left(\sum_{i=1}^n y_i \alpha_i h_i \right).$$

For such finite linear combinations let $\psi : \mathbb{R}' \otimes \mathcal{H} \rightarrow \mathcal{H}$ be $\psi(v) = \sum_{i=1}^n y_i \alpha_i h_i$. The map ψ is independent of the particular representation of v because if

$$\sum_{i=1}^n \alpha_i (\langle y_i, \cdot \rangle_{\mathbb{R}} \otimes h_i) = v = \sum_{i=1}^m \beta_i (\langle z_i, \cdot \rangle_{\mathbb{R}} \otimes g_i)$$

for a suitable $m \in \mathbb{N}$ and corresponding $\beta_i, z_i \in \mathbb{R}, g_i \in \mathcal{H}$ for all $i \leq m$, then

$$0 = \left\| \langle 1, \cdot \rangle_{\mathbb{R}} \otimes \left(\sum_{i=1}^n y_i \alpha_i h_i - \sum_{i=1}^m z_i \beta_i g_i \right) \right\|_{\rho}^2 = \left\| \sum_{i=1}^n y_i \alpha_i h_i - \sum_{i=1}^m z_i \beta_i g_i \right\|^2$$

We can also observe that $\|v\|_{\rho}^2 = \|\sum_{i=1}^n y_i \alpha_i h_i\|^2 = \|\psi(v)\|^2$. Furthermore, ψ is linear and therefore an isometry. Since the finite linear combinations lie dense in $\mathbb{R}' \otimes \mathcal{H}$ and \mathcal{H} , and both $\mathbb{R}' \otimes \mathcal{H}$ and \mathcal{H} are complete, we can extend ψ to a surjective isometry between $\mathbb{R}' \otimes \mathcal{H}$ and \mathcal{H} [Eng89, Corollary 4.3.18]. In particular, any $v \in \mathbb{R}' \otimes \mathcal{H}$ can be represented as $\psi^{-1}(h)$ with a unique $h \in \mathcal{H}$.

The width of C_{\oplus} can now be lower bounded in the following way: choose $\alpha > 0$, let $b_{\alpha} = \sup\{b : \Pr(\epsilon \geq b) > \alpha \text{ and } \Pr(\epsilon \leq -b) > \alpha\}$ and $I = [-b_{\alpha}, b_{\alpha}]$. Then

$$\begin{aligned} & \text{width}_{(\widehat{g} \upharpoonright S_f, \psi^{-1}(h) \upharpoonright S_f)} C_{\oplus} \\ &= \sup_{x \in \mathcal{X}_S, z \in I} (g(x) + (f_0(x) + z)h(x)) - \inf_{x \in \mathcal{X}_S, z \in I} (g(x) + (f_0(x) + z)h(x)) \end{aligned}$$

whenever $(\widehat{g} \upharpoonright S_f, \psi^{-1}(h) \upharpoonright S_f)$ has norm one, $g \in \mathcal{H} \odot \mathcal{H}$ and $h \in \mathcal{H}$. In particular, when choosing the same point x and using z to move to absolute values, we gain

$$(3.24) \quad \text{width}_{(\widehat{g} \upharpoonright S_f, \psi^{-1}(h) \upharpoonright S_f)} C_{\oplus} \geq 2\|h\|b_{\alpha} \sup_{x \in \mathcal{X}_S} \frac{|h(x)|}{\|h\|} \geq b_{\alpha}\|h\| \text{width}_{h/\|h\|}(C_{\mathcal{X}_S}),$$

where $C_{\mathcal{X}_S}$ is the usual convex set for the kernel $k \upharpoonright \mathcal{X}_S \times \mathcal{X}_S$.

We need to complement this bound with a bound that is based on g to deal with cases where $\|h\|$ is small. When $\|f_0\|_\infty$ is smaller than b_α then there is a simple way to get a lower bound that involves g . For two points $x_1, x_2 \in \mathcal{X}_S$ and any $h \in \mathcal{H}$, we can choose $z_1, z_2 \in \mathbb{R}_S$ such that $(f_0(x_1) + z_1)h(x_1) \geq 0 \geq (f_0(x_2) + z_2)h(x_2)$ and, hence,

$$(3.25) \quad \text{width}_{(\hat{g}|_{S_f}, \psi^{-1}(h)|_{S_f})} C_\oplus \geq \sup_{x \in \mathcal{X}_S} g(x) - \inf_{x \in \mathcal{X}_S} g(x) \geq \|g\|_{\mathcal{H} \odot \mathcal{H}} \text{width}_{g/\|g\|_{\mathcal{H} \odot \mathcal{H}}} C_\odot.$$

We can combine (3.24) and (3.25) to gain a lower bound on the width of C_\oplus in terms of the widths of C and C_\odot .

THE LOW NOISE SETTING. The situation gets more complicated when $|f_0|$ attains values that are significantly larger than b_α . For instance, when there is no noise, i.e. $\epsilon = 0$ (a.s.), and there exists some $h \in \mathcal{H}, g \in \mathcal{H} \odot \mathcal{H}$ such that $h(x) \neq 0$ and $f_0(x) = -g(x)/h(x)$ on \mathbb{R}_S , and $(\hat{g}|_{S_f}, \psi^{-1}(h)|_{S_f})$ has norm one, it holds that $\text{width}_{(\hat{g}|_{S_f}, \psi^{-1}(h)|_{S_f})} C_\oplus = 0$. For f_0 to be equal or close to $-g/h$ it is necessary that f_0 attains large values when $\|h\|$ is small. For example, when $\mathcal{H} \odot \mathcal{H}$ is finite dimensional with dimension d , $\lambda_d > 0$ is the smallest eigenvalue of a suitable kernel matrix based on the kernel κ , and k is a bounded kernel, then $f_0(x) = -g(x)/h(x)$ can only happen if

$$\sup_{x \in \mathcal{X}} |f_0(x)| \geq \frac{\|g\|_\infty}{\|h\|_\infty} \geq \frac{\|g\| \lambda_d^{1/2}}{d^{1/2} \|h\| \|k\|_\infty^{1/2}}.$$

For a small value of $\|h\|$ this implies that $\|g\|$ will be close to 1 and $|f_0|$ has to attain a large value at some locations $x \in \mathcal{X}$.

INTERPOLATION AND ANOTHER LOOK AT THE LOW NOISE SETTING. We look now at the case where there is no noise at all, that is $Y = f_0(X)$,

$$C_\oplus = \text{cch}\{(\kappa_y((f_0(x), x), \cdot), \langle f_0(x), \cdot \rangle_{\mathbb{R}^r} \otimes k(x, \cdot)) : x \in \mathcal{X})\}$$

and we are interested in interpolating f_0 . In particular, we are controlling the width of C_\oplus depending on how f_0 is related to \mathcal{H} and $\mathcal{H} \odot \mathcal{H}$. The direct sum approach is useful to gain a deeper understanding of how well $(\mathfrak{C}_y, \mathfrak{m}_y^\otimes)$ can be approximated. The width of C_\oplus in this interpolation setting has a simple form. Assume that the support of the marginal measure is all of \mathcal{X} and since there is no noise it then follows that the support of the measure P is $S = \{(x, f_0(x)) : x \in \mathcal{X}\}$. For $g \in \mathcal{H} \odot \mathcal{H}, h \in \mathcal{H}$,

$$\text{width}_{(\hat{g}, \psi^{-1}(h))} C_\oplus = \sup_{x \in \mathcal{X}} (g(x) + f_0(x)h(x)) - \inf_{x \in \mathcal{X}} (g(x) + f_0(x)h(x))$$

The functions $f_0 \times h$ lie in the RKHS \mathcal{H}_{f_0} which has the kernel function $k_0(x, y) = f_0(x)k(x, y)f_0(y)$. According to [PR16, Proposition 5.20] the RKHS \mathcal{H}_{f_0} is equal to $\{f_0 \times h : h \in \mathcal{H}\}$ and the inner product on \mathcal{H}_{f_0} is given by $\langle f_0 \times h_1, f_0 \times h_2 \rangle_{\mathcal{H}_{f_0}} = \langle h_1, h_2 \rangle$ whenever $h_1, h_2 \in \mathcal{H}$. If $\mathcal{H}_{f_0} \cap (\mathcal{H} \odot \mathcal{H}) = \{0\}$ then we can embed both \mathcal{H}_{f_0} and $\mathcal{H} \odot \mathcal{H}$ in the direct sum $\mathcal{G} = (\mathcal{H} \odot \mathcal{H}) \oplus \mathcal{H}_{f_0}$ such that for any $f \in \mathcal{H}_{f_0}, h \in \mathcal{H} \odot \mathcal{H}$ it holds that $\|f\|_{\mathcal{H}_{f_0}} = \|(0, f)\|_{\mathcal{G}}$ and $\|h\|_{\mathcal{H} \odot \mathcal{H}} = \|(h, 0)\|_{\mathcal{G}}$. As in Lemma 3.4 it holds that $\mathcal{G} \cong \mathcal{H}_{k_0 + \kappa}$ and, therefore, it also holds that $\|f\|_{k_0} = \|f\|_{k_0 + \kappa}$ and $\|h\|_{\mathcal{H} \odot \mathcal{H}} = \|h\|_{k_0 + \kappa}$. In this case,

$$\text{width}_{(\hat{g}, \psi^{-1}(h))} = \text{width}_{(g, f_0 \times h)} C_{\mathcal{G}}$$

whenever $g \in \mathcal{H} \odot \mathcal{H}, h \in \mathcal{H}$, where $C_{\mathcal{G}} = \text{cch}\{(\kappa(x, \cdot), k_0(x, \cdot)) : x \in \mathcal{X}\} \subset \mathcal{G}$. This follows directly from

$$g(x) + f_0(x)h(x) = \langle (g, f_0 \times h), (\kappa(x, \cdot), k_0(x, \cdot)) \rangle_{\mathcal{G}}.$$

We can now follow the approach from Section 2.3.5 and, in particular, apply Proposition 2.10 to the RKHS with kernel $k_{f_0} + \kappa$. Assumptions on f_0 imply then lower bounds on the width.

3.4.2. Linearly dependent spaces

The setting above where we approximate \mathfrak{C}_n and $\mathfrak{m}_{y,n}^{\otimes}$ simultaneously is easy to deal with because the corresponding RKHSs are linearly independent. On the other hand, when the spaces over which we want to optimize are linearly dependent then the RKHS is not isometrically isomorphic to the direct product space and the approach needs to be modified. This can happen, for example, when we try to approximate \mathfrak{C} simultaneously to \mathfrak{m} . In this context the corresponding spaces $\mathcal{H} \odot \mathcal{H}$ and \mathcal{H} can overlap. For instance, when k is a polynomial kernel of order two then \mathcal{H} and $\mathcal{H} \odot \mathcal{H}$ are not linearly independent.

Whenever $\mathcal{H} \odot \mathcal{H}$ and \mathcal{H} are not linearly independent it is natural to identify elements like $(h, 0)$ and $(0, h)$, $h \in (\mathcal{H} \odot \mathcal{H}) \cap \mathcal{H}$. One way to do so is to consider the subspace $U = \{(-h, h) : h \in (\mathcal{H} \odot \mathcal{H}) \cap \mathcal{H}\}$ of $\mathcal{K} := (\mathcal{H} \odot \mathcal{H}) \oplus \mathcal{H}$. The subspace is closed: let $\{(-h_n, h_n)\}_{n \in \mathbb{N}}$ be a convergent sequence in U . This sequence is also a Cauchy sequence and for any $\epsilon > 0$ there exists an $N \in \mathbb{N}$ such that for all $n, m \geq N$,

$$\epsilon > \|(-h_n, h_n) - (-h_m, h_m)\|_{\oplus}^2 = \|h_m - h_n\|^2 + \|h_n - h_m\|_{\mathcal{H} \odot \mathcal{H}}^2$$

and $\{h_n\}_{n \in \mathbb{N}}$ is a Cauchy sequence both in \mathcal{H} and $\mathcal{H} \odot \mathcal{H}$. Hence, it converges in both spaces. Let f be its limit in $\mathcal{H} \odot \mathcal{H}$ and g its limit in \mathcal{H} then for any $x \in \mathcal{X}$ there exists an $n \in \mathbb{N}$ such that $|f(x) - g(x)| \leq \epsilon + |h_n(x) - h_n(x)| = \epsilon$ and $f = g$. It also follows right away that $\lim_{n \rightarrow \infty} \|(-h_n, h_n) - (-f, g)\|_{\oplus} = 0$ and the sequence has its limit in U .

Consider the quotient space \mathcal{K}/U with co-sets $f^{\bullet} = f + U$, $f \in \mathcal{K}$, and the quotient norm $\|f^{\bullet}\|_{\mathcal{K}/U} = \inf\{\|f + h\|_{\mathcal{K}} : h \in U\}$. The space \mathcal{K}/U is again a Hilbert space since U is closed (e.g. [RS72, Section III.4]), and it is isometrically isomorphic to the Hilbert space $\mathcal{H} \odot \mathcal{H} + \mathcal{H}$ when the latter is equipped with the norm $\|f\|_{+}^2 = \inf\{\|g\|_{\mathcal{H} \odot \mathcal{H}}^2 + \|h\|^2 : f = g + h, g \in \mathcal{H} \odot \mathcal{H}, h \in \mathcal{H}\}$; in particular, a co-set $(g, h) + U \in \mathcal{K}/U$ is mapped to the function $f = g + h$. This map is well defined since if $(g_1, h_1) \in (g, h)^{\bullet}$ then there is some h_2 such that $g_1 + h_1 = g - h_2 + h + h_2 = f$. Furthermore, by the choice of U , there are no two elements $u^{\bullet}, v^{\bullet} \in \mathcal{K}/U$, $u^{\bullet} \neq v^{\bullet}$, that are mapped to the same function f . Assume otherwise, then there is some f such that $f = g_1 + h_1 = g_2 + h_2$ and, therefore, $(g_2 + g_1 - g_2, h_2 - g_1 + g_2) = (g_1, h_1)$. Since $g_1 - g_2 \in \mathcal{H} \odot \mathcal{H}$ and $g_1 - g_2 = h_1 - h_2 \in \mathcal{H}$ it follows that $(g_2, h_2)^{\bullet} = (g_1, h_1)^{\bullet}$ which contradicts the assumption. Finally, any element in $\mathcal{H} \odot \mathcal{H} + \mathcal{H}$ can be represented this way since if $f = g + h$, $g \in \mathcal{H} \odot \mathcal{H}, h \in \mathcal{H}$ then $(g, h)^{\bullet}$ is mapped to f . Using again [Aro50, Theorem p. 353] we can conclude that \mathcal{K}/U and $\mathcal{H}_{\kappa+k}$ are isometrically isomorphic.

While \mathcal{K}/U and $\mathcal{H}_{\kappa+k}$ are isometrically isomorphic it does not hold in general that \mathcal{K} and \mathcal{K}/U are isometrically isomorphic to $\mathcal{H}_{\kappa+k}$. Hence, when mapping an element $u \in \mathcal{K}$ to $u^\bullet \in \mathcal{K}/U$, then finding an approximation v^\bullet of u^\bullet in \mathcal{K}/U , we generally cannot invert the \bullet operation to gain an approximation of u . Selecting an arbitrary element in v^\bullet does not work either since a small value of $\|u^\bullet - v^\bullet\|_{\mathcal{K}/U}$ does not imply that all elements in the corresponding co-sets have small distances, i.e. there is no reason why $\sup_{w \in v^\bullet} \|u - w\|_{\mathcal{K}}$ should be small. However, we are not trying to approximate arbitrary elements in \mathcal{K} but only elements

$$(\mathfrak{C}, \mathfrak{m}) = \frac{1}{n} \sum_{i=1}^n (\kappa(X_i, \cdot), k(X_i, \cdot))$$

and we are optimizing the approximation over $\tilde{C} = \text{cch}\{(\kappa(x, \cdot), k(x, \cdot)) : x \in \mathcal{X}\} \subset \mathcal{K}$. The important observation is that for any non-zero element $(-h, h) \in U$, that is $h \in (\mathcal{H} \odot \mathcal{H}) \cap \mathcal{H}$, we have

$$\langle (-h, h), (\kappa(x, \cdot), k(x, \cdot)) \rangle_{\mathcal{K}} = -h(x) + h(x) = 0,$$

and \tilde{C} is a subset of U^\perp .

The subspace U^\perp together with the inner product inherited from \mathcal{K} is isometrically isomorphic to $\mathcal{H}_{\kappa+k}$. This follows since \mathcal{K}/U and $\mathcal{H}_{\kappa+k}$ are isometrically isomorphic and U^\perp and \mathcal{K}/U are isometrically isomorphic. The latter holds since every co-set corresponds to exactly one element in U^\perp , and for $u \in U^\perp$, $\|u^\bullet\|_{\mathcal{K}/U} = \inf\{\|u + v\|_{\mathcal{K}} : v \in U\} = \|u\|_{\mathcal{K}}$.

Also, $\overline{\text{span}} \tilde{C} = U^\perp$. We know already that $\overline{\text{span}} \tilde{C} \subset U^\perp$. To show that they are equal let $\mathcal{K} = \text{span}((\overline{\text{span}} \tilde{C}) \cup U)$. Observe that this space is closed since $\overline{\text{span}} \tilde{C}$ and U are, and because they are orthogonal. It is sufficient to show that $(f, 0) \in \mathcal{K}, (0, g) \in \mathcal{K}$ for all $f \in \mathcal{H} \odot \mathcal{H}$ and $g \in \mathcal{H}$ since the smallest closed subspace that contains all these elements is $(\mathcal{H} \odot \mathcal{H}) \oplus \mathcal{H}$.

For $f = \sum_{i=1}^n \beta_i (\kappa(x_i, \cdot) + k(x_i, \cdot)) \in \mathcal{H}_{\kappa+k}$ define $\psi(f) = \sum_{i=1}^n \beta_i (\kappa(x_i, \cdot), k(x_i, \cdot)) \in \overline{\text{span}} \tilde{C} \subset (\mathcal{H} \odot \mathcal{H}) \oplus \mathcal{H}$. The operator $\psi : \mathcal{H}_{\kappa+k} \rightarrow \overline{\text{span}} \tilde{C}$ is linear and defined on a dense subset of $\mathcal{H}_{\kappa+k}$. It is furthermore norm preserving since

$$\|\psi(f)\|_{\oplus}^2 = \sum_{i,j=1}^n \beta_i \beta_j \kappa(x_i, x_j) + \sum_{i,j=1}^n \beta_i \beta_j k(x_i, x_j) = \|f\|_{\kappa+k}^2.$$

Hence, it can be extended to a linear isometry, which we will also denote by ψ , between $\mathcal{H}_{\kappa+k}$ and $\overline{\text{span}} \tilde{C}$ with the norm inherited from $(\mathcal{H} \odot \mathcal{H}) \oplus \mathcal{H}$.

For any $h \in (\mathcal{H} \odot \mathcal{H}) \cap \mathcal{H}$ we can infer that it lies in the RKHS with kernel $\kappa + k$ due to [Aro50, Theorem, p. 353] and $\psi(h)$ lies in $\overline{\text{span}} \tilde{C}$. Write $\psi(h)$ as (h_1, h_2) , $h_1 \in \mathcal{H} \odot \mathcal{H}, h_2 \in \mathcal{H}$, then for all $x \in \mathcal{X}$,

$$\begin{aligned} h_1(x) + h_2(x) &= \langle (h_1, h_2), (\kappa(x, \cdot), k(x, \cdot)) \rangle_{\oplus} = \langle \psi(h), \psi(\kappa(x, \cdot) + k(x, \cdot)) \rangle_{\oplus} \\ &= \langle h, \kappa(x, \cdot) + k(x, \cdot) \rangle_{\kappa+k} = h(x). \end{aligned}$$

In other words, for any $h \in (\mathcal{H} \odot \mathcal{H}) \cap \mathcal{H}$ we have $h_1 \in \mathcal{H} \odot \mathcal{H}, h_2 \in \mathcal{H}$ such that $h = h_1 + h_2$ and $(h_1, h_2) \in \overline{\text{span}} \tilde{C}$. Since $h, h_1 \in \mathcal{H} \odot \mathcal{H}$ it follows that

$h_2 = h - h_1 \in (\mathcal{H} \odot \mathcal{H}) \cap \mathcal{H}$ and $(h_2, -h_2) \in U$. Thus, $(h, 0) = (h_1, h_2) + (h_2, -h_2) \in \mathcal{K}$. Similarly, we can observe that $(0, h) \in \mathcal{K}$.

For $f \in \mathcal{H} \odot \mathcal{H}$ let $\psi(f) = (f_1, f_2)$ with $f_1 \in \mathcal{H} \odot \mathcal{H}$ and $f_2 \in \mathcal{H}$. In other words, $f = f_1 + f_2$ and since $\mathcal{H} \odot \mathcal{H}$ is a linear space, we know that $f_2 = f - f_1 \in (\mathcal{H} \odot \mathcal{H}) \cap \mathcal{H}$. And, as above, we can conclude that $(f, 0)$ also lies in \mathcal{K} . The same argument also shows that for any $g \in \mathcal{H}$ we have $(0, g) \in \mathcal{K}$. Hence, $\mathcal{K} = (\mathcal{H} \odot \mathcal{H}) \oplus \mathcal{H}$ and $\overline{\text{span}} \tilde{C} = U^\perp$.

LEMMA 3.5. — *Let \mathcal{X} be a measurable space and k a measurable kernel function on \mathcal{X} with corresponding RKHS \mathcal{H} then $\overline{\text{span}}\{(\kappa(x, \cdot), k(x, \cdot)) : x \in \mathcal{X}\} \subset (\mathcal{H} \odot \mathcal{H}) \oplus \mathcal{H}$ equipped with the inner product of $(\mathcal{H} \odot \mathcal{H}) \oplus \mathcal{H}$ is isometrically isomorphic to $\mathcal{H}_{\kappa+k}$.*

3.4.3. Simultaneous least-squares risk approximation for unbounded Y

Often it is unnecessary to include the $(1/n) \sum_{i=1}^n Y_i^2$ term in the simultaneous approximation since many methods only rely on the terms that include f (e.g. the ridge regressor) and also $(1/n) \sum_{i=1}^n Y_i^2 \in \mathbb{R}$ itself can be represented by a single real number and does not need to be compressed. However, when selecting points $(X_{\iota(1)}, Y_{\iota(m)})$, $m \ll n$, for a coresnet then

$$\frac{1}{m} \sum_{i=1}^m (f^2(X_{\iota(i)}) - 2Y_{\iota(i)}f(X_{\iota(i)})) + \frac{1}{n} \sum_{i=1}^n Y_i^2$$

is not the mean squared error of f given the sample $X_{\iota(1)}, Y_{\iota(1)}, \dots, X_{\iota(m)}, Y_{\iota(m)}$ and might even be negative. An easy way to remedy this problem is to move to $(1/m) \sum_{i=1}^m Y_{\iota(i)}^2$ but then we do not have any guarantee that this is close to $(1/n) \sum_{i=1}^n Y_i^2$. An alternative is to include the Y_i 's in the simultaneous approximation problem. This can be done by, for instance, defining a kernel on $\mathbb{R} \times \mathcal{X}$ through $r((y_1, x_1), (y_2, x_2)) = \langle y_1, y_2 \rangle_{\mathbb{R}}$ and by considering the direct sum

$$\widehat{\mathcal{H} \odot \mathcal{H}} \oplus (\mathbb{R}' \otimes \mathcal{H}) \oplus \mathcal{H}_r.$$

Alternatively, we can restrict the functions to the support S of the underlying measure and consider

$$(\widehat{\mathcal{H} \odot \mathcal{H}})_S \oplus (\mathbb{R}' \otimes \mathcal{H})_S \oplus \mathcal{H}_{r|S_f \times S_f}.$$

If there is *no constant function in the RKHS $\mathcal{H}_{\mathcal{X}_S}$* then $(\mathbb{R}' \otimes \mathcal{H})_S \cap \mathcal{H}_{r|S_f \times S_f} = \{0\}$ and

$$(\widehat{\mathcal{H} \odot \mathcal{H}})_S \cap ((\mathbb{R}' \otimes \mathcal{H})_S \cap \mathcal{H}_{r|S_f \times S_f}) = \{0\}$$

if, furthermore, \mathbb{R}_S contains at least two different values: any function in $\widehat{\mathcal{H} \odot \mathcal{H}}$ is of the form $g^2(x), x \in \mathcal{X}_S, g \in \mathcal{H}$ and functions in $(\mathbb{R}' \otimes \mathcal{H})_S \cap \mathcal{H}_{r|S_f \times S_f}$ are of the form $(y, x) \mapsto yh(x) + cy$ for some constant c . For any functions $g, h \in \mathcal{H}$, choose $y_1, y_2 \in \mathbb{R}_S, y_1 \neq y_2$, and $x_1, x_2 \in \mathcal{X}_S$ are such that $h(x_1) \neq h(x_2)$. For g^2 to be equal to $yh(x) + cy$ it has to hold that $g^2(x_1)$ is equal to $y_1h(x_1) + cy_1$ and it also has to be equal to $y_2h(x_1) + cy_2$. In other words, $(y_1 - y_2)h(x_1) = c(y_2 - y_1)$ and $h(x_1) = -c$ and similarly for $h(x_2)$, that is $h(x_1) = h(x_2)$ with a contradiction to the choice of x_1

and x_2 . Hence, under these conditions we can identify the direct sum with an RKHS corresponding to a sum of kernels,

$$(\widehat{\mathcal{H} \odot \mathcal{H}})_S \oplus (\mathbb{R}' \otimes \mathcal{H})_S \oplus \mathcal{H}_{r|S_f \times S_f} \cong \mathcal{H}_{(\kappa_y + \rho + r)|S_f \times S_f}.$$

When the RKHS $\mathcal{H}_{\mathcal{X}_S}$ contains the constant function then the intersection $(\mathbb{R}' \otimes \mathcal{H})_S \cap \mathcal{H}_{r|S_f \times S_f}$ is not empty since the function $(y, x) \mapsto y$, with domain S , lies in $(\mathbb{R}' \otimes \mathcal{H})_S$ and in $\mathcal{H}_{r|S_f \times S_f}$. We can follow the same approach as in Section 3.4.2 and consider the one dimensional subspace $U = \{(-h, h) : h \in (\mathbb{R}' \otimes \mathcal{H})_S \cap \mathcal{H}_{r|S_f \times S_f}\}$ of $(\mathbb{R}' \otimes \mathcal{H})_S \oplus \mathcal{H}_{r|S_f \times S_f}$ and consider the quotient space $\mathcal{Q} = ((\mathbb{R}' \otimes \mathcal{H})_S \oplus \mathcal{H}_{r|S_f \times S_f})/U$ with the usual quotient norm. The space \mathcal{Q} is a Hilbert space [RS72, Section III.4]. By the same argument as in Section 3.4.2 we can infer that

$$\mathcal{Q} \cong \mathcal{H}_{(\rho+r)|S_f \times S_f}$$

and for $(x, y) \in S$, $(\rho((y, x), \cdot), r((y, x), \cdot))$ lies in $U^\perp \subseteq (\mathbb{R}' \otimes \mathcal{H})_S \oplus \mathcal{H}_{r|S_f \times S_f}$. Also, the space U^\perp , with the inherited inner product, is isometric isomorphic to $\mathcal{H}_{(\rho+r)|S_f \times S_f}$. When \mathbb{R}_S contains at least two elements then $(\widehat{\mathcal{H} \odot \mathcal{H}})_S \cap ((\mathbb{R}' \otimes \mathcal{H})_S \cap \mathcal{H}_{r|S_f \times S_f}) = \{0\}$ and

$$(\widehat{\mathcal{H} \odot \mathcal{H}})_S \oplus \mathcal{Q} \cong \mathcal{H}_{(\kappa_y + \rho + r)|S_f \times S_f}.$$

We will apply these results to the problem of ridge regression and it is convenient to have a proposition which provides guarantees on the approximation in the ridge regression context. Since we do not need to approximate the sum of the Y_i^1 terms to compute the ridge regression estimator we will consider the space $(\widehat{\mathcal{H} \odot \mathcal{H}})_S \cap (\mathbb{R}' \otimes \mathcal{H})_S$. We make the assumption that \mathcal{H} does not contain the constant functions, which removes the need to consider quotient spaces. Furthermore, we will assume *sub-Gaussian noise* and that f_0 is bounded but we will allow for *unbounded Y_i random variables*. Recall the definitions of the kernel functions ρ in (3.21), $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, $\kappa = k^2$, and its extension κ_y (see just below (3.21)). Before stating a result on the compression, we need to modify the arguments that we used to control the difference between $\mathbf{m}_{y,n}^\otimes$ and $\widehat{\mathbf{m}}_{y,n}^\otimes$, and the difference between $\widetilde{\mathfrak{C}}_c^{\otimes, S}$ and $\widehat{\mathfrak{C}}_c^{\otimes, S_n}$. This is necessary since the kernel function, which we will denote by τ below, is $(\kappa_y + \rho)((y, x), (y', x')) = \kappa(x, x') + yy'k(x, x')$ in the current context, and this kernel function is not of the form $yy'\widetilde{k}(x, x')$, where \widetilde{k} is some kernel on \mathcal{X} . Since we assumed that latter form in Section 3.3.3 when we derived the bounds on the differences, we cannot simply reuse the earlier results. Fortunately, the necessary modifications are minor: let $\mathbf{m}_{\tau,n}$ be the empirical mean element corresponding to the kernel function τ on $\mathbb{R} \times \mathcal{X}$ and

$$\widehat{\mathbf{m}}_{\tau,n} = \frac{1}{n} \sum_{i \leq n} \tau(\widehat{Y}_i^{(n)}, X_i, \cdot) = \frac{1}{n} \sum_{i \leq n} (\kappa(X_i, \cdot) + \langle \widehat{Y}_i^{(n)}, \cdot \rangle_{\mathbb{R}} k(X_i, \cdot)).$$

The distance between the empirical mean element and its capped version is

$$\begin{aligned} \|\mathbf{m}_{\tau,n} - \widehat{\mathbf{m}}_{\tau,n}\|_{\tau}^2 &= \frac{1}{n^2} \left\| \sum_{i \leq n} \langle Y_i - \widehat{Y}_i^{(n)}, \cdot \rangle_{\mathbb{R}} k(X_i, \cdot) \right\|_{\tau}^2 \\ &\leq \frac{1}{n^2} \left\| \sum_{i \leq n} \langle Y_i - \widehat{Y}_i^{(n)}, \cdot \rangle_{\mathbb{R}} k(X_i, \cdot) \right\|_{\rho}^2 \end{aligned}$$

since the κ terms cancel and because the function inside the norm lies within \mathcal{H}_{ρ} (apply [PR16, Theorem 5.4] to get the inequality). This implies that we can reuse the bound in (3.16), and with probability $1 - \delta$, $\delta \in (0, 1)$, we have that

$$(3.26) \quad \|\mathbf{m}_{\tau,n} - \widehat{\mathbf{m}}_{\tau,n}\|_{\tau}^2 \leq \sqrt{2\nu} \|k\|_{\infty}^{1/2} e^{-r_n^2/4\nu} \left(2/n + \pi e^{-r_n^2/2\nu}\right)^{1/2} \delta^{-1/2},$$

where $\nu > 0$ is the variance factor corresponding to the sub-Gaussian noise terms. We also need control over the covariance operators corresponding to the capped Y_i 's. We proceed as in Section 3.3.3. Assuming that the law P of (Y, X) is a Radon measure let S be its support and let S_n be the support of $(\widehat{Y}^{(n)}, X)$ (which is well defined as the law of this random variable is again a Radon measure). Let $\widetilde{\mathfrak{C}}_{c,\tau}^S : \mathcal{H}_{\tau} \rightarrow \mathcal{H}_{\tau}$ be the covariance operator corresponding to the original random variable and $\widehat{\mathfrak{C}}_{c,\tau}^{S_n} : \mathcal{H}_{\tau} \rightarrow \mathcal{H}_{\tau}$ the covariance operator corresponding to the capped random variable. We start by bounding the difference between these covariance operators in the operator norm,

$$\left\| \widetilde{\mathfrak{C}}_{c,\tau}^S - \widehat{\mathfrak{C}}_{c,\tau}^{S_n} \right\|_{\text{op}} = \sup_{\|h_1\|_{S_f} = 1} \sup_{\|h_2\|_{S_f} = 1} \left\langle \left(\widetilde{\mathfrak{C}}_{c,\tau}^S - \widehat{\mathfrak{C}}_{c,\tau}^{S_n} \right) h_1 \upharpoonright_{S_f}, h_2 \upharpoonright_{S_f} \right\rangle_{\tau, S_f}.$$

Due to Lemma 3.4 the space \mathcal{H}_{τ, S_f} is isometrically isomorphic to a direct sum space and, as above (3.22), we can write $h_i = f_i + g_i$ for $i \in \{1, 2\}$, $f_1, f_2 \in \mathcal{H}_{\kappa_y \upharpoonright_{S_f \times S_f}}$, $g_1, g_2 \in \mathcal{H}_{\rho \upharpoonright_{S_f \times S_f}}$ and such that the squared norms of the h_i equals the sum of the squared norms of the f_i and g_i . We can proceed by expanding the h_i 's and observing that $f_i(\widehat{y}, x) = f_i(y, x)$ for all $y \in \mathbb{R}$ and $x \in \mathcal{X}$ since f_i is a function of the second coordinate only,

$$\begin{aligned} &\left\langle \left(\widetilde{\mathfrak{C}}_{c,\tau}^S - \widehat{\mathfrak{C}}_{c,\tau}^{S_n} \right) h_1 \upharpoonright_{S_f}, h_2 \upharpoonright_{S_f} \right\rangle_{\tau, S_f} \\ &= E(h_1(Y, X)h_2(Y, X)) - E(h_1(Y, X))E(h_2(Y, X)) \\ &\quad - E(h_1(\widehat{Y}^{(n)}, X)h_2(\widehat{Y}^{(n)}, X)) + E(h_1(\widehat{Y}^{(n)}, X))E(h_2(\widehat{Y}^{(n)}, X)). \end{aligned}$$

The difference of the bias terms becomes

$$\begin{aligned} &E(h_1(\widehat{Y}^{(n)}, X))E(h_2(\widehat{Y}^{(n)}, X)) - E(h_1(Y, X))E(h_2(Y, X)) \\ &= E(g_1(\widehat{Y}^{(n)}, X))E(g_2(\widehat{Y}^{(n)}, X)) - E(g_1(Y, X))E(g_2(Y, X)) \\ &\quad + E(f_1(Y, X))(E(g_2(\widehat{Y}^{(n)}, X)) - E(g_2(Y, X))) \\ &\quad + E(f_2(Y, X))(E(g_1(\widehat{Y}^{(n)}, X)) - E(g_1(Y, X))). \end{aligned}$$

Since f_1, f_2 have norm one it follows that $E(f_1(Y, X))$ and $E(f_2(Y, X))$ are upper bounded by $\|\kappa\|_\infty^{1/2} = \|k\|_\infty$. Also recall that $g_i, i \in \{1, 2\}$, can be written as \check{u}_i for some $u_i \in \mathcal{H}_k$ and u_i has norm one (see (3.13)). Hence

$$\left| E\left(g_i(\widehat{Y}^{(n)}, X)\right) - E\left(g_i(Y, X)\right) \right| \leq \|k\|_\infty^{1/2} E\left(|\widehat{Y}^{(n)} - Y|\right)$$

and the bound (3.14) can be used. Similarly,

$$\begin{aligned} & \left| E\left(g_1(\widehat{Y}^{(n)}, X)\right) E\left(g_2(\widehat{Y}^{(n)}, X)\right) - E\left(g_1(Y, X)\right) E\left(g_2(Y, X)\right) \right| \\ & \leq \left| E\left(g_1(\widehat{Y}^{(n)}, X)\right) \right| \left| E\left(g_2(\widehat{Y}^{(n)}, X)\right) - E\left(g_2(Y, X)\right) \right| \\ & \quad + \left| E\left(g_2(Y, X)\right) \right| \left| E\left(g_1(\widehat{Y}^{(n)}, X)\right) - E\left(g_1(Y, X)\right) \right| \\ & \leq 2\|k\|_\infty E\left(|\widehat{Y}^{(n)} - Y|\right) \left(E(|Y - f_0(X)|) + E(|f_0(X)|) \right) \\ & \leq 2\|k\|_\infty (\sigma + \|f_0\|_\infty) E\left(|\widehat{Y}^{(n)} - Y|\right), \end{aligned}$$

where we assume that the noise term has variance $\sigma^2 > 0$ and f_0 is bounded and measurable. Hence,

$$\begin{aligned} & \left| E\left(h_1(\widehat{Y}^{(n)}, X)\right) E\left(h_2(\widehat{Y}^{(n)}, X)\right) - E\left(h_1(Y, X)\right) E\left(h_2(Y, X)\right) \right| \\ & \leq 2\|k\|_\infty \left(\sigma + \|f_0\|_\infty + \|k\|_\infty^{1/2} \right) E\left(|\widehat{Y}^{(n)} - Y|\right). \end{aligned}$$

We can deal with the covariance terms in the same way,

$$\begin{aligned} & E\left(h_1(Y, X) h_2(Y, X)\right) - E\left(h_1(\widehat{Y}^{(n)}, X) h_2(\widehat{Y}^{(n)}, X)\right) \\ & = E\left(g_1(Y, X) g_2(Y, X)\right) - E\left(g_1(\widehat{Y}^{(n)}, X) g_2(\widehat{Y}^{(n)}, X)\right) \\ & \quad + E\left(f_1(Y, X) \left(g_2(Y, X) - g_2(\widehat{Y}^{(n)}, X)\right)\right) \\ & \quad + E\left(f_2(Y, X) \left(g_1(Y, X) - g_1(\widehat{Y}^{(n)}, X)\right)\right) \\ & \leq \|k\|_\infty^{1/2} \left(E\left(|Y| \left| g_1(Y, X) - g_1(\widehat{Y}^{(n)}, X) \right| \right) + E\left(|Y| \left| g_2(Y, X) - g_2(\widehat{Y}^{(n)}, X) \right| \right) \right) \\ & \quad + 2\|k\|_\infty^{3/2} E\left(|Y - \widehat{Y}^{(n)}|\right) \\ & \leq 2\|k\|_\infty E\left(|Y| |Y - \widehat{Y}^{(n)}|\right) + 2\|k\|_\infty^{3/2} E\left(|Y - \widehat{Y}^{(n)}|\right) \\ & \leq 2\|k\|_\infty \left(E(|Y - f_0|^2)^{1/2} E\left((Y - \widehat{Y}^{(n)})^2\right)^{1/2} + \|f_0\|_\infty E\left(|Y - \widehat{Y}^{(n)}|\right) \right) \\ & \quad + 2\|k\|_\infty^{3/2} E\left(|Y - \widehat{Y}^{(n)}|\right) \\ & = 2\sigma\|k\|_\infty E\left((Y - \widehat{Y}^{(n)})^2\right)^{1/2} + 2\|k\|_\infty \left(\|f_0\|_\infty + \|k\|_\infty^{1/2} \right) E\left(|Y - \widehat{Y}^{(n)}|\right) \end{aligned}$$

and we can apply (3.14) and (3.15). Combining the above bounds and substituting (3.14) and (3.15) yields the following bound,

$$\begin{aligned}
 (3.27) \quad & \left\| \tilde{\mathfrak{C}}_{c,\tau}^S - \widehat{\mathfrak{C}}_{c,\tau}^{S_n} \right\|_{\text{op}} \\
 & \leq 2\sigma \|k\|_\infty E\left((Y - \widehat{Y}^{(n)})^2\right)^{1/2} + 4\|k\|_\infty (\sigma/2 + \|f_0\|_\infty + \|k\|_\infty^{1/2}) E(|Y - \widehat{Y}^{(n)}|) \\
 & \leq 4\sigma\nu^{1/2} \|k\|_\infty e^{-r_n^2/4\nu} + \sqrt{32\pi\nu} \|k\|_\infty (\sigma/2 + \|f_0\|_\infty + \|k\|_\infty^{1/2}) e^{-r_n^2/2\nu} \\
 & = 4\sigma^2 \|k\|_\infty e^{-r_n^2/4\sigma^2} + \sqrt{32\pi}\sigma \|k\|_\infty (\sigma/2 + \|f_0\|_\infty + \|k\|_\infty^{1/2}) e^{-r_n^2/2\sigma^2},
 \end{aligned}$$

where we used $\nu = \sigma^2$ in the last line. To make use of this bound we need a lower bound on the smallest eigenvalue of $\tilde{\mathfrak{C}}_{c,\tau}^S$. We proceed as in (3.17). Instead of imposing such an assumption directly we can also use an assumption on the covariance operator corresponding to the kernel k and the marginal distribution on \mathcal{X} , which seems more natural. To see this, fix $h \in \mathcal{H}_{\tau,S_f}$, $\|h\|_{\tau,S_f} = 1$, and let $f \in \mathcal{H}_{\kappa_y \upharpoonright S_f \times S_f}$, $g \in \mathcal{H}_{k \upharpoonright \tilde{S} \times \tilde{S}}$ be such that $h = f + \tilde{g}$ and the squared norm of h equals the sum of the squared norms of the f and g (see the discussion around (3.17)).

$$\begin{aligned}
 & \left\langle \tilde{\mathfrak{C}}_{c,\tau}^S h, h \right\rangle_{\tau,S_f} \\
 & = E\left((f(X) + f_0(X)g(X) + \epsilon g(X))^2\right) - E^2(f(X) + f_0(X)g(X)) \\
 & = E\left((f(X) + f_0(X)g(X))^2\right) - E^2(f(X) + f_0(X)g(X)) + \sigma^2 E(g^2(X)) \\
 & \geq \sigma^2 \left\langle \tilde{\mathfrak{C}}^{\tilde{S}} g, g \right\rangle_{k,\tilde{S}}.
 \end{aligned}$$

If k is continuous, we are guaranteed the smallest eigenvalue $\bar{\lambda}_{\star,\tilde{S}}$ of $\tilde{\mathfrak{C}}_k^{\tilde{S}}$ is bounded away from zero (see below (3.17)). In particular, if we choose r_1 such that the last display in (3.27) is upper bounded by $\sigma^2 \bar{\lambda}_{\star,\tilde{S}}/2$ then the smallest eigenvalue $\widehat{\lambda}_\star^{(n)}$ of the capped covariance operator is at least half the smallest eigenvalue $\bar{\lambda}_\star$ of $\tilde{\mathfrak{C}}_{c,\tau}^S$ and is lower bounded by $\sigma^2 \bar{\lambda}_{\star,\tilde{S}}/2$. To guarantee this, we can define the sequence $\{r_n\}_{n \geq 1}$ similarly to before, starting with

$$(3.28) \quad r_1 = 1 \vee 2\sigma \log^{1/2} \left(\frac{22\|k\|_\infty(\sigma + \|f_0\|_\infty + \|k\|_\infty^{1/2})}{\sigma^2 \bar{\lambda}_{\star,\tilde{S}}} \right)$$

and by assuring that the sequence is non-decreasing.

We are now in a position to state a compression result in the regression context along the lines of Proposition 3.3, but in the case where we compress the data simultaneously for the kernel κ_y and ρ . For conciseness we will use the notation $\|\tau\|_{S_{f,n},\infty}$ for $\|\tau \upharpoonright S_{f,n} \times S_{f,n}\|_\infty$ where $S_{f,n} = \{(y, x) : (x, y) \in S_n\}$ and S_n is the support of measure corresponding to the capped random variables.

PROPOSITION 3.6. — *Let $(\mathcal{X} \times \mathbb{R}, \mathfrak{T}, \mathcal{A}, P)$ be a topological measure space such that P is a Radon probability measure which has support S . Let k be continuous bounded kernel function defined on \mathcal{X} such that the corresponding RKHS is finite dimensional and does not contain the constant functions. Let $(X_1, Y_1), \dots, (X_n, Y_n)$ be i.i.d. random variables with law P , and assume that $Y_i = f_0(X_i) + \epsilon_i$, for all*

$i \leq n$, where f_0 is a measurable and bounded function and $\epsilon_1, \dots, \epsilon_n$ are i.i.d. sub-Gaussian random variables with variance $0 < \sigma^2$ which are independent of X_1, \dots, X_n . Consider the kernel function $\tau = \kappa_y + \rho$ on $\mathbb{R} \times \mathcal{X}$ and let $\bar{\lambda}_{*,\tilde{S}}$ be the smallest eigenvalue of the covariance operator $\tilde{\mathfrak{C}}_k^{\tilde{S}}$ corresponding to the kernel function $k|_{\tilde{S} \times \tilde{S}}$, $\tilde{S} = \overline{\{x : (x, y) \in S\}}$. Furthermore, let $\bar{\lambda}_*$ be the smallest eigenvalue of $\tilde{\mathfrak{C}}_{c,\tau}^S$, then $\bar{\lambda}_* > 0$. Choose $q \in (0, 1)$ and define the sequence $\{r_n\}_{n \geq 1}$ in the following way: choose r_1 as in (3.28) and for $n > 2$ let $r_n = r_1 \vee \sqrt{2}\sigma \vee \sqrt{2}\sigma \log^{1/2}(16n^{1/2}\sigma^2\|k\|_\infty/q)$. For the smallest eigenvalue $\widehat{\lambda}_*^{(n)}$ of $\widehat{\mathfrak{C}}_{c,\tau}^{S_n}$ it holds that $\widehat{\lambda}_*^{(n)} \geq \bar{\lambda}_*/2 \geq \sigma^2\bar{\lambda}_{*,\tilde{S}}/2 > 0$, where S_n is the support of the law of $P^{(n)}$ corresponding to $(X, \widehat{Y}^{(n)})$. There exists a ball of radius $\widehat{\delta}^{(n)} = \widehat{\lambda}_*^{(n)}/2\|\tau\|_{S_{f,n},\infty}^{1/2} \geq \widehat{\lambda}_*^{(n)}/2(\|k\|_\infty + (\|f_0\|_\infty + r_n)^{1/2}\|k\|_\infty^{1/2})$ around $\widehat{\mathbf{m}}_\tau$ within the affine space spanned by $\widehat{C}_\tau = \{\tau((y, x), \cdot) : (x, y) \in S_n\}$ as a subset of \mathcal{H}_τ . Furthermore, for any $q \in (0, 1)$ and whenever n is (strictly) greater than

$$\left(\frac{8\|\tau\|_{S_{f,n},\infty} \left(\sqrt{2 \log(12/q)} + 192\|\tau\|_{S_{f,n},\infty} / \widehat{\lambda}_*^{(n)} \right)^2}{\left(\widehat{\lambda}_*^{(n)} \right)^2} \right)^2 \vee \left(\frac{16\|\tau\|_{S_{f,n},\infty}^{1/2} + \sqrt{288 \log(4/q)}}{\widehat{\delta}^{(n)}} \right)^2$$

then with probability $1 - q$ there exists a ball of radius $\widehat{\delta}^{(n)}/4$ around $\widehat{\mathbf{m}}_{\tau,n}$ in $\widehat{C}_{\tau,n}$ within the affine subspace spanned by \widehat{C}_τ and

$$\|\widehat{\mathbf{m}}_{\tau,n} - \mathbf{m}_{\tau,n}\|_\tau \leq n^{-1/2}.$$

Proof. — Most of the statement has already been derived. Just note that

$$\|\tau\|_{S_{f,n},\infty}^{1/2} \leq (\|k^2\|_\infty + (\|f_0\|_\infty + r_n)\|k\|_\infty)^{1/2} \leq \|k\|_\infty + (\|f_0\|_\infty + r_n)^{1/2}\|k\|_\infty^{1/2}.$$

For the definition of r_n and the bound on the difference between the mean and the capped mean we could use essentially the same bound as in Proposition 3.3. Instead we use here a slightly different bound to demonstrate how the arguments can be varied: when $r_n \geq \sqrt{2}\sigma$ it follows from (3.26) that with probability \tilde{q} ,

$$\|\mathbf{m}_{\tau,n} - \widehat{\mathbf{m}}_{\tau,n}\|_\tau^2 \leq \sqrt{8}\sigma\|k\|_\infty^{1/2} e^{-r_n^2/4\nu} \tilde{q}^{-1/2}.$$

Setting the right side equal to $n^{-1/4}$, setting $\tilde{q} = q/2$ and solving for r_n gives

$$r_n = \sqrt{2}\sigma \log^{1/2}(16n^{1/2}\sigma^2\|k\|_\infty/q). \quad \square$$

4. Compression using the CGM and related approaches

We discuss in this section two methods to compress \mathbf{m}_n . The bottleneck in both algorithms is the computation of the vector

$$s = (\langle k(X_1, \cdot), \mathbf{m}_n \rangle, \dots, \langle k(X_n, \cdot), \mathbf{m}_n \rangle)^\top.$$

If s is available then the remaining parts of the two algorithms that we analyze have a runtime of nl , where l is the number of iterations the algorithms are run for.

In particular, for large n , l will be in the order of $\log(n)$ when using the classical CGM and of order $n^{1/2}$ when using the kernel herding algorithm. This then results in a runtime of $O(n \log(n))$ and $O(n^{3/2})$ respectively to gain a representation of \mathbf{m}_n . The CGM achieves in this context a compression down to $\log(n)$ many points and the kernel herding algorithm down to $n^{1/2}$, that is, if we have a ball of sufficient size around \mathbf{m}_n in C_n .

A naive algorithm to compute s has a runtime of $O(n^2)$. In fact, a brute-force computation needs to compute all pairs $k(X_i, X_j)$, $i, j \leq d$, and the computational complexity is the same as the computational complexity of computing the kernel matrix itself (though the algorithm only needs $O(n)$ memory instead of $O(n^2)$). However, there is hope for faster algorithms. For instance, when we have a finite dimensional RKHS with dimension d then we can represent \mathbf{m}_n as a linear combination of d points and s can be computed in $O(nd)$ time. Computing the representation of \mathbf{m}_n needs another d^2 steps. In practice this is not useful because we would derive an exact representation of \mathbf{m}_n based on d data points to find an approximation of \mathbf{m}_n using more than d data points. Ideally, we would hope for an algorithm that can compute, or approximate, s in $n \log n$ steps independently of the dimension d of the Hilbert space. Alternatively, we could try to modify the main algorithms itself to mitigate the complexity of computing s . There are some standard ways to deal with large scale data in the context of the CGM as summarized in [BCC⁺22]. However, they do not lead to computational benefits in our particular setting. We discuss a promising alternative that is based on a divide and conquer approach in some detail below (Section 4.1.1 and 4.2.2) after analyzing the standard algorithms. We also include short discussions on how to adapt these methods when aiming for compressing other quantities like the covariance and how to use the CGM to compress the data in the case of kernel regression.

4.1. Kernel herding and subset selection

Let us start by stating a version of the kernel herding algorithm [CWS10] for compressing the empirical measure.

Algorithm 1 (The kernel herding algorithm for compressing \mathbf{m}_n)

Input: sample X_1, \dots, X_n , kernel k , number of points in the coreset T .

Initialise: let $w_1 = k(X_1, \cdot) - \mathbf{m}_n$ and $\iota(1) = 1$, iterate through $t \geq 2$:

choose $i^* \in \arg \max_{i \leq n} \langle w_t, k(X_i, \cdot) \rangle$

set $\iota(t) = i^*$, $w_{t+1} = w_t - (k(X_{i^*}, \cdot) - \mathbf{m}_n)$, and $\widehat{\mathbf{m}}_t = \frac{1}{t} \sum_{u=1}^t k(X_{\iota(u)}, \cdot)$

Stop when $t = T$ and return the approximation $\widehat{\mathbf{m}}_T$.

The index function $\iota : \{1, \dots, T\} \rightarrow \{1, \dots, n\}$ tracks the samples that we include in the coreset and the elements w_t measure the error between \mathbf{m}_n and $\widehat{\mathbf{m}}_t$ as $\|w_t\| =$

$t\|\mathbf{m}_n - \widehat{\mathbf{m}}_t\|$. The algorithm converges with a rate of $1/t$ if, and only if, the sequence of weights w_t is bounded. In other words, if the sequence diverges then the algorithm converges with a slow rate. It is easy to show that $\|w_t\|$ stays bounded when a ball of radius $\delta > 0$ exists around \mathbf{m}_n in C_n and that $\|\mathbf{m}_n - \widehat{\mathbf{m}}_t\| \leq \alpha_{KH}/\delta t$ for a constant α_{KH} and all t . In particular, we can choose

$$\alpha_{KH} = 8\|k\|_\infty.$$

Also, notice that the same bound holds when a ball of radius δ exists around \mathbf{m}_n in the affine span of C_n .

Instead of running the algorithm for T iterations independent of the approximation error we can also use the approximation error as a stopping criterion. The approximation error $(1/t)\|w_t\| = (\|\widehat{\mathbf{m}}_t\|^2 - 2\langle \widehat{\mathbf{m}}_t, \mathbf{m}_n \rangle + \|\mathbf{m}_n\|^2)^{1/2}$ can be computed exactly in $O(n^2)$ if we prevent the algorithm from running for more than n iterations. In detail, pre-computing $\|\mathbf{m}_n\|$ and s can be done in $O(n^2)$. Also, $\|\widehat{\mathbf{m}}_t\|$ can be computed in $O(t)$ given $\|\widehat{\mathbf{m}}_{t-1}\|$ by using that $\|\widehat{\mathbf{m}}_t\|^2 = \|\widehat{\mathbf{m}}_{t-1}\|^2 + 2\langle k(X_{i(t)}, \cdot), \widehat{\mathbf{m}}_{t-1} \rangle + k(X_{i(t)}, X_{i(t)})$. Similarly, $\langle \widehat{\mathbf{m}}_t, \mathbf{m}_n \rangle$ can be easily gained from $\langle \widehat{\mathbf{m}}_{t-1}, \mathbf{m}_n \rangle$ by using $\langle \widehat{\mathbf{m}}_t, \mathbf{m}_n \rangle = \langle \widehat{\mathbf{m}}_{t-1}, \mathbf{m}_n \rangle + s_{i(t)}$. A natural stopping point for the algorithm is an approximation error of $n^{-1/2}$ which guarantees that $\|\widehat{\mathbf{m}}_t - \mathbf{m}\|$ will be of the same order as $\|\mathbf{m}_n - \mathbf{m}\|$.

The compression of this algorithm is sub-optimal but it has the advantage that it returns a coreset. The CGM, which we discuss below, achieves a significantly better compression but does not return a coreset of samples.

It is easy to gain high probability guarantees for the approximation error of a compression that uses $n^{1/2}$ many points, under the conditions stated in Section 2.5. With a bit more work it is also possible to control the expected approximation error. We summarize these in the following proposition under a Lipschitz assumption on the kernel function, assuming that we have a Mercer kernel and that the constant functions are in the RKHS; in particular, we assume that k is a continuous kernel on $[0, 1]^l$, which is a sufficient assumption for Mercer’s theorem to hold. When discussing the CGM we give a similar proposition which uses instead an assumption on the covariance operator; the aim is to highlight how the various assumptions can be combined with the algorithms.

PROPOSITION 4.1. — *Let X_1, \dots, X_n be i.i.d. random variables on some probability space (Ω, \mathcal{A}, P) , which attain values in $\mathcal{X} = [0, 1]^l, l \geq 1$, and let k be a continuous kernel function on \mathcal{X} such that the corresponding RKHS \mathcal{H} is d -dimensional, $1 \leq d < \infty$, functions $h \in \mathcal{H}$ are Lipschitz continuous in the sense of (2.13) with Lipschitz constant $L > 0$, and $\mathbf{1} \in \mathcal{H}$. Furthermore, assume that the law of X_1 has a density p on \mathcal{X} and $\inf_{x \in \mathcal{X}} p(x) \geq c > 0$ for some constant c . Let $\widetilde{\lambda}_d$ be the smallest eigenvalue of the Mercer decomposition. With probability $1 - q, q \in (0, 1)$,*

$$\|\widehat{\mathbf{m}}_{\lceil n^{1/2} \rceil} - \mathbf{m}_n\| \leq \frac{32\|k\|_\infty}{\delta} n^{-1/2}$$

whenever

$$n > \left(\frac{\sqrt{2 \log(4/q)} + 96\|k\|_\infty^{1/2}/\delta}{c\beta_l(\delta/8L)^l} \right)^2 \vee \left(\frac{\sqrt{2}\|k\|_\infty^{1/2}(1 + \sqrt{\log(2/q)})}{\delta/4} \right)^2$$

and where

$$\delta = \frac{c(\tilde{\lambda}_d^{1/2} \wedge L)^{l+1} \beta_l}{(l+1)(2L)^l}.$$

Furthermore, for any $t \geq 1$ and $n \geq 1$,

$$E(\|\widehat{\mathbf{m}}_t - \mathbf{m}_n\|) \leq 32\|k\|_\infty/t\delta + 4\|k\|_\infty^{1/2} \exp(-(c_1 n^{1/2} - c_2)_+^2)/t^{1/2},$$

where c_1 and c_2 can be chosen as

$$c_1 = \frac{\delta}{\sqrt{32}\|k\|_\infty^{1/2}} \wedge c\beta_l(\delta/8L)^l \quad \text{and} \quad c_2 = 1 \vee 96\|k\|_\infty^{1/2}/\delta$$

If the stopping criterion of the algorithm is an error of $\|\widehat{\mathbf{m}}_{\widehat{t}} - \mathbf{m}_n\| \leq n^{-1/2}$, i.e. $\widehat{t} = \inf_{t \geq 1} \|\widehat{\mathbf{m}}_t - \mathbf{m}_n\| \leq n^{-1/2}$, and if the infimum is greater than n , then $\widehat{t} = n$ and $\widehat{\mathbf{m}}_{\widehat{t}} = \mathbf{m}_n$, then

$$E(\widehat{t}) \leq \left[32\|k\|_\infty n^{1/2}/\delta \right] + 2n \exp(-(c_1 n^{1/2} - c_2)_+^2)$$

whenever $n \geq N$.

Proof. — The *first statement* follows directly from Theorem 2.16 and the bound on the error of the kernel herding algorithm. For the *second statement*, observe that

$$\Pr(\|\mathbf{m}_n - \mathbf{m}\| \geq \delta/4) \leq \exp\left(-\left(\frac{\delta n^{1/2}}{\sqrt{32}\|k\|_\infty^{1/2}} - 1\right)^2\right)$$

whenever $n \geq (\sqrt{32}\|k\|_\infty^{1/2}/\delta)^2$, follows by the same argument as in Theorem 2.16. Similarly, there is a ball of radius $\delta/2$ around \mathbf{m} in C_n in the affine span of C_n with probability at least

$$1 - (1/2) \exp(-(n^{1/2} c\beta_l(\delta/8L)^l - 96\|k\|_\infty^{1/2}/\delta)^2)$$

whenever $n \geq (96(8L)^l\|k\|_\infty^{1/2}/c\beta_l\delta^{l+1})^2$. In particular, at the event that there also exists a ball of radius $\delta/2$ around \mathbf{m} and $\|\mathbf{m}_n - \mathbf{m}\| < \delta/4$ we have a ball of radius $\delta/4$ around \mathbf{m}_n in C_n and $\|\widehat{\mathbf{m}}_t - \mathbf{m}_n\| \leq 32\|k\|_\infty/\delta t$. The probability that either \mathbf{m}_n is not $\delta/4$ close to \mathbf{m} or that there is no ball of radius $\delta/2$ around \mathbf{m} in C_n is upper bounded by

$$(4.1) \quad \exp\left(-\left(\frac{\delta n^{1/2}}{\sqrt{32}\|k\|_\infty^{1/2}} - 1\right)^2\right) + (1/2) \exp(-(n^{1/2} c\beta_l(\delta/8L)^l - 96\|k\|_\infty^{1/2}/\delta)^2).$$

To simplify the bound consider the expressions for the constants c_1, c_2 in the theorem statement

$$c_1 = \frac{\delta}{\sqrt{32}\|k\|_\infty^{1/2}} \wedge c\beta_l(\delta/8L)^l \quad \text{and} \quad c_2 = 1 \vee 96\|k\|_\infty^{1/2}/\delta.$$

The constant c_1 is simply the minimum of the left hand terms in the two exponentials in Eq. (4.1) and c_2 the maximum of the right hand terms. Using these we can bound the sum of the two exponentials as long as $c_1 \geq c_2$. We deal with the case that

$c_2 > c_1$ by using the operator $(a)_+ = a \vee 0$ in the following expression upper bound on Eq. (4.1)

$$2 \exp\left(- (c_1 n^{1/2} - c_2)_+^2\right)$$

Notice that even when there is no ball around \mathbf{m}_n , for $t \geq 1$ it holds that

$$\|w_{t+1}\|^2 = \|k(X_{\iota(t)}, \cdot) - \mathbf{m}_n\|^2 + \|w_t\|^2 - 2\langle k(X_{\iota(t)}, \cdot) - \mathbf{m}_n, \mathbf{m}_n \rangle.$$

Now, $\langle k(X_{\iota(t)}, \cdot), \mathbf{m}_n \rangle \geq \|\mathbf{m}_n\|^2$ since \mathbf{m}_n lies in C_n which has extremes $k(X_1, \cdot), \dots, k(X_n, \cdot)$, and

$$t^2 \|\widehat{\mathbf{m}}_t - \mathbf{m}_n\|^2 = \|w_{t+1}\|^2 \leq \sum_{i=1}^t \|k(X_{\iota(i)}, \cdot) - \mathbf{m}_n\|^2 \leq 4t \|k\|_\infty.$$

Hence, in the event that there is no ball we still have that $\|\widehat{\mathbf{m}}_t - \mathbf{m}_n\| \leq 2t^{-1/2} \|k\|_\infty^{1/2}$.

Combining the bound for the case that there is a ball with the bound for the case that there is no ball, we find that

$$E(\|\widehat{\mathbf{m}}_t - \mathbf{m}_n\|) \leq 32 \|k\|_\infty / t\delta + 4 \|k\|_\infty^{1/2} \exp(- (c_1 n^{1/2} - c_2)_+^2) t^{-1/2}$$

whenever n is large enough and with c_1, c_2 as in the theorem statement.

The *third statement* follows along similar lines. In the event that we have a ball of size $\delta/4$ it follows that $\|\widehat{\mathbf{m}}_t - \mathbf{m}_n\| \leq 32 \|k\|_\infty / t\delta$. Setting the right hand side to $n^{-1/2}$ leads to $\widehat{t} \leq \lceil 32 \|k\|_\infty n^{1/2} / \delta \rceil$. If this event does not occur then $\|\widehat{\mathbf{m}}_t - \mathbf{m}_n\| \leq 2 \|k\|_\infty^{1/2} t^{-1/2}$ and $\widehat{t} \leq \lceil 4 \|k\|_\infty n \rceil$, but the algorithm stops when $\widehat{t} > n$ and the trivial upper bound $\widehat{t} \leq n$ is more useful. Combining these we find that

$$E(\widehat{t}) \leq \lceil 32 \|k\|_\infty n^{1/2} / \delta \rceil + 2n \exp(- (c_1 n^{1/2} - c_2)_+^2)$$

when n is large enough. □

4.1.1. Avoiding the explicit computation of s

There are various ways one can try to reduce the computation time. For instance, the stochastic conditional gradient method seems like a promising candidate. An alternative way to mitigate the cost of computing s is to split the data into batches of size about $n^{1/2}$, which implies that for each batch the corresponding vector can be computed in $O(n)$. Algorithm 2 implements such a version of kernel herding. There are a number of interesting observations that can be made when following this route. We will discuss a few such observations in this section and in Section 4.2.2 below.

In terms of Algorithm 2, notice that the number of samples per batch can always be chosen in $\{\ell - 1, \ell, \ell + 1, \ell + 2\}$ to guarantee that $n = \sum_{i=1}^\ell \ell_i$ because $\ell(\ell - 1) \leq (n^{1/2} + 1/2)(n^{1/2} - 1/2) < n \leq (\ell + 1/2)^2 \leq \ell(\ell + 2)$ and by a version of Algorithm 1 we mean the obvious modification where instead of $k(X_1, \cdot), \dots, k(X_n, \cdot)$ we use $\widehat{\mathbf{m}}_1, \dots, \widehat{\mathbf{m}}_\ell$ to approximate \mathbf{m}_n . The algorithm works by specifying the number of iterations for the kernel herding algorithm. Alternatively, it makes sense to run the first ℓ optimization algorithms as well as the last optimization step until an error of $n^{1/4}$ is attained.

The runtime of the algorithm can be computed in the following way: observe that, initially, the standard kernel herding algorithm is applied ℓ times to about ℓ

many samples and the overall order of runtime for the first part is $O(\ell^3) = O(n^{3/2})$. Also, observe that, given $\widehat{\mathbf{m}}_1, \dots, \widehat{\mathbf{m}}_\ell$, an approximation of $\|\mathbf{m}_n\|$ can be computed in $\ell \lceil n^{1/4+\alpha} \rceil^2$, which is of order $O(n^{1+2\alpha})$, by using the approximation $\widetilde{\mathbf{m}}_n = (1/n)(\ell_1 \widehat{\mathbf{m}}_1 + \dots + \ell_\ell \widehat{\mathbf{m}}_\ell)$. Similarly, $\widehat{K} = (\langle \widehat{\mathbf{m}}_i, \widehat{\mathbf{m}}_j \rangle)_{i,j \leq \ell}$ can be computed in $\ell^2 \lceil n^{1/4+\alpha} \rceil^2$ which is of order $O(n^{3/2+2\alpha})$ and, given \widehat{K} , the vector $\widehat{s} = (\langle \widehat{\mathbf{m}}_1, \widetilde{\mathbf{m}}_n \rangle, \dots, \langle \widehat{\mathbf{m}}_\ell, \widetilde{\mathbf{m}}_n \rangle)^\top$ can be computed in ℓ^2 steps, which is of order $O(n)$. Given \widehat{s} and $\|\widetilde{\mathbf{m}}_n\|$ the second application of the kernel herding algorithm can be run in $\lceil n^{1/4+\alpha} \rceil \ell$, which is of order $O(n^{3/4+\alpha})$, and Algorithm 2 has an overall order of $O(n^{3/2+2\alpha})$.

Algorithm 2 (A version of kernel herding that avoids the explicit computation of s)

Input: sample X_1, \dots, X_n , kernel $k, \alpha > 0$.

Initialise: let $\ell = \lfloor n^{1/2} + 1/2 \rfloor$.

Split sample into ℓ disjoint batches:

$$X_{j_1}, \dots, X_{j_{\ell_j}}, j \leq \ell, \quad \text{with each } \ell_j \in \{\ell - 1, \ell, \ell + 1, \ell + 2\}.$$

Apply Algorithm 1 with $T = \lceil n^{1/4+\alpha} \rceil$ to each batch to get $\widehat{\mathbf{m}}_1, \dots, \widehat{\mathbf{m}}_\ell$.

Compute $\|\widetilde{\mathbf{m}}_n\|$ and $\widehat{s} = (\langle \widehat{\mathbf{m}}_1, \widetilde{\mathbf{m}}_n \rangle, \dots, \langle \widehat{\mathbf{m}}_\ell, \widetilde{\mathbf{m}}_n \rangle)^\top$, where $\widetilde{\mathbf{m}}_n = \frac{1}{n} \sum_{i=1}^{\ell} \ell_i \widehat{\mathbf{m}}_i$.

Apply a version of Algorithm 1 with $T = \lceil n^{1/4+\alpha} \rceil$ to $\widetilde{\mathbf{m}}_n$ using $\|\widetilde{\mathbf{m}}_n\|$ and \widehat{s} to get ι .

Return the approximation $\frac{1}{T} \sum_{i=1}^T \widehat{\mathbf{m}}_{\iota(i)}$.

Quantifying the approximation error of this algorithm is more difficult and in the following we only highlight some of the challenges that one has to address to control the approximation error. For n large enough the difference $\|\widehat{\mathbf{m}}_i - (1/\ell_i) \sum_{j=1}^{\ell_i} k(X_{ij}, \cdot)\|$ is with high probability of order $n^{-1/4}$ for all $i \leq \ell$. Furthermore,

$$\|\mathbf{m}_n - \widetilde{\mathbf{m}}_n\|^2 = \left\| \mathbf{m}_n - \frac{1}{n} \sum_{i=1}^{\ell} \ell_i \widehat{\mathbf{m}}_i \right\|^2 = \left\| \frac{1}{n} \sum_{i=1}^{\ell} \sum_{j=1}^{\ell_i} (k(X_{ij}, \cdot) - \widehat{\mathbf{m}}_i) \right\|^2.$$

Notice that the Bochner integral is

$$E \left(\sum_{j=1}^{\ell_i} (k(X_{ij}, \cdot) - \widehat{\mathbf{m}}_i) \right) = \ell_i (\mathbf{m} - E(\widehat{\mathbf{m}}_i)).$$

Furthermore, observe that the $\widehat{\mathbf{m}}_1, \dots, \widehat{\mathbf{m}}_\ell$ are independent random variables since they are functions of separate samples and that $E(\langle \mathbf{X}, \mathbf{Y} \rangle) = \langle E(\mathbf{X}), E(\mathbf{Y}) \rangle$ for

independent random variables in $\mathcal{L}^2(P; \mathcal{H})$. Hence,

$$\begin{aligned} & E\left(\left\|\mathbf{m}_n - \frac{1}{n} \sum_{i=1}^{\ell} \ell_i \widehat{\mathbf{m}}_i\right\|^2\right) \\ &= \frac{1}{n^2} \sum_{i=1}^{\ell} E\left(\left\|\sum_{j=1}^{\ell_i} (k(X_{ij}, \cdot) - \widehat{\mathbf{m}}_i)\right\|^2\right) + \frac{1}{n^2} \sum_{i \neq j}^{\ell} \ell_i \ell_j \langle \mathbf{m} - E(\widehat{\mathbf{m}}_i), \mathbf{m} - E(\widehat{\mathbf{m}}_j) \rangle \\ &\approx n^{-3/2} E\left(\left\|\sum_{j=1}^{\ell_1} (k(X_{1j}, \cdot) - \widehat{\mathbf{m}}_1)\right\|^2\right) + \|\mathbf{m} - E(\widehat{\mathbf{m}}_1)\|^2, \end{aligned}$$

where we have an approximation in the last line since the ℓ_i 's are not necessarily all equal. A first difficulty is to determine the bias $\|\mathbf{m} - E(\widehat{\mathbf{m}}_1)\|$ that the kernel herding algorithm introduces. A simple bound on the bias is gained by using $\|\mathbf{m} - E(\widehat{\mathbf{m}}_1)\| \leq E(\|\mathbf{m}_n - \widehat{\mathbf{m}}_1\|)$ and Proposition 4.1 can be used to bound this by about $n^{-1/4}$ which implies a bound on the squared bias of order $n^{-1/2}$. This bound is of no use since we need a bias of order $1/n$ or less. The other term behaves approximately as $1/n$ if there is a ball of size $\delta > 0$ around \mathbf{m} in C and n is large enough. In particular, under the conditions of Proposition 4.1,

$$E\left(\left\|\frac{1}{\ell_1} \sum_{j=1}^{\ell_1} (k(X_{1j}, \cdot) - \widehat{\mathbf{m}}_1)\right\|^2\right) \leq 16\alpha^2/n^{1/2}\delta^2 + 8\|k\|_{\infty} \exp(-(c_1 n^{1/4} - c_2)^2) n^{-1/4}.$$

Up to the exponential term on the right, we have that

$$n^{-3/2} E\left(\left\|\sum_{j=1}^{\ell_1} (k(X_{1j}, \cdot) - \widehat{\mathbf{m}}_1)\right\|^2\right) \lesssim n^{-3/2} \ell_1^2 n^{-1/2} \approx 1/n.$$

If the bias is also of order $O(1/n)$ then

$$\Pr\left(\left\|\mathbf{m}_n - \frac{1}{n} \sum_{i=1}^{\ell} \ell_i \widehat{\mathbf{m}}_i\right\| \geq t\right) \lesssim \frac{1}{nt^2}.$$

In particular, for any $\beta > 0$,

$$\Pr\left(\|\mathbf{m}_n - \widetilde{\mathbf{m}}_n\| \geq n^{-1/2+\beta}\right) = \Pr\left(\left\|\mathbf{m}_n - \frac{1}{n} \sum_{i=1}^{\ell} \ell_i \widehat{\mathbf{m}}_i\right\| \geq n^{-1/2+\beta}\right) \lesssim n^{-2\beta}.$$

To summarize, if the bias is of order $O(1/n)$ we will have with high probability an approximation of \mathbf{m}_n that has an error of order $n^{-1/2+\beta}$ and this approximation consists of approximately $T\ell \approx n^{3/4}$ many points.

The second application of the kernel herding algorithm aims to compress this further. In particular, if with high probability there is a ball around \mathbf{m}_n in the convex set $\text{ch}\{\widehat{\mathbf{m}}_i : i \leq \ell\}$, then we can hope that $n^{1/4}$ many of the $\widehat{\mathbf{m}}_i$ are sufficient to approximate \mathbf{m}_n with an error of order $n^{-1/2+\beta}$. This would imply that an approximation with $n^{1/2}$ many elements is sufficient. However, since $\widehat{\mathbf{m}}_i$ converges to \mathbf{m}_n as n goes to infinity, the size of such a ball has to be a function of n and will shrink

with n . This itself does not imply that the algorithm will converge slowly since the smaller δ might be set-off by a smaller size of the convex set. In any case, a detailed analysis of the interplay between $\text{ch}\{\widehat{\mathbf{m}}_i : i \leq \ell\}$ and \mathbf{m}_n is necessary to understand the compression that can be achieved by this algorithm and variations thereof.

Let us conclude our discussion of these algorithms with a final simple observation. The elements $\bar{\mathbf{m}}_i = (1/\ell_i) \sum_{j=1}^{\ell_i} k(X_{ij}, \cdot)$, which we are approximating with $\widehat{\mathbf{m}}_i$, can be interpreted as a sequence of independent and identically distributed (up to differences in the ℓ_i 's) random variables whose second moment is given by

$$\begin{aligned} E(\|\bar{\mathbf{m}}_i\|^2) &= \frac{1}{\ell_i^2} \sum_{j=1}^{\ell_i} E(k(X_{ij}, X_{ij})) + \frac{1}{\ell_i^2} \sum_{u \neq v}^{\ell_i} \langle E(k(X_{iu}, \cdot)), E(k(X_{iv}, \cdot)) \rangle \\ &= \frac{1}{\ell_i} + \frac{\ell_i - 1}{\ell_i} \|\mathbf{m}\|^2, \end{aligned}$$

whenever $k(x, x) = 1$ for all $x \in \mathcal{X}$. Hence,

$$E(\|\bar{\mathbf{m}}_i - \mathbf{m}\|^2) = E(\|\bar{\mathbf{m}}_i\|^2) - \|\mathbf{m}\|^2 \approx n^{-1/2} \quad \text{and} \quad E(\|\bar{\mathbf{m}}_i - \mathbf{m}\|) \lesssim n^{-1/4}.$$

4.2. Better compression with the CGM

A significantly better compression can be attained by using the CGM. The downside of using the CGM is that no coreset of datapoints is generated but some convex combination of the images of the data points in \mathcal{H} that approximates \mathbf{m}_n well. The standard CGM for compressing \mathbf{m}_n is given below.

Algorithm 3 (The CGM for compressing \mathbf{m}_n .)

Input: sample X_1, \dots, X_n , kernel k , number of iterations T .

Initialise: let $\widehat{\mathbf{m}}_1 = k(X_1, \cdot)$, $\alpha_{11} = 1$ and $\iota(1) = 1$, iterate through $t \geq 2$:

choose $i^* \in \arg \max_{i \leq n} \langle k(X_i, \cdot), \widehat{\mathbf{m}}_{t-1} - \mathbf{m}_n \rangle$,

let $\alpha^* = \frac{\langle k(X_{i^*}, \cdot) - \widehat{\mathbf{m}}_{t-1}, \widehat{\mathbf{m}}_{t-1} - \mathbf{m}_n \rangle}{\|k(X_{i^*}, \cdot) - \widehat{\mathbf{m}}_{t-1}\|^2} \wedge 1$,

set $\iota(t) = i^*$, $\alpha_{tt} = \alpha^*$ and for all $u \leq t-1$, $\alpha_{tu} = (1 - \alpha^*)\alpha_{t-1,u}$,

and let $\widehat{\mathbf{m}}_t = \sum_{u=1}^t \alpha_{tu} k(X_{\iota(u)}, \cdot)$.

Stop when $t = T$ and return the approximation $\widehat{\mathbf{m}}_T$.

Notice that $\alpha^* \geq 0$ since $k(X_{i^*}, \cdot)$ maximizes the inner product between any element in C_n and $\widehat{\mathbf{m}}_{t-1} - \mathbf{m}_n$. This algorithm guarantees that the error is bounded by

$$\|\mathbf{m}_n - \widehat{\mathbf{m}}_t\| \leq 2\|k\|_\infty^{1/2} \exp\left(-\frac{\delta(t-1)}{6\|k\|_\infty^{1/2}}\right),$$

when a ball of size δ exists around \mathbf{m}_n in C_n within the affine subspace spanned by C_n [BT04, Proposition 3.2] and with S denoting the support of the law of X_1 .

The run-time of this algorithm is again dominated by the $O(n^2)$ run-time cost needed to compute s . When s is available the run-time reduces to $O(Tn)$: the arg max step can be performed in $O(n)$ given s and when the inner products $\langle k(X_i, \cdot), \widehat{\mathbf{m}}_{t-1} \rangle$ are available. Similarly, if s , the inner products $\langle k(X_i, \cdot), \widehat{\mathbf{m}}_{t-1} \rangle$, $\|\widehat{\mathbf{m}}_{t-1}\|$, $\langle \widehat{\mathbf{m}}_{t-1}, \mathbf{m}_n \rangle$ and $\|k(X_{i^*}, \cdot) - \widehat{\mathbf{m}}_{t-1}\|$ are available, it is possible to compute α^* in $O(1)$. The norm term in the denominator can be computed in $O(1)$ from $\|\widehat{\mathbf{m}}_{t-1}\|$ and the inner products $\langle k(X_i, \cdot), \widehat{\mathbf{m}}_{t-1} \rangle$. The coefficients α_{tu} can be computed in $O(T^2)$. Updating the elements $\langle k(X_i, \cdot), \widehat{\mathbf{m}}_{t-1} \rangle$ to

$$\langle k(X_i, \cdot), \widehat{\mathbf{m}}_t \rangle = (1 - \alpha^*) \langle k(X_i, \cdot), \widehat{\mathbf{m}}_{t-1} \rangle + \alpha^* \langle k(X_i, \cdot), k(X_{i^*}, \cdot) \rangle$$

can be done in $O(n)$. Furthermore, $\|\widehat{\mathbf{m}}_t\|^2 = (1 - \alpha^*)^2 \|\widehat{\mathbf{m}}_{t-1}\|^2 + (\alpha^*)^2 k(X_{i^*}, X_{i^*}) + 2\alpha^*(1 - \alpha^*) \langle k(X_{i^*}, \cdot), \widehat{\mathbf{m}}_{t-1} \rangle$ and $\langle \widehat{\mathbf{m}}_t, \mathbf{m}_n \rangle = (1 - \alpha^*) \langle \widehat{\mathbf{m}}_{t-1}, \mathbf{m}_n \rangle + \alpha^* \langle k(X_{i^*}, \cdot), \mathbf{m}_n \rangle$ can both be updated in $O(1)$. In particular, if we aim for a compression down to $T = \log(n)$ elements then the run-time of the algorithm is $O(n \log(n))$, if s is available.

As for the kernel herding algorithm, it is easy to bound, with high probability, the approximation error, as well as the expected error and the number of data points that are needed for the approximation when the stopping criterion is a pre-specified error. In the following proposition, we bound the approximation error given that the algorithm is run for $\lceil 12\|k\|_\infty^{1/2} \log(n)/\delta \rceil$ many iterations. Alternatively, it is possible to use $\lceil \log^\gamma(n) \rceil$, with $\gamma > 1$, as a stopping criterion that does not depend on the unknown quantity δ . For large enough n , $\lceil 12\|k\|_\infty^{1/2} \log(n)/\delta \rceil \leq \lceil \log^\gamma(n) \rceil$ and the guarantees will carry over to that setting.

PROPOSITION 4.2. — *Let $(\mathcal{X}, \mathcal{A}, P)$ be some probability space with P being a topological measure that is τ -additive, and with measurable kernel function k defined on \mathcal{X} such that the corresponding RKHS \mathcal{H} is finite dimensional. Furthermore, let X_1, \dots, X_n be i.i.d. random variables attaining values in \mathcal{X} and with law P . Assume that $\|k\|_\infty < \infty$, and that the centered covariance operator $\widetilde{\mathfrak{C}}_c$ has an eigen-decomposition with smallest non-zero eigenvalue being $\bar{\lambda}_d$. Let $\beta = 3\|k\|_\infty^{1/2}/\delta$ then with probability $1 - q$, $q \in (0, 1)$,*

$$\left\| \widehat{\mathbf{m}}_{\lceil 2\beta \log(n) \rceil} - \mathbf{m}_n \right\| \leq 2\|k\|_\infty^{1/2} n^{-1/2}$$

whenever n is (strictly) greater than

$$\left(\frac{8\|k\|_\infty \left(\sqrt{2 \log(4/q)} + 192\|k\|_\infty / \bar{\lambda}_d \right)}{\bar{\lambda}_d^2} \right)^2 \vee \left(\frac{12\|k\|_\infty \left(1 + \sqrt{\log(2/q)} \right)}{\bar{\lambda}_d} \right)^2$$

Let $\delta = \bar{\lambda}_d/2\|k\|_\infty^{1/2}$. For any $t \geq 1, n \geq 1$,

$$E(\|\widehat{\mathbf{m}}_t - \mathbf{m}_n\|) \leq \exp(-\delta(t-1)/24\|k\|_\infty^{1/2}) + 6\|k\|_\infty^{1/2} \exp(-(c_3 n^{1/2} - c_4)_+^2),$$

where $c_3 = (\bar{\lambda}_d^2/32\|k\|_\infty) \wedge (\bar{\lambda}_d/12\|k\|_\infty)$ and $c_4 = (48\|k\|_\infty/\bar{\lambda}_d) \vee 1$ are possible choices.

If the stopping criterion of the algorithm is an error of $\|\widehat{\mathbf{m}}_{\hat{t}} - \mathbf{m}_n\| \leq n^{-1/2}$, i.e. $\hat{t} = \inf_{t \geq 1} \|\widehat{\mathbf{m}}_t - \mathbf{m}_n\| \leq n^{-1/2}$, and if the infimum is greater than n , then $\hat{t} = n$ and $\widehat{\mathbf{m}}_{\hat{t}} = \mathbf{m}_n$, then

$$E(\hat{t}) \leq \lceil 1 + 12\|k\|_{\infty}^{1/2} \log(n)/\delta \rceil + 3n \exp(-(c_3 n^{1/2} - c_4)_+^2).$$

Proof. — The first statement follows directly from Theorem 2.17 and the bound on the error of the CGM. For the other statements let us consider the space \mathcal{H}_S corresponding to the kernel $k_S = k|_S \times S$, where S is the support of the law P , and with corresponding objects $\mathbf{m}_S, \mathbf{m}_{S,n}$ and $\tilde{\mathfrak{C}}_c^S$. As in the proof of Proposition 4.1, we have that

$$\begin{aligned} \Pr(\|\mathbf{m}_{S,n} - \mathbf{m}_S\|_S \geq \delta/4) \\ \leq \exp\left(-\left(\frac{\delta n^{1/2}}{\sqrt{32}\|k\|_{\infty}^{1/2}} - 1\right)^2\right) \leq \exp\left(-\left(\frac{n^{1/2}\bar{\lambda}_d}{12\|k\|_{\infty}} - 1\right)^2\right) \end{aligned}$$

whenever $n \geq (\sqrt{32}\|k\|_{\infty}^{1/2}/\delta)^2$. Furthermore, there is a ball or radius $\delta/2$ around \mathbf{m}_S in C_n (as a subset of the affine span of C_S) with probability at least

$$1 - 2 \exp\left(-\left(\frac{n^{1/2}\bar{\lambda}_d^2}{32\|k\|_{\infty}} - \frac{48\|k\|_{\infty}}{\bar{\lambda}_d}\right)^2\right),$$

for large enough n . Hence, with probability at least

$$1 - 3 \exp(-(c_3 n^{1/2} - c_4)_+^2)$$

there is a ball or radius $\delta/4$ around $\mathbf{m}_{S,n}$ in C_n (as a subset of the affine span of C_S). The second result follows since the CGM reduces the error in each step and the initial error is bounded by $\|\widehat{\mathbf{m}}_1 - \mathbf{m}_n\| \leq 2\|k\|_{\infty}^{1/2}$.

The third statement follows along similar lines. In the event that we have a ball of size $\delta/4$ it follows that $\|\widehat{\mathbf{m}}_t - \mathbf{m}_n\| \leq \exp(-\delta(t-1)/24\|k\|_{\infty}^{1/2})$. Setting the right hand side to $n^{-1/2}$ leads to $\hat{t} \leq \lceil 1 + 12\|k\|_{\infty}^{1/2} \log(n)/\delta \rceil$. \square

4.2.1. Compression for kernel regression

We can also apply Algorithm 3 to compress the data for kernel regression. The only thing that we need to do is to use the kernel function $\tau((y, x), (y', x')) = (\kappa_y + \rho)((y, x), (y', x')) = \kappa(x, x') + yy'k(x, x')$ that we used in Section 3.4.3 and where k is some kernel function on the space \mathcal{X} , and to cap the response variables Y . We state the corresponding result for the compression of the mean element in high probability below. One can obviously also derive bounds on the deviation in expectation and the expected number of points in the core-set.

PROPOSITION 4.3. — *Let $(\mathcal{X} \times \mathbb{R}, \mathfrak{T}, \mathcal{A}, P)$ be a topological measure space such that P is a Radon probability measure which has support S . Let k be continuous bounded kernel function defined on \mathcal{X} such that the corresponding RKHS is finite dimensional and does not contain the constant functions. Let $(X_1, Y_1), \dots, (X_n, Y_n)$ be i.i.d. random variables with law P , and assume that $Y_i = f_0(X_i) + \epsilon_i$, for all $i \leq n$,*

where f_0 is a measurable and bounded function and $\epsilon_1, \dots, \epsilon_n$ are i.i.d. sub-Gaussian random variables with variance $0 < \sigma^2$ which are independent of X_1, \dots, X_n . Let $\bar{\lambda}_{\star, \tilde{S}}$ be the smallest eigenvalue of the covariance operator $\tilde{\mathfrak{C}}_k^{\tilde{S}}$ corresponding to the kernel function $k|_{\tilde{S} \times \tilde{S}}$, $\tilde{S} = \overline{\{x : (x, y) \in S\}}$. Chose $q \in (0, 1)$ and define the sequence $\{r_n\}_{n \geq 1}$ in the following way:

$$r_1 = 1 \vee 2\sigma \log^{1/2} \left(\frac{22\|k\|_\infty(\sigma + \|f_0\|_\infty + \|k\|_\infty^{1/2})}{\sigma^2 \bar{\lambda}_{\star, \tilde{S}}} \right)$$

and for $n \geq 2$, let

$$r_n = r_1 \vee \sqrt{2}\sigma \vee \sqrt{2}\sigma \log^{1/2} (16n^{1/2}\sigma^2\|k\|_\infty/q).$$

Define $\widehat{Y}^{(n)} = (Y \wedge (r_n + \|f_0\|_\infty)) \vee -(r_n + \|f_0\|_\infty)$, let $\widehat{\mathbf{m}}_{\tau, n}$ be the empirical mean element corresponding to the kernel τ and the data $(X_1, \widehat{Y}_1^{(n)}), \dots, (X_n, \widehat{Y}_n^{(n)})$ and let $\widehat{\mathbf{m}}_{\tau, t}$ be the output of the algorithm when applied to the capped data and $\widehat{\mathbf{m}}_{\tau, n}$. Let $\beta = 48\|\tau\|_{S_{f, n, \infty}}/\sigma^2 \bar{\lambda}_{\star, \tilde{S}}$ then with probability $1 - q$

$$\left\| \widehat{\mathbf{m}}_{\tau, \lceil \beta \log(4n\|\tau\|_{S_{f, n, \infty}}) \rceil + 1} - \mathbf{m}_{\tau, n} \right\| \leq 2n^{-1/2}$$

whenever n is (strictly) greater than

$$\left(\frac{16\|\tau\|_{S_{f, n, \infty}} \left(\sqrt{2 \log(12/q)} + 384\|\tau\|_{S_{f, n, \infty}}/\sigma^2 \bar{\lambda}_{\star, \tilde{S}} \right)}{\sigma^4 \bar{\lambda}_{\star, \tilde{S}}^2} \right)^2 \vee \left(\frac{64\|\tau\|_{S_{f, n, \infty}} + 68\|\tau\|_{S_{f, n, \infty}}^{1/2} \log^{1/2}(4/q)}{\sigma^2 \bar{\lambda}_{\star, \tilde{S}}} \right)^2.$$

Proof. — The statement follows from Proposition 3.6. In particular, under the states conditions and with probability at least $1 - q$, simultaneously $\|\widehat{\mathbf{m}}_{\tau, n} - \mathbf{m}_{\tau, n}\|_\tau \leq n^{-1/2}$ and

$$\left\| \widehat{\mathbf{m}}_{\tau, t} - \widehat{\mathbf{m}}_{\tau, n} \right\|_\tau \leq 2\|\tau\|_{S_{f, n, \infty}}^{1/2} \exp \left(-\frac{\widehat{\delta}^{(n)}}{4} \frac{t-1}{6\|\tau\|_{S_{f, n, \infty}}^{1/2}} \right).$$

Setting the right side of the last equation equal to $n^{-1/2}$ yields

$$t - 1 = \left\lceil \frac{12\|\tau\|_{S_{f, n, \infty}}^{1/2} \log(4n\|\tau\|_{S_{f, n, \infty}})}{\widehat{\delta}^{(n)}} \right\rceil.$$

Replacing $\widehat{\delta}^{(n)}$ by its lower bound $\sigma^2 \bar{\lambda}_{\star, \tilde{S}}/4\|\tau\|_{S_{f, n, \infty}}^{1/2}$ gives the constant β stated in the Proposition 4.3. \square

Remark 4.4. — The eigenvalue in the definition of the sequence r_n can be replaced in that definition by a lower bound on this eigenvalue. Similarly, the term $\|f_0\|_\infty$ in the definition of \widehat{Y} can be replaced by an upper bound. We also used here the lower bound $\sigma^2 \bar{\lambda}_{\star, \tilde{S}}$ on $\widehat{\lambda}_\star^{(n)}$ instead of using $\widehat{\lambda}_\star^{(n)}$ directly. This affects, in particular, the number n from which point onward the compression results apply.

4.2.2. Avoiding the explicit computation of s

Mitigating the cost of computing s is more difficult when the CGM is used. The main problem is that we are aiming for a run-time of $O(n \log(n))$ and there is not much leeway in each iteration. For instance, if, like for kernel herding, we split the data into \sqrt{n} batches of size \sqrt{n} then we have an overall run-time of $\sqrt{n} \times (\sqrt{n})^2 = n^{3/2}$ because computing s per batch incurs a quadratic cost in the sample size. One way to reduce that computational cost is to make the quadratic term smaller but then we have many batches. For example, if we aim for a $\log(n)$ batch size then we have $n/\log(n)$ many batches and the reduction in sample size is minuscule. In particular, we could not just run the CGM directly on the $n/\log(n)$ many approximations since that would result in an $n^2/\log^2(n)$ run-time cost. One way around this problem is to apply the process iteratively: in the first iteration use about $n/\log(n) =: T_1$ many batches and compute $\widehat{\mathbf{m}}_{11}, \dots, \widehat{\mathbf{m}}_{1T_1}$. This can be done in about $(n/\log(n)) \times \log^2(n) = n \log(n)$ time, resulting in approximations that consist of $\log \log(n)$ many elements each. If we want to allow a run-time of $O(n \log(n))$ per iteration then in the second iteration we can use $T_2 := T_1/\log^2(n)$ many batches since $(T_1/\log^2(n)) \times \log^4(n) = n \log(n)$ (ignoring the $\log \log(n)$ terms). Continuing this process, at iteration 3, we have $T_3 = T_2/\log^4(n)$ many batches, and, more generally, for $i \geq 2$, we have $T_i = T_{i-1}/\log^{2^{i-1}}(n)$ many batches. We can stop the iterations when we are down to \sqrt{n} many batches since we can apply the CGM then directly. To get down to \sqrt{n} many batches we need about

$$\ell = \log\left(\frac{\log(n)}{\log \log(n)}\right) \approx \log \log(n)$$

many iterations since

$$\sqrt{n} \approx T_i = \frac{n}{\prod_{i=1}^{\ell} \log^{2^{i-1}}(n)} = \frac{n}{\log^{2^{\ell}}(n)}.$$

This then implies an overall run-time of this algorithm of about $O(n \log(n) \log \log(n))$.

A major concern with this algorithm is that we have many optimization problems that have to be solved simultaneously and we need to be lucky in each case to have a ball of sufficient size around the corresponding \mathbf{m}_n in C_n . It seems rather unlikely that we can guarantee for each of these optimization problems the existence of such a ball. A better way to approach this compression problem might be to work instead with fixed error bounds that have to be achieved in each optimization problem. The hope with this approach is that we can then guarantee a sufficient compression but the number of sample points needed might be larger than $\log(n)$. Algorithm 4 implements this idea.

In the algorithm $\mathbf{m}_{1,1}, \dots, \mathbf{m}_{1,T_1}$ denote the mean elements corresponding to the initial T_1 batches. For the analysis of the algorithm it is useful to also have the mean elements corresponding to all the samples entering into the j 'th batch in iteration i ; denote this element by $\mathbf{m}_{i,j}$. The idea of the algorithm is to approximate $\mathbf{m}_{i,j}$ in iteration i and batch j . Working directly with $\mathbf{m}_{i,j}$ is not possible if we try to stay around $n \log(n)$ computation time per iteration since $\mathbf{m}_{i,j}$ will consist eventually of about \sqrt{n} many samples in each batch which implies a cost of n per batch. Therefore,

we approximate $\mathbf{m}_{i,j}$ first by $\tilde{\mathbf{m}}_{i,j}$ which will consist, under suitable conditions, of far fewer sample points. The approximation $\tilde{\mathbf{m}}_{i,j}$ is then further compressed into $\hat{\mathbf{m}}_{i,j}$ which consists of even fewer sample points. The variables $M_{i,j}$ keep track over how many sample points $\mathbf{m}_{i,j}$ is averaged. Hence, $\mathbf{m}_{i,j} = (1/M_{i,j}) \sum_{u \in \mathcal{I}_{i,j}} M_{i-1,u} \mathbf{m}_{i-1,u}$ for all $j \leq T_i$ and $2 \leq i \leq \ell$.

Algorithm 4 (A compression algorithm for \mathbf{m}_n that uses the CGM and avoids the explicit computation of s .)

Input: sample X_1, \dots, X_n , kernel k .

In the follow let $\ell = \left\lceil \frac{1}{\log(2)} \log\left(\frac{\log(n)}{\log \log(n)}\right) - 1 \right\rceil$.

Split the sample into $T_1 := \lceil n/\log(n) \rceil$ batches.

Let $\mathcal{I}_{1,1}, \dots, \mathcal{I}_{1,T_1}$ be the corresponding indices of the sample points.

Apply the CGM to each batch to approximate $\mathbf{m}_{1,1}, \dots, \mathbf{m}_{1,T_1}$ by using sample points indexed by $\mathcal{I}_{1,1}, \dots, \mathcal{I}_{1,T_1}$ until the error of all approximations $j \leq T_1$ is below $\varepsilon_{1,j} = |\mathcal{I}_{1,j}|^{-1/2}$.

Store the approximations in $\hat{\mathbf{m}}_{1,1}, \dots, \hat{\mathbf{m}}_{1,T_1}$ and let

$$M_{1,1} = |\mathcal{I}_{1,1}|, \dots, M_{1,T_1} = |\mathcal{I}_{1,T_1}|.$$

Iterate through $i = 2, \dots, \ell$:

Split the approximations $\hat{\mathbf{m}}_{i-1,1}, \dots, \hat{\mathbf{m}}_{i-1,T_{i-1}}$ into $T_i := \lceil T_{i-1}/\log^{2^{i-1}}(n) \rceil$ batches.

Let $\mathcal{I}_{i,1}, \dots, \mathcal{I}_{i,T_i}$ be the corresponding indices and for all $j \leq T_i$ let

$$M_{i,j} = \sum_{u \in \mathcal{I}_{i,j}} M_{i-1,u}.$$

For each batch $j \leq T_i$ average the old approximations

$$\tilde{\mathbf{m}}_{i,j} := \frac{1}{M_{i,j}} \sum_{u \in \mathcal{I}_{i,j}} M_{i-1,u} \hat{\mathbf{m}}_{i-1,u}.$$

Apply the CGM to each batch $j \leq T_i$, approximating $\tilde{\mathbf{m}}_{i,j}$ by convex combinations of the elements $\hat{\mathbf{m}}_{i-1,u}$, $u \in \mathcal{I}_{i,j}$, with an error of at most

$$\varepsilon_{i,j} = M_{i,j}^{-1/2}.$$

Store the approximations in $\hat{\mathbf{m}}_{i,1}, \dots, \hat{\mathbf{m}}_{i,T_i}$.

Apply the CGM a final time to $\hat{\mathbf{m}}_{\ell,1}, \dots, \hat{\mathbf{m}}_{\ell,T_\ell}$ to compress $\frac{1}{T_\ell} \sum_{j=1}^{T_\ell} \hat{\mathbf{m}}_{\ell,j}$

with an approximation error of at most $n^{-1/2}$ and return the approximation.

We left out a few details in the algorithm. In particular, the usual vector s that consists of inner products between $k(X_i, \cdot)$ and \mathbf{m}_n has to be replaced by vectors with entries of the form $\langle \widehat{\mathbf{m}}_{i-1,u}, \widetilde{\mathbf{m}}_{i,j} \rangle$, $u \in \mathcal{I}_{i,j}$, when $i \geq 2$. The element $\widetilde{\mathbf{m}}_{i,j}$ corresponds to an average over the $\widehat{\mathbf{m}}_{i-1,u}$ terms and there are $|\mathcal{I}_{i,j}|$ many terms over which this average is taken. The quantity $|\mathcal{I}_{i,j}|$ is not of major concern when bounding the computational complexity. The computational complexity of calculating these inner product vectors is rather dominated by how many points are contained in the approximations $\widehat{\mathbf{m}}_{i,j}$. Another point worth noting is that the final approximation will ideally be given in terms of convex combinations of the original sample points $k(X_1, \cdot), k(X_2, \cdot), \dots$. Roughly speaking, this convex combination can be computed by multiplying the weights in the different iterations. Finally, observe that we can keep track of how well $\widetilde{\mathbf{m}}_{i,j}$ is approximated if $\widehat{\mathbf{m}}_{i,j}$ does not consist of too many points since the $|\mathcal{I}_{i,j}|$ are chosen small enough that we can compute and store the corresponding kernel matrices

$$\left(\langle \widehat{\mathbf{m}}_{i-1,u}, \widehat{\mathbf{m}}_{i-1,v} \rangle \right)_{u,v \in \mathcal{I}_{i,j}}$$

and from these kernel matrices we can compute the approximation errors.

Bounding the size of the set which is used in the resulting approximation in high probability or expectation is a major challenge that we will not address here. However, it is easier to say something about the resulting *approximation error* by refining the analysis of the kernel herding algorithm: the philosophy of the algorithm is to guarantee in high probability in each iteration that \mathbf{m}_n is approximated with an error of $n^{1/2}$. In detail, observe that for any $1 \leq i \leq \ell$, $\mathbf{m}_n = (1/n) \sum_{j=1}^{T_i} M_{i,j} \mathbf{m}_{i,j}$, where we use that $\sum_{j=1}^{T_i} M_{i,j} = n$. We can use the link between $\widehat{\mathbf{m}}_{i,j}$ and $\mathbf{m}_{i,j}$ to measure in each iteration the error when approximating \mathbf{m}_n by $(1/n) \sum_{j=1}^{T_i} M_{i,j} \widehat{\mathbf{m}}_{i,j}$. The naive approach of using the triangular inequality does not lead to useful results since

$$\left\| \mathbf{m}_n - (1/n) \sum_{j=1}^{T_i} M_{i,j} \widehat{\mathbf{m}}_{i,j} \right\| \leq \frac{1}{n} \sum_{j=1}^{T_i} M_{i,j} \|\mathbf{m}_{i,j} - \widehat{\mathbf{m}}_{i,j}\| \leq \varepsilon_i$$

and we would need to set ε_i to $n^{-1/2}$ to guarantee a low enough approximation error. But aiming in each batch for an error of $n^{-1/2}$ when only $\log(n)$ sample points are in each batch is not useful. As for the kernel herding analysis, a better approach might be to consider the variance of the error and to make use of the independence of the sample points. Let us first look at the case $i = 1$,

$$\begin{aligned} E \left(\left\| \mathbf{m}_n - \frac{1}{n} \sum_{j=1}^{T_i} M_{i,j} \widehat{\mathbf{m}}_{i,j} \right\|^2 \right) &= \frac{1}{n^2} E \left(\left\| \sum_{j=1}^{T_i} M_{i,j} (\mathbf{m}_{i,j} - \widehat{\mathbf{m}}_{i,j}) \right\|^2 \right) \\ &= \frac{1}{n^2} \sum_{j=1}^{T_i} E \left(\|M_{i,j} (\mathbf{m}_{i,j} - \widehat{\mathbf{m}}_{i,j})\|^2 \right) \end{aligned}$$

$$\begin{aligned}
 & + \frac{1}{n^2} \sum_{j_1=1}^{T_i} \sum_{j_2 \neq j_1}^{T_i} M_{i,j_1} M_{i,j_2} E(\langle \mathbf{m}_{i,j_1} - \widehat{\mathbf{m}}_{i,j_1}, \mathbf{m}_{i,j_2} - \widehat{\mathbf{m}}_{i,j_2} \rangle) \\
 & = \frac{1}{n^2} \sum_{j=1}^{T_i} M_{i,j}^2 E(\| \mathbf{m}_{i,j} - \widehat{\mathbf{m}}_{i,j} \|^2) \\
 & + \frac{1}{n^2} \sum_{j_1=1}^{T_i} \sum_{j_2 \neq j_1}^{T_i} M_{i,j_1} M_{i,j_2} \langle \mathbf{m} - E(\widehat{\mathbf{m}}_{i,j_1}), \mathbf{m} - E(\widehat{\mathbf{m}}_{i,j_2}) \rangle.
 \end{aligned}$$

As for the kernel herding algorithm we can control the bias term in a crude manner by using that

$$\begin{aligned}
 |\langle \mathbf{m} - E(\widehat{\mathbf{m}}_{i,j_1}), \mathbf{m} - E(\widehat{\mathbf{m}}_{i,j_2}) \rangle| & \leq \max_{j \leq T_i} \| \mathbf{m} - E(\widehat{\mathbf{m}}_{i,j}) \|^2 \\
 & \leq \max_{j \leq T_i} E(\| \mathbf{m}_{i,j} - \widehat{\mathbf{m}}_{i,j} \|^2) \leq \varepsilon_i^2.
 \end{aligned}$$

However, this is not leading to an improvement since in the first iteration

$$E\left(\left\| \mathbf{m}_n - \frac{1}{n} \sum_{j=1}^{T_1} M_{1,j} \widehat{\mathbf{m}}_{1,j} \right\|^2 \right) \leq \frac{\varepsilon_1^2 \max_{j' \leq T_1} M_{1,j'}^2 (T_1 + T_1^2)}{n^2} \approx \varepsilon_1^2 (1 + 1/n)$$

and

$$\Pr\left(\left\| \mathbf{m}_n - \frac{1}{n} \sum_{j=1}^{T_1} M_{1,j} \widehat{\mathbf{m}}_{1,j} \right\| \geq n^{-1/2} \right) \leq \varepsilon_1^2 (n + 1).$$

implies that ε_1 would have to be of order $n^{-1/2}$. A central question at this point is *of what order is the bias term*. In particular, is the upper bound of ε_1^2 for the squared bias term overly pessimistic? A natural threshold for the error in each batch is $\log^{-1/2}(n)$ in the first iteration since there are about $\log(n)$ many samples in each batch. For $\log^{-1/2}(n)$ to be sufficiently low we need a bound on the bias term of about $cn^{-1/2}$, $c \in (0, 1)$, since then

$$\Pr\left(\left\| \mathbf{m}_n - \frac{1}{n} \sum_{j=1}^{T_1} M_{1,j} \widehat{\mathbf{m}}_{1,j} \right\| \geq n^{-1/2} \right) \leq c + \varepsilon_1^2 \log(n) \leq 2c,$$

when a threshold of $c^{1/2} \log^{-1/2}(n)$ is used in the optimization. In other words, the bias term has to fall exponentially fast to allow for a threshold that is proportional to the sample size, i.e. the bias has to be below $\exp(-m/2)$, where $m = \log(n)$ is the sample size in each batch in the first iteration.

The error in the successive approximations can be treated in a similar way and since there are only about $\log \log(n)$ many iterations a simple union bound argument suffices to control the error simultaneously over all iterations. To demonstrate how

the error evolves consider $i = 2$, then

$$E \left(\left\| \mathbf{m}_n - \frac{1}{n} \sum_{j=1}^{T_2} M_{2,j} \widehat{\mathbf{m}}_{2,j} \right\|^2 \right) \leq 3 \left(E \left(\left\| \mathbf{m}_n - \frac{1}{n} \sum_{j=1}^{T_1} M_{1,j} \widehat{\mathbf{m}}_{1,j} \right\|^2 \right) + E \left(\left\| \frac{1}{n} \sum_{j=1}^{T_1} M_{1,j} \widehat{\mathbf{m}}_{1,j} - \frac{1}{n} \sum_{j=1}^{T_2} M_{2,j} \widehat{\mathbf{m}}_{2,j} \right\|^2 \right) \right)$$

and

$$n^2 E \left(\left\| \frac{1}{n} \sum_{j=1}^{T_1} M_{1,j} \widehat{\mathbf{m}}_{1,j} - \frac{1}{n} \sum_{j=1}^{T_2} M_{2,j} \widehat{\mathbf{m}}_{2,j} \right\|^2 \right) = E \left(\left\| \sum_{j=1}^{T_2} (M_{2,j} \widetilde{\mathbf{m}}_{2,j} - M_{2,j} \widehat{\mathbf{m}}_{2,j}) \right\|^2 \right) = \sum_{j=1}^{T_2} M_{2,j}^2 E(\|\widetilde{\mathbf{m}}_{2,j} - \widehat{\mathbf{m}}_{2,j}\|^2) + 2 \sum_{j=1}^{T_2} \sum_{u \neq j} M_{i,j} M_{i,u} \langle E(\widetilde{\mathbf{m}}_{2,j} - \widehat{\mathbf{m}}_{2,j}), E(\widetilde{\mathbf{m}}_{2,u} - \widehat{\mathbf{m}}_{2,u}) \rangle$$

where we can move the expectation inside the inner product since $\widetilde{\mathbf{m}}_{2,j}$ and $\widehat{\mathbf{m}}_{2,j}$ are independent of $\widetilde{\mathbf{m}}_{2,u}$ and $\widehat{\mathbf{m}}_{2,u}$. The bias term that is now important is $\|E(\widetilde{\mathbf{m}}_{2,j}) - E(\widehat{\mathbf{m}}_{2,j})\|$ and we need a similar fast decay of the bias as for $i = 1$. The other term is easier to deal with,

$$\frac{1}{n^2} \sum_{j=1}^{T_2} M_{2,j}^2 E(\|\widetilde{\mathbf{m}}_{2,j} - \widehat{\mathbf{m}}_{2,j}\|^2) \leq \frac{1}{n^2} \sum_{j=1}^{T_2} M_{2,j} = \frac{1}{n}$$

by the choice of $\varepsilon_{i,j}$ in the algorithm.

There are a few open problems concerning this algorithm, and variations thereof. The algorithm is set up to enforce tighter and tighter error bounds in each iteration, i.e. the error threshold changes approximately from $\log^{-1/2}(n)$ in the first iteration to $\log^{-3/2}(n)$ in the second iteration and $\log^{-7/2}(n)$ in the third iteration. The hope is that good approximations in the first iteration allow us to get even better approximations in the second round and so forth. But it is by no means obvious that this intuition is correct and in all likelihood these choices are not optimal.

The next major obstacles in controlling the error of the algorithm are obviously the bias terms. If there is an exponential decrease in the bias then we are in a very fortunate situation and can control the approximation error. If the bias term decreases slower then it might be worth to consider alternatives of the CGM which incorporate bias reduction techniques and are not focusing solely on the approximation error.

The biggest challenge when studying this algorithm is in all likelihood the problem of controlling the size of the ball around the various elements $\mathbf{m}_{i,j}$ simultaneously over all iterations and batches. In fact, a uniform bound might even be suboptimal for analyzing the performance of the algorithm since small ball sizes can be compensated for by batches that have a larger ball around their corresponding $\mathbf{m}_{i,j}$ and which need less sample points than suggested by a worst case bound. In other words, we might need to control the fluctuations or the distribution of the ball sizes.

5. Applications

In the following, we look at how our compression techniques can be combined with machine learning methods. In particular, we are looking at the two sample problem, at kernel ridge regression and at kernel PCA. These applications fall roughly within the area of kernel mean embeddings (see [MFSS17] for a substantial review of that area). We work with *finite dimensional RKHSs* to be able to make use of our theoretical results. In particular, we will be using Proposition 4.2 and variations thereof. Beside the assumption that the RKHS is finite dimensional, the main assumption of the proposition is that the *kernel function is bounded*. For large n , with high probability, one can compress a dataset of size n to a weighted average of $12\|k\|_\infty \log(n)/\bar{\lambda}_d$ points, where $\bar{\lambda}_d$ is the smallest non-zero eigenvalue of the associated covariance operator.

We are interested in understanding how competitive the CGM based approaches are in terms of rates of compression and run-time. There are two caveats that we want to discuss upfront:

- (1) the assumption that the RKHS is finite dimensional is closely related to the scaling of the compression as a function of $1/\bar{\lambda}_d$. If $\bar{\lambda}_d$ is small compared to n then the compression quality can suffer. It is likely that the effect of small $\bar{\lambda}_d$ can be weakened by suitable regularization ideas or other approaches. But developing these approaches will require further research.
- (2) The upfront cost for running the CGM algorithm that underlies Proposition 4.2 is $O(n^2)$. Without this upfront cost the algorithm would have a run-time in the order of $O(n \log(n))$. We discussed various ideas of how to reduce this upfront cost in Section 4, but, again, significant further research into such ideas is needed.

Since it is currently unclear what compression rates can be achieved when avoiding the upfront cost of $O(n^2)$, we formulate the runtime statements as functions of $\psi_{\text{comp}}(n)$ and $\psi_{\text{size}}(n)$, where $O(\psi_{\text{comp}}(n))$ is the computational cost for calculating the compression and $O(\psi_{\text{size}}(n))$ is the order of the number of points that are needed in the compression to guarantee, with high probability, that the compression is no more than $cn^{-1/2}$, $c > 0$, away from the empirical mean embedding. Our current results guarantee that $O(\psi_{\text{comp}}(n)) = n^2$ and $O(\psi_{\text{size}}(n)) = \log(n)$. Generally, the hope is that these can be improved to something in the order of $O(\psi_{\text{comp}}(n)) = n \log^\alpha(n)$ and $O(\psi_{\text{size}}(n)) = \log^\alpha(n)$, for some $\alpha \geq 1$.

5.1. Two Sample Test

The first application we consider is the two sample test problem [GBR⁺12]. In the two sample test problem i.i.d. data X_1, \dots, X_n and Y_1, \dots, Y_m attaining values in \mathcal{X} are given, the X_i 's are furthermore independent from the Y_i 's but it is unknown if the X_i 's have the same distribution as the Y_i 's. The null-hypothesis is that the distributions are equal. One way to build a test statistic for this hypothesis testing problem is to consider $\|\mathbf{m}_{X,n} - \mathbf{m}_{Y,m}\|$, where k is a kernel function on \mathcal{X} , $k(X, \cdot)$,

$k(Y, \cdot) \in \mathcal{L}^1(P)$, $\mathbf{m}_{X,n} = (1/n) \sum_{i=1}^n k(X_i, \cdot)$ and $\mathbf{m}_{Y,m} = (1/m) \sum_{i=1}^m k(Y_i, \cdot)$. Calculating the norm can be done in $O((n \vee m)^2)$ by using that

$$\|\mathbf{m}_{X,n} - \mathbf{m}_{Y,m}\|^2 = \frac{1}{n^2} \sum_{i,j=1}^n k(X_i, X_j) - \frac{2}{nm} \sum_{i=1}^n \sum_{j=1}^m k(X_i, Y_j) + \frac{1}{m^2} \sum_{i,j=1}^m k(Y_i, Y_j).$$

When using one of the compression approaches this turns into a run-time of the order $O((\psi_{\text{comp}}(n) \vee \psi_{\text{comp}}(m)) \vee (\psi_{\text{size}}(n) \vee \psi_{\text{size}}(m))^2)$. In particular, we can simply replace $\mathbf{m}_{Y,m}$ and $\mathbf{m}_{X,n}$ by their approximations. Furthermore, with high probability, the rate of convergence of $\|\mathbf{m}_{X,n} - \mathbf{m}_{Y,n}\|$ to $\|\mathbf{m}_X - \mathbf{m}_Y\|$, where $\mathbf{m}_X = \int k(X, \cdot) dP$ and $\mathbf{m}_Y = \int k(Y, \cdot) dP$, will be preserved when moving to the compression.

5.2. Kernel ridge regression

Let us consider now the regression problem with data $(X_1, Y_1), \dots, (X_n, Y_n)$, where we assume that the pairs are independent and that the Y_i are bounded. In Section 1.8 we discussed the case of infinite dimensional RKHSs and potential benefits of coresets methods compared to (or in conjunction with) the Nyström method in that setting. In this section, we will focus on the finite dimensional case and compare results for coresets to the bounds for Nyström method [LLW23]. When the CGM is used to approximate \mathbf{c}_y , \mathbf{m}_y and $(1/n) \sum_{i \leq n} Y_i$ simultaneously we get a single index function $\iota : \{1, \dots, l\} \rightarrow \{1, \dots, n\}$ and corresponding approximations $\widehat{\mathbf{c}}_{y,l} = \sum_{i=1}^l w_i \kappa(X_{\iota(i)}, \cdot)$ and $\widehat{\mathbf{m}}_{y,l} = \sum_{i=1}^l w_i \langle Y_{\iota(i)}, \cdot \rangle \otimes k(X_{\iota(i)}, \cdot)$ with strictly positive w_i 's such that $w_1 + \dots + w_l = 1$. The approximation of the least-squares error for a function $h \in \mathcal{H}$ is

$$\sum_{i=1}^l w_i (Y_{\iota(i)} - h(X_{\iota(i)}))^2 = \langle \widehat{\mathbf{c}}_{y,l}, \widetilde{h} \rangle_{\widehat{\mathcal{H} \odot \mathcal{H}}} - 2 \langle \widehat{\mathbf{m}}_{y,l}^{\otimes}, \check{h} \rangle_{\mathbb{R}' \otimes \mathcal{H}} + \sum_{i=1}^l w_i Y_{\iota(i)}^2,$$

where we denote the function $(x, y) \mapsto h^2(x)$ with \widetilde{h} . Due to the representer theorem we can write the solution to the ridge regression problem in the form $h_{\star} = \sum_{i=1}^l \alpha_i k(X_{\iota(i)}, \cdot)$ for suitable $\alpha_i \in \mathbb{R}$. Substituting this into the equation for the least-squares error and ignoring the last term (which is irrelevant for finding the solution) leads to

$$\langle \widehat{\mathbf{c}}_{y,l}, \widetilde{h}_{\star} \rangle_{\widehat{\mathcal{H} \odot \mathcal{H}}} - 2 \sum_{i=1}^l \alpha_i \langle \widehat{\mathbf{m}}_{y,l}^{\otimes}, \widehat{k}(X_{\iota(i)}, \cdot) \rangle_{\mathbb{R}' \otimes \mathcal{H}}.$$

Let C_l be an $l \times l$ matrix with the entry in row i and column j being

$$\begin{aligned} \sum_{u=1}^l w_u \langle k(X_{\iota(u)}, \cdot) \otimes k(X_{\iota(u)}, \cdot), k(X_{\iota(i)}, \cdot) \otimes k(X_{\iota(j)}, \cdot) \rangle_{\otimes} \\ = \sum_{u=1}^l w_u k(X_{\iota(u)}, X_{\iota(i)}) k(X_{\iota(u)}, X_{\iota(j)}). \end{aligned}$$

then

$$\langle \widehat{\mathbf{c}}_{y,l}, \widetilde{h}_{\star} \rangle_{\widehat{\mathcal{H} \odot \mathcal{H}}} = \alpha^{\top} C_l \alpha.$$

Also, let K_l be the kernel matrix for samples $X_{\iota(1)}, \dots, X_{\iota(l)}$ and let m_l be an l -dimensional vector with entry i being

$$\left\langle \widehat{\mathbf{m}}_{y,l}^{\otimes}, \widehat{k}(X_{\iota(i)}, \cdot) \right\rangle_{\mathbb{R}' \otimes \mathcal{H}} = \sum_{u=1}^l w_u Y_{\iota(u)} k(X_{\iota(i)}, X_{\iota(u)}).$$

With these in place the solution of the ridge-regression problem with regularization parameter $\lambda > 0$ is found by minimizing

$$\alpha^\top C_l \alpha - 2\alpha^\top m_l + \lambda \alpha^\top K_l \alpha$$

with respect to $\alpha \in \mathbb{R}^l$. Taking the gradient with respect to α and setting it to zero yields

$$(5.1) \quad 2C_l \alpha - 2m_l + 2\lambda K_l \alpha = 0.$$

Observe that $C_l = K_l W K_l$, where W is a diagonal matrix with $W_{uu} = w_u$ for all $u \leq l$. Similarly, $m_l = K_l W y = K_l W (K_l W)^\dagger K_l W y$, where y is a vector with entries $y_u = Y_{\iota(u)}$ for all $u \leq l$. Hence, we can rewrite (5.1) as

$$K_l ((W K_l + \lambda I_l) \alpha - W (K_l W)^\dagger K_l W y) = 0,$$

where I_l is the $l \times l$ identity matrix. Since W has strictly positive entries on the diagonal we can rewrite this as

$$(5.2) \quad K_l W ((K_l + \lambda W^{-1}) \alpha - (K_l W)^\dagger K_l W y) = 0,$$

for which a solution is given by

$$(5.3) \quad \alpha = (K_l + \lambda W^{-1})^{-1} (K_l W)^\dagger K_l W y.$$

The inverse is well defined because K_l is p.s.d. and W^{-1} is (strictly) positive definite; the sum of a p.s.d. and strictly positive definite matrix is strictly positive definite and, therefore, has an inverse. Also $(K_l W)^\dagger = (K_l W)^{-1}$ whenever K_l is of full rank and in this case

$$(5.4) \quad \alpha = (K_l + \lambda W^{-1})^{-1} y.$$

This α is also a solution to (5.2) in the general case when K_l is not full rank since $(K_l W)^\dagger K_l W y$ can be replaced by y in this equation.

In terms of the runtime, if we use an $O(n^3)$ algorithm for deriving the inverse then, after compression, the runtime is $O(\psi_{\text{comp}}(n) \vee (\psi_{\text{size}}(n))^3)$. For example, if we work with a finite dimensional RKHS and use the standard CGM then we attain a runtime of $O(n^2)$. Beside the reduction in runtime the storage demand also goes down since only a matrix of size $\psi_{\text{size}}(n) \times \psi_{\text{size}}(n)$ has to be stored for calculating α , and this can be as small as $\log^2(n)/(\bar{\lambda}_d)^2$. The CGM itself needs memory in the order of $O(n)$. In this setting the Nyström method attains a reduction to \sqrt{n} many sample points [LLW23, Theorem 3.9] instead of $\log(n)/\bar{\lambda}_d$ that the coreset approach attains. The smallest non-zero eigenvalue $\bar{\lambda}_d$ of the covariance operator is the key factor that controls which methods attains a better compression.

5.3. Kernel PCA

Kernel PCA [SSM97, BBZ07] is another popular kernel method. For kernel PCA it is common to use the empirical covariance operator as a surrogate. It is worth mentioning that this plug-in approach has a large bias when working in infinite dimensional RKHSs and does not achieve the minimax optimal rate of convergence [Kol18]. However, in finite dimensional RKHSs this is not of major concern and we can use the eigenfunction of $\tilde{\mathfrak{C}}_n \in L(\mathcal{H}, \mathcal{H})$ as an estimate of the eigenfunctions of the covariance operator. In this context, we want to approximate $\tilde{\mathfrak{C}}_n$, which is given by

$$\frac{1}{n} \sum_{i=1}^n k(X_i, \cdot) \widehat{\otimes} k(X_i, \cdot),$$

by using the CGM. We like to compare this approach to the well established approach of using the Nyström method. When using the Nyström method a reconstruction error of order $1/\sqrt{n}$ can be achieved with \sqrt{n} samples [SSRR20, Corollary 3].

As discussed in Section 3.2, we can apply the CGM to the RKHS with the kernel function $\kappa(x, y) = k^2(x, y)$ to approximate \mathfrak{C}_n with some convex combination $\widehat{\mathfrak{C}}_t = \sum_{i=1}^t \alpha_i \kappa(X_{\iota(i)}, \cdot)$, where $\alpha_i \geq 0$ for all $i \leq t$, $\alpha_1 + \dots + \alpha_t = 1$, and ι is some selection of data points. The element \mathfrak{C}_n is closely related to the operator $\tilde{\mathfrak{C}}_n$ and a natural approximation of $\tilde{\mathfrak{C}}_n$ is $\widehat{\mathfrak{C}}_t = \sum_{i=1}^t \alpha_i k(X_{\iota(i)}, \cdot) \widehat{\otimes} k(X_{\iota(i)}, \cdot)$. Note that for any $f, g \in \mathcal{H}$,

$$\langle \widehat{\mathfrak{C}}_t, (f \otimes g) \circ \psi \rangle_{\odot} = \sum_{i=1}^t \alpha_i f(X_{\iota(i)}) g(X_{\iota(i)}) = \langle \widehat{\mathfrak{C}}_t f, g \rangle,$$

where $\psi : \mathcal{X} \rightarrow \mathcal{X} \times \mathcal{X}$, $\psi(x) = (x, x)$. The operator $\widehat{\mathfrak{C}}_t$ is clearly symmetric and, hence, self-adjoint since the RKHS is finite dimensional. Furthermore, all eigenvalues are non-negative since if $e \in \mathcal{H}$ is an eigenfunction of $\widehat{\mathfrak{C}}_t$ then

$$\langle \widehat{\mathfrak{C}}_t e, e \rangle = \sum_{i=1}^t \alpha_i e^2(X_{\iota(i)}) \geq 0.$$

The main question is now if we can quantify the difference between eigenfunctions of $\widehat{\mathfrak{C}}_t$ and $\tilde{\mathfrak{C}}_n$. Let us assume that there are no multiple eigenvalues and that $\widehat{\lambda}_1 \geq \dots \geq \widehat{\lambda}_d > 0$ are the eigenvalues of $\widehat{\mathfrak{C}}_t$ and e_1, \dots, e_d are the corresponding eigenfunctions. Similarly, let $\mu_1, \dots, \mu_d > 0$ be the eigenvalues of $\tilde{\mathfrak{C}}_n$ and f_1, \dots, f_d the corresponding eigenfunctions. Furthermore, assume that the CGM is run until

$$\|\widehat{\mathfrak{C}}_t - \mathfrak{C}_n\|_{\odot} \leq \epsilon.$$

Since

$$\|\widehat{\mathfrak{C}}_t - \tilde{\mathfrak{C}}_n\|_{\text{op}} = \sup_{\|h\|=1} \|\widehat{\mathfrak{C}}_t h - \tilde{\mathfrak{C}}_n h\| = \sup_{\|h\|=1} \sup_{\|g\|=1} \langle (\widehat{\mathfrak{C}}_t - \tilde{\mathfrak{C}}_n)(h), g \rangle = \langle \widehat{\mathfrak{C}}_t - \mathfrak{C}_n, (h \otimes g) \circ \psi \rangle_{\odot}$$

and $\|(h \otimes g) \circ \psi\|_{\otimes} \leq \|h \otimes g\|_{\otimes} = \|h\| \|g\| \leq 1$ it follows from the Cauchy–Schwarz inequality that

$$\|\widehat{\mathfrak{C}}_t - \tilde{\mathfrak{C}}_n\|_{\text{op}} \leq \|\widehat{\mathfrak{C}}_t - \mathfrak{C}_n\|_{\odot} \leq \epsilon.$$

From this bound on the operator norm it follows right away that $\|(\widehat{\mathfrak{C}}_t - \widetilde{\mathfrak{C}}_n)(e_i)\|$ and $\|(\widehat{\mathfrak{C}}_t - \widetilde{\mathfrak{C}}_n)(f_i)\|$ are less than ϵ for all $i \leq d$. In particular,

$$|\lambda_i - \langle e_i, \widetilde{\mathfrak{C}}_n e_i \rangle| \leq \epsilon \quad \text{and} \quad |\langle f_i, \widehat{\mathfrak{C}}_t f_i \rangle - \mu_i| \leq \epsilon$$

for all $i \leq d$. In particular,

$$\lambda_1 \leq \langle e_1, \widetilde{\mathfrak{C}}_n e_1 \rangle + \epsilon \leq \sup_{\|h\|=1} \|\widetilde{\mathfrak{C}}_n h\| \epsilon \leq \mu_1 + \epsilon.$$

By symmetry of the argument it follows that $|\lambda_1 - \mu_1| \leq \epsilon$. The difference between e_1 and f_1 can now also be controlled: let $a_1, \dots, a_d \in \mathbb{R}$ be such that $e_1 = a_1 f_1 + \dots + a_d f_d$ then $1 = \|e_1\|^2 = a_1^2 + \dots + a_d^2$,

$$\mu_1 - 2\epsilon \leq \langle e_1, \widetilde{\mathfrak{C}}_n e_1 \rangle = \sum_{i=1}^d a_i^2 \mu_i$$

and from $(1 - a_1^2)\mu_1 - 2\epsilon \leq (1 - a_1^2)\mu_2$ we can infer that for sufficiently small $\epsilon > 0$,

$$a_1^2 \geq 1 - \frac{2\epsilon}{\mu_1 - \mu_2} \quad \text{and} \quad \|e_1 - f_1\|^2 = 2 - 2a_1 \leq 2 - 2\left(1 - \frac{2\epsilon}{\mu_1 - \mu_2}\right)^{1/2}.$$

The other eigenfunctions can be treated in a similar way by moving to the subspaces that are orthogonal to the already covered eigenfunctions e_1, \dots, e_l , $l \leq d$. Notice the factor $\mu_1 - \mu_2$. Such eigengaps are common when estimating eigenfunctions or projection operators and are closely linked to the *Davis–Kahan theorem* which controls the stability of eigenspaces under perturbations. In particular, such factors appear when controlling the difference between the empirical and true projection operator in the operator norm. In kernel PCA it is common to quantify the difference with a weaker norm and, therefore, results as in [SSRR20] do not depend on the eigengap. The eigengap aside, we can note that the approximation error falls at a rate of $\epsilon^{1/4}$, where ϵ is controlled by our compression algorithms. In particular, Proposition 4.2 is applicable if we use the kernel $\kappa(x, y) = k^2(x, y)$. To guarantee an approximation error of

$$\|e_i - f_i\| \leq \sqrt{2} \left(1 - \left(1 - \frac{2\epsilon}{\mu_i - \mu_{i+1}}\right)^{1/2}\right)^{1/2} \lesssim \frac{\epsilon^{1/4}}{(\mu_i - \mu_{i+1})^{1/4}} \lesssim n^{-1/2}$$

we need an ϵ in the order of n^{-2} . In expectation we achieve an error of order n^{-2} whenever

$$n^{-2} \approx \exp(-t/\bar{\lambda}_{d,\kappa}),$$

where $\bar{\lambda}_{d,\kappa}$ denotes the smallest non-zero eigenvalue of the covariance operator that corresponds to the kernel κ . In other words, we need a coresnet of size $t \approx 2 \log(n)/\bar{\lambda}_{d,\kappa}$. This is potentially favorable to the \sqrt{n} size that Nyström methods need when $\bar{\lambda}_{d,\kappa}$ is not too small. A crude bound on the computational complexity of an eigendecomposition is $O(n^3)$. If we use this crude bound then we get the following runtime bound for kernel PCA when using coresets, $O(\psi_{\text{comp}}(n) \vee (\psi_{\text{size}}(n))^3)$.

6. Example: Slow rate of convergence in infinite dimensions

The last section of this paper is dedicated to the construction of the example for which the kernel herding algorithm performs strictly worse than in finite dimensions when the density function of the data distribution has a density that is bounded away from zero. The corresponding theorem is the following.

THEOREM 6.1. — *There exists an initialization, a continuous kernel, and a Borel probability measure on $[0, 1]$ which assigns non-zero probability to open intervals for which the kernel herding algorithm does not converge fast, i.e. there exists no constant b such that $\|\mathbf{m}_t - \mathbf{m}\| \leq b/t$ for all $t \geq 1$.*

The proof of Proposition 6.1 is split into two parts. In the first part, we construct a Hilbert space, a map $\phi : [0, 1] \rightarrow \mathcal{H}$, and an element $\mathbf{m} \in \mathcal{H}$ such that the algorithm does not converge fast. We then use this Hilbert space to construct an RKHS for which the algorithm behaves in exactly the same way as when acting on the Hilbert space, and, consequently, the algorithm does not converge fast when applied to the RKHS.

The construction idea

Before getting into the technical details we like to outline the basic intuition of the construction: let the mean element $\mathbf{m} = 0$. Then, given an infinite dimensional Hilbert space \mathcal{H} , choose an orthonormal sequence $\{e_n\}_{n \geq 1}$ and elements $\{a_n\}_{n \geq 1}$ in \mathcal{H} such that each a_n is a multiple of e_n . Initialize the algorithm with an element $c \in \mathcal{H}$ which is of small magnitude compared to the a_n and has a positive inner product with each a_n . The idea is that the different a_n 's will be chosen at one point by the algorithm and will add to the (rescaled) approximation error w_t of \mathbf{m} (t is the iteration number of the algorithm). In fact, we like to show that its norm will diverge to infinity.

This initial construction has a few problems which have to be addressed to make this construction work. The first problem with this construction is that $\langle a_n, c \rangle$ is positive. In fact, $\langle a_n, e_n \rangle, \langle c, e_n \rangle > 0$ for all $n \geq 1$. But, we want the mean element \mathbf{m} to be 0. Hence, we will need probability mass on the negative side to counter the mass accumulated by the a_n and c . We can achieve this by introducing another set of elements $\{b_n\}_{n \geq 1}$ which are lying opposite to the a_n . Therefore, each b_n is a negative multiple of e_n . These b_n need to be further constraint in magnitude. If they are of a similar order like the a_n then they can cancel the weight added to w_t by the a_n 's. We are using here sequences with values in the order of $1/\ln(n+1)$ for a_n and -2^{-n} for b_n .

Even though the b_n 's are of small magnitude compared to the a_n 's it is not directly obvious why these b_n 's should not be chosen many times by the algorithm to cancel step-by-step the weight accumulated by the a_n 's. Here is an argument why this does not happen: the a_n 's are constructed such that each a_n is chosen exactly once and they are selected in order by the algorithm. At a given iteration there is then an element a_m which has not yet been chosen and our construction assures that in this

case $\langle a_m, w_t \rangle$ equals the initial value $\langle a_m, c \rangle$, which is of magnitude $1/(m \ln(m+1))$. Since the algorithm chooses the element $h \in \phi[\mathcal{X}]$ that maximizes the inner product with w_t we can infer that this inner product must be larger than $1/(m \ln(m+1))$. Or put differently, an element b_n will only be chosen if $\langle b_n, w_t \rangle \geq 1/(m \ln(m+1))$, that is $\langle e_n, w_t \rangle \geq 2^n/(m \ln(m+1))$. If the algorithm chooses b_n then we are at least assured that $\langle e_n, w_{t+1} \rangle \geq 2^n/(m \ln(m+1)) - 2^{-m}$ (Figure 1.3 p. 484 visualizes these bounds for different m). We do not need this extra scaling of 2^n and we only use in the proof that there are sufficiently many e_n for which $|\langle e_n, w_t \rangle|$ is larger than $1/\ln(m+1)$. The number of elements for which the inner product is at least of this size grows in m and the sum over these inner products gives us a diverging number that approaches infinity in m . This is then sufficient to show that the norm of w_t diverges.

Interlacing

In the above discussion we assume $\mathbf{m} = 0$. However, constructing the probability measure such that $\mathbf{m} = 0$ is not straightforward. The problem is that the scaling on the positive side (the a_n 's and the c) is exponentially larger than the scaling on the negative side (the b_n 's). To get $\mathbf{m} = 0$ we would need the probability mass for the a_n 's and c times the magnitude of these elements to be scaled so that it equals the probability mass of the b_n 's times the scale of the b_n 's. The exponential difference in scale implies that the probability mass of the b_n 's needs to grow exponentially in n and the sum of all this mass has to add up to infinity.

By closer inspection, one can observe that the a_n 's pose no serious problem since one can just downscale the probability assigned to them by an exponential factor. However, the c poses a more serious problem. Let $p > 0$ be the probability corresponding to c . We use $c = \sum_{n=1}^{\infty} n^{-1} e_n$ and we thus have a factor of p/n pulling the mean element towards the positive direction in dimension n . Hence, we will need a probability of $p_n = p2^n/n$ for the b_n elements to counter this pull. Since p does not change with n we are left with p_n 's that grow rapidly in n .

Using an initialization c is in a way too rigid and does not allow us to assign lower probability mass as n increases. One way to overcome this problem is to break the initialization up and add probability mass to the different dimensions while the algorithm is running. We do this by replacing the single c with infinitely many elements, one for each dimension e_n . Since we do not want to alter the overall behavior of the algorithm these different elements will need to be of a low scale and we need to sum multiple elements to regain the $1/n$ value that c would have assigned. Therefore, for each dimension e_n , we are left with a finite sequence of elements $c_{n,1}, c_{n,2}, \dots$ which takes the role of the original c .

The question is then how we can guarantee that all these $c_{n,i}$ elements are chosen to simulate the initialization through c before the algorithm proceeds as usual. We guarantee this by introducing dimensions $\tilde{e}_{n,i}$ which are orthogonal to all the e_n . These dimensions are used to force the algorithm to choose $c_{n,i+1}$ after $c_{n,i}$ until the final element of the sequence is chosen and we have a weight of $1/n$ in dimension e_n .

We still have not addressed the problem of assigning different probabilities to the different dimensions. But, since c is now broken into many small pieces, it is easy to “lose” probability in n .

PROPOSITION 6.2. — *For any infinite-dimensional Hilbert space \mathcal{H} there exists a continuous function $\phi : \mathcal{X} \rightarrow \mathcal{H}$, $\mathcal{X} := [0, 1]$, a probability measure P on $\mathcal{B}_{[0,1]}$ which assigns positive measure to any open subset of \mathcal{X} , and an initialization $w_1 \in \phi[\mathcal{X}]$ such that the kernel herding algorithm when applied to $\int \phi(x) dP(x)$ generates a sequence $\{w_t\}_{t \geq 1}$ that is unbounded and the algorithm does not converge with a $1/t$ rate to $\mathbf{m} = \int \phi(x) dP(x) \in \mathcal{H}$.*

Proof of Proposition 6.2.

(a): Definition of the convex set. — Let $\{N_i\}_{i=1}^\infty$ be a set of natural numbers to be defined below, pick a countable infinite orthonormal sequence $\{e'_n\}_{n \geq 1}$ in \mathcal{H} and split this sequence into $\{e_n\}_{n \geq 1}$ and the sequences $\tilde{e}_{n,1}, \dots, \tilde{e}_{n,N_n}$ where n goes through $2, 3, \dots$. This can be done since these are countable many sequences of $N_n + 1$ elements and since countable unions of countable sets are again countable. Furthermore, define the sequences $\{a_n\}_{n \geq 1}, \{b_n\}_{n \geq 1}, \{c_{n,m} : 1 \leq n, 1 \leq m \leq N_n\}, \{d_n : 2 \leq n\} \subseteq \mathcal{H}$ by

$$\begin{aligned} a_n &:= \left(a'_n + \frac{1}{n}\right)e_n \quad \text{with} \\ a'_n &:= C \left\lceil \frac{2^n}{\ln(n+1)} \right\rceil 2^{-n} \quad \text{and} \quad C = 4 \left\lceil 3 + \frac{4 \ln(9)}{\ln(2)} \right\rceil = 64, \\ b_n &:= -2^{-n}e_n, \\ N_1 &:= 1 \quad \text{and for } n \geq 2, \quad N_n := \left\lceil \frac{2}{n \langle -b_n, e_n \rangle} \right\rceil, \\ c_{1,1} &:= e_1 + \alpha_{2,1} \tilde{e}_{2,1}, \quad \text{and for } n \geq 2 : \\ c_{n,1} &:= \beta_n e_n + \alpha_{n,1} \tilde{e}_{n,1} - \alpha_{n,2} \tilde{e}_{n,2}, \dots \\ c_{n,N_n-1} &:= \beta_n e_n + \alpha_{n,N_n-1} \tilde{e}_{n,N_n-1} - \alpha_{n,N_n} \tilde{e}_{n,N_n}, \\ c_{n,N_n} &:= \beta_n e_n + \alpha_{n,N_n} \tilde{e}_{n,N_n} - \alpha_{n+1,1} \tilde{e}_{n+1,1}, \\ d_2 &:= -(1/2)\alpha_{2,1} \tilde{e}_{2,1} \quad \text{and for all } 2 \leq n \text{ let} \\ d_n &:= (1/2)\alpha_{n,1} \tilde{e}_{n,1}, \\ \beta_n &:= -\frac{1}{nN_n}, \quad \text{for which } -\beta_n < \langle -b_n, e_n \rangle \text{ holds,} \\ \alpha_{n,1} &:= \sqrt{\frac{\langle e_n, a_n \rangle}{n}}, \quad \text{and} \quad \alpha_{n,i} := \sqrt{\alpha_{n,1}^2 + (i-1)\beta_n^2} \quad \text{for } 2 \leq i \leq N_n. \end{aligned}$$

$-\beta_n$ is smaller than $\langle -b_n, e_n \rangle$ because $-\beta_n = 1/(nN_n) \leq \langle -b_n, e_n \rangle/2$. Also observe that the sequence a'_n is non-increasing in n since

$$\left\lceil \frac{2^{n+1}}{\ln(n+2)} \right\rceil \frac{1}{2^{n+1}} \leq \left\lceil 2 \frac{2^n}{\ln(n+1)} \right\rceil \frac{1}{2^{n+1}} \leq \left\lceil 2 \left\lceil \frac{2^n}{\ln(n+1)} \right\rceil \right\rceil \frac{1}{2^{n+1}} = \left\lceil \frac{2^n}{\ln(n+1)} \right\rceil \frac{1}{2^n},$$

where we used that the function $\lceil \cdot \rceil$ is monotonically increasing.

(b): Construction of a continuous map ϕ . — We construct a continuous function $\phi : [0, 1] \rightarrow \mathcal{H}$ which goes through the points $\{a_n\}_{n \in \mathbb{N}}, \{b_n\}_{n \in \mathbb{N}}, \{c_{n,i} : 1 \leq i \leq n, 1 \leq n \leq N_n\}$ and $\{d_n : 2 \leq n\}$. We split the construction into three separate functions, ϕ_1 for the a_n, b_n elements, ϕ_2 for the $c_{n,i}$ and ϕ_3 for the d_n elements.

For ease of reading let $y_n = 1/(n + 1)$ and $z_n = (y_n + y_{n+1})/2$ for all $n \geq 1$. Define $\phi_1 : [0, 1] \rightarrow \mathcal{H}$, with n going through $1, 2, 3, \dots$, by

$$\phi_1(x) := \begin{cases} \frac{1-x}{1-y_1} a_1 & \text{if } y_1 < x \leq 1, \\ \frac{x-\xi}{y_n-\xi} a_n & \text{if } \xi := \frac{y_n+z_n}{2} < x \leq y_n, \\ \frac{\xi-x}{\xi-z_n} b_n & \text{if } z_n < x \leq \frac{y_n+z_n}{2} =: \xi, \\ \frac{x-\xi}{z_n-\xi} b_n & \text{if } \xi := \frac{y_{n+1}+z_n}{2} < x \leq z_n, \\ \frac{\xi-x}{\xi-y_{n+1}} a_{n+1} & \text{if } y_{n+1} < x \leq \frac{y_{n+1}+z_n}{2} =: \xi, \\ 0 & \text{if } x = 0. \end{cases}$$

The function is continuous on $(0, 1]$ as it is piecewise linear and the end points of the lines are connected. The only critical point is 0. For continuity at 0 it suffices that for any $\epsilon > 0$ we can pick a δ such that $x < \delta$ implies $\|\phi(x)\| < \epsilon$. We restrict the search for a δ to points $1/n, n \in \mathbb{N}$. For such a δ the maximum of $\phi(x)$ in an interval $[0, \delta]$ is either attained on an a_n or a b_n . As we have that $\lim_{n \rightarrow \infty} \|a_n\| = \lim_{n \rightarrow \infty} \|b_n\| = 0$ there is for every $\epsilon > 0$ an $N \in \mathbb{N}$ such that for all $n > N$ we have $\max(\|a_n\|, \|b_n\|) < \epsilon$ and, consequently for $\delta = 1/(N + 1)$ we have that $\|\phi(x)\| < \epsilon$ for any $0 \leq x \leq \delta$.

In the following let $N_1 := 1$. Furthermore, let $\tilde{y}_n := 1/n, \Delta_n := (\tilde{y}_n - \tilde{y}_{n+1})/N_n, u_{n,m} := \tilde{y}_n - m\Delta_n, u_{n,0} := \tilde{y}_n$ and let $\tilde{z}_{n,m} := (u_{n,m-1} + u_{n,m})/2$, for all $n \geq 1$ and $1 \leq m \leq N_n - 1$. With n going through all of $1, 2, \dots$, define $\phi_2 : [0, 1] \rightarrow \mathcal{H}$ by

$$\phi_2(x) := \begin{cases} \frac{\tilde{y}_n-x}{\tilde{y}_n-\tilde{z}_{n,1}} c_{n,1} & \text{if } \tilde{z}_{n,1} < x \leq \tilde{y}_n, \\ \frac{x-u_{n,m}}{\tilde{z}_{n,m}-u_{n,m}} c_{n,m} & \text{if } u_{n,m} < x \leq \tilde{z}_{n,m}, 1 \leq m \leq N_n - 1, \\ \frac{u_{n,m}-x}{u_{n,m}-\tilde{z}_{n,m+1}} c_{n,m+1} & \text{if } \tilde{z}_{n,m+1} < x \leq u_{n,m}, 1 \leq m \leq N_n - 1, \\ \frac{x-\tilde{y}_{n+1}}{\tilde{z}_{n,N_n}-\tilde{y}_{n+1}} c_{n,N_n} & \text{if } \tilde{y}_{n+1} < x \leq \tilde{z}_{n,N_n}, \\ 0 & \text{if } x = 0. \end{cases}$$

Similarly, by going through all $n \geq 2$ define $\phi_3 : [0, 1] \rightarrow \mathcal{H}$ by

$$\phi_3(x) := \begin{cases} 0 & \text{if } 1/2 < x \leq 1, \\ \frac{\tilde{y}_n-x}{\tilde{y}_n-\tilde{z}_{n,1}} d_n & \text{if } \tilde{z}_{n,1} < x \leq \tilde{y}_n, \\ \frac{x-u_{n,1}}{\tilde{z}_{n,1}-u_{n,1}} d_n & \text{if } u_{n,1} < x \leq \tilde{z}_{n,1}, \\ 0 & \tilde{y}_{n+1} \leq x \leq u_{n,1}, \\ 0 & \text{if } x = 0. \end{cases}$$

With the same reasoning as for ϕ_1 one can infer that ϕ_2 and ϕ_3 are continuous. Define

$$\phi(x) := \begin{cases} \phi_1(3x - 2) & \text{if } 2/3 < x \leq 1, \\ \phi_2(3x - 1) & \text{if } 1/3 < x \leq 2/3, \\ \phi_3(3x) & \text{if } 0 \leq x \leq 1/3. \end{cases}$$

The function ϕ is continuous since ϕ_1, ϕ_2, ϕ_3 are continuous and $\phi_1(0) = \phi_2(1) = \phi_2(0) = \phi_3(1) = 0$. This implies also that $\phi : [0, 1] \rightarrow \mathcal{H}$ is measurable. It is also Bochner integrable with respect to any probability measure defined on the Borel sets of \mathbb{R} as $\|\phi(\cdot)\| : [0, 1] \rightarrow \mathbb{R}$ is continuous and, hence, bounded, i.e. $\int \|\phi(x)\| dP(x) < \infty$.

(c): Definition of the probability measure. — We construct a Borel measure by defining a density p on $[0, 1]$. Using the variables defined for ϕ_1, ϕ_2, ϕ_3 , constants $\mathbf{a}_1, \mathbf{b}_1, \dots$ and n going through $1, 2, \dots$ we set

$$p_1(x) = \begin{cases} \mathbf{a}_1 & \text{if } y_1 < x \leq 1, \\ \mathbf{a}_n & \text{if } \frac{y_n+z_n}{2} < x \leq y_n, \\ \mathbf{b}_n & \text{if } \frac{y_{n+1}+z_n}{2} < x \leq \frac{y_n+z_n}{2}, \\ \mathbf{a}_{n+1} & \text{if } y_{n+1} < x \leq \frac{y_{n+1}+z_n}{2}, \\ 0 & \text{if } x = 0. \end{cases}$$

and using constants \mathbf{c}_n we furthermore define

$$p_2(x) := \begin{cases} \mathbf{c}_n, & \text{if } \tilde{y}_{n+1} < x \leq \tilde{y}_n, \\ 0 & \text{if } x = 0. \end{cases}$$

Finally, going through all $n \geq 2, 1 \leq m \leq N_n$ and with the constants $\mathbf{d}_{n,m}$ let

$$p_3(x) := \begin{cases} 1 & \text{if } 1/2 < x \leq 1, \\ \mathbf{d}_n & \text{if } u_{n,1} \leq x \leq \tilde{y}_n, \\ 1 & \text{if } \tilde{y}_{n+1} \leq x \leq u_{n,1} \\ 0 & \text{if } x = 0 \end{cases}$$

and combine these to define the density p by

$$p(x) := \begin{cases} p_1(3x - 2) & \text{if } 2/3 < x \leq 1, \\ p_2(3x - 1) & \text{if } 1/3 < x \leq 2/3, \\ p_3(3x) & \text{if } 0 \leq x \leq 1/3. \end{cases}$$

Now, $\mathbf{m} = 0$ iff $\langle e_n, \mathbf{m} \rangle = E\langle e_n, \phi \rangle = 0 = \langle e'_{n,i}, \mathbf{m} \rangle = E\langle e'_{n,i}, \phi \rangle$ for all $n \geq 1, 1 \leq i \leq N_n$.

Observe that in general, if $a, b \in [0, 1], a \leq b$, the density p is constant on $[a, b]$ with value $\mu \in [0, \infty), h \in \mathcal{H}$ and $\psi : [0, 1] \rightarrow \mathcal{H}$ is defined by

$$\psi(x) = \begin{cases} (x - a)/(b - a)\mu h & x \in [a, b], \\ 0 & \text{otherwise} \end{cases}$$

then for any e_n (and $e'_{n,i}$)

$$\langle e_n, E\psi \rangle = E\langle e_n, \psi \rangle = \int_{[a,b]} \frac{x - a}{b - a} \mu \langle e_n, h \rangle = (1/2)\mu \langle e_n, h \rangle (b - a)$$

and if

$$\psi(x) = \begin{cases} (b - x)/(b - a)\mu h & x \in [a, b], \\ 0 & \text{otherwise} \end{cases}$$

then

$$\langle e_n, E\psi \rangle = (1/2)\mu\langle e_n, h \rangle(b - a).$$

So,

$$\langle e_1, \mathbf{m} \rangle = \frac{1}{6} \left\langle e_1, \left(1 - \frac{y_1 + z_1}{2}\right) \mathbf{a}_1 a_1 + \left(\frac{y_1 - y_2}{2}\right) \mathbf{b}_1 b_1 + (\tilde{y}_1 - \tilde{y}_2) \mathbf{c}_1 c_{1,1} \right\rangle$$

will be zero by setting

$$\mathbf{b}_1 := 2 \cdot 6 \left(\left(1 - \frac{y_1 + z_1}{2}\right) \mathbf{a}_1 \frac{\langle a_1, e_1 \rangle}{\langle -b_1, e_1 \rangle} + \left(1 - \frac{1}{2}\right) \mathbf{c}_1 \frac{\langle c_{1,1}, e_1 \rangle}{\langle -b_1, e_1 \rangle} \right)$$

and $\langle e_n, \mathbf{m} \rangle = 0$ by setting

$$\mathbf{b}_n = 2(n + 1)(n + 2) \left(\frac{1}{4} \left(\frac{1}{n} - \frac{1}{n + 2} \right) \mathbf{a}_n \frac{\langle a_n, e_n \rangle}{\langle -b_n, e_n \rangle} + \left(\frac{1}{n} - \frac{1}{n + 1} \right) \frac{\mathbf{c}_n \beta_n}{\langle -b_n, e_n \rangle} \right).$$

Also, for any $n \geq 1$ we have that $\mathbf{b}_n > 0$ if $\mathbf{a}_n, \mathbf{c}_n > 0$. Let, \mathcal{N} be a normalising constant to be defined below and let

$$\mathbf{a}_1 := \frac{\mathcal{N} \langle -b_1, e_1 \rangle}{24 \langle a_1, e_1 \rangle} \Big/ \left(1 - \frac{y_1 + z_1}{2}\right) > 0 \quad \text{and} \quad \mathbf{c}_1 := \frac{\mathcal{N} \langle -b_1, e_1 \rangle}{12 \langle c_{1,1}, e_1 \rangle} > 0$$

such that $\mathbf{b}_1 = \mathcal{N}$. Also set for all $n \geq 2$

$$\mathbf{a}_n := \frac{n}{n + 1} \frac{\mathcal{N} \langle -b_n, e_n \rangle}{\langle a_n, e_n \rangle} > 0 \quad \text{and} \quad \mathbf{c}_n := \frac{n}{2(n + 2)} \frac{\mathcal{N} \langle b_n, e_n \rangle}{\beta_n} \frac{\Delta_{n+1}}{\Delta_n} > 0$$

which makes $\mathbf{b}_n = \mathcal{N}(1 - \Delta_{n+1}/\Delta_n)$ and all $\langle e_n, \mathbf{m} \rangle = 0$. For the elements $\tilde{e}_{n,i}$ we have that

$$6 \langle \tilde{e}_{2,1}, \mathbf{m} \rangle = \Delta_1 \mathbf{c}_1 \langle c_{1,1}, \tilde{e}_{2,1} \rangle + \Delta_2 \mathbf{c}_2 \langle c_{2,1}, \tilde{e}_{2,1} \rangle + \Delta_2 \mathbf{d}_2 \langle d_2, \tilde{e}_{2,1} \rangle$$

and we set

$$\mathbf{d}_2 := \frac{\Delta_1}{\Delta_2} \mathbf{c}_1 \frac{\langle c_{1,1}, \tilde{e}_{2,1} \rangle}{\langle -d_2, \tilde{e}_{2,1} \rangle} + \mathbf{c}_2 \frac{\langle c_{2,1}, \tilde{e}_{2,1} \rangle}{\langle -d_2, \tilde{e}_{2,1} \rangle} > 0.$$

Furthermore, for all $n > 2$ let

$$\mathbf{d}_n := \frac{\Delta_{n-1}}{\Delta_n} \mathbf{c}_{n-1} \frac{\langle -c_{n-1, N_{n-1}}, \tilde{e}_{n,1} \rangle}{\langle d_n, \tilde{e}_{n,1} \rangle} - \mathbf{c}_n \frac{\langle c_{n,1}, \tilde{e}_{n,1} \rangle}{\langle d_n, \tilde{e}_{n,1} \rangle}.$$

$\mathbf{d}_n > 0$ for $n > 2$ since

$$\begin{aligned} \frac{2 \langle -d_n, \tilde{e}_{n,1} \rangle}{\alpha_{n,1} \mathcal{N}} \mathbf{d}_n &= \frac{(n - 1)^2}{n + 1} \langle -b_{n-1}, e_{n-1} \rangle \left[\frac{2}{(n - 1) \langle -b_{n-1}, e_{n-1} \rangle} \right] \\ &\quad - \frac{\Delta_{n+1}}{\Delta_n} \frac{n^2}{n + 2} \langle -b_n, e_n \rangle \left[\frac{2}{n \langle -b_n, e_n \rangle} \right] \\ &\geq 2 \frac{n - 1}{n + 1} - 3 \frac{n}{n + 2} \frac{\Delta_{n+1}}{\Delta_n} = 2 \frac{n - 1}{n + 1} - 3 \frac{n^2}{(n + 2)^2} \frac{N_n}{N_{n+1}} \\ &\geq 2 \frac{n - 1}{n + 1} - \frac{3 n^2 (n + 1)}{2 (n + 2)^2} \left(\frac{1}{n} + \frac{1}{2^{n+1}} \right) \end{aligned}$$

which is strictly greater zero if

$$4(n - 1)(n + 2)^2 - 4n(n + 1)^2 = 4(n^2 + n - 4)$$

is. But this is obvious for $n \geq 3$. For all remaining $n, i \geq 2$ we can observe that

$$\langle \tilde{e}_{n,i}, \mathbf{m} \rangle = (1/6)\Delta_n \mathbf{c}_n \langle c_{n,i-1} - c_{n,i}, \tilde{e}_{n,i} \rangle = 0$$

for all $m \geq 2$. Hence, $\mathbf{m} = 0$ and the density is strictly greater 0 on all but three points. It remains to set \mathcal{N} such that the density integrates to one. We have for any $\mathcal{N} > 0$ that

$$\begin{aligned} 0 < \int_{[0,1]} p &= \mathbf{a}_1(1 - (y_1 + z_1)/2)/3 + (1/3) \sum_{n=2}^{\infty} \mathbf{a}_n((y_n - z_{n-1}) - (y_n - z_n))/2 \\ &+ (1/3) \sum_{n=1}^{\infty} \mathbf{b}_n((y_n - z_n) - (y_{n+1} - z_n))/2 \\ &+ \sum_{n=1}^{\infty} \frac{1}{3n(n+1)} \mathbf{c}_n + 1/6 + (1/3) \sum_{n=1}^{\infty} \Delta_n \mathfrak{d}_n + (1/3) \sum_{n=1}^{\infty} (u_{n,1} - \tilde{y}_{n+1}). \end{aligned}$$

The first sum is a finite multiple of \mathcal{N} since $\mathbf{a}_n \approx 2^{-n}$ and the sum over $((y_n - z_{n-1}) - (y_n - z_n))/2$ is bounded by 1. Similarly, the \mathbf{c}_n sum is bounded since \mathbf{c}_n itself is upper bounded by \mathcal{N} and the rest is quadratic in n . Furthermore, \mathbf{b}_n is upper bounded by \mathcal{N} and the sum of the intervals cannot exceed 1. Finally, \mathfrak{d}_n is upper bounded since

$$\frac{2\langle -d_n, \tilde{e}_{n,1} \rangle}{\alpha_{n,1}\mathcal{N}} \mathfrak{d}_n \leq \frac{2(n-1)}{n+1}$$

$\alpha_{n,1}$ is bounded and so is $\langle -d_n, \tilde{e}_{n,1} \rangle$. Hence, the sum is a finite multiple of \mathcal{N} and we have in total a term that is a finite multiple of \mathcal{N} plus a constant that is smaller than 1/2. Therefore, we can choose \mathcal{N} such that $\int p = 1$.

(d): Behaviour of the algorithm. — Initialize the algorithm with $w_1 := c_{1,1} \in \phi[X]$ and let x_t be the element which is chosen at stage t . The algorithm behaves as follows:

- (1) For any $t \geq 1$, if $w_t \neq 0$ then $x_t \in \{a_n\}_{n \geq 1} \cup \{b_n\}_{n \geq 1} \cup \{c_{n,m} : 1 \leq n, 1 \leq m \leq N_n\} \cup \{d_{n,m} : 2 \leq n, 1 \leq m \leq N_n\}$.
- (2) Let $n = \min\{m : a_m \text{ has not been chosen in steps } 1, \dots, t-1\}$. If $t \geq 2$ then either the smallest element of $\{(m, j) : c_{m,j} \text{ has not been chosen in steps } 1, \dots, t-1\} \setminus \{c_{1,1}\}$ in the lexicographic order is (n, i) with $1 \leq i \leq N_n$ and

$$w_t = -\gamma_1 e_1 - \dots - \gamma_{n-1} e_{n-1} + \gamma_n e_n + \alpha_{n,i} \tilde{e}_{n,i},$$

where

$$\gamma_j = (2^j a'_j - l) 2^{-j}, l \in \mathbb{N} \quad \text{and} \quad a'_j \geq \gamma_j \geq \min \left\{ a'_j, \max \left\{ \frac{2^j \langle a_n, e_n \rangle}{n} - 2^{-j}, 0 \right\} \right\}$$

for $1 \leq j \leq n-1$ and $\gamma_n = -(i-1)\beta_n$ (first case), or the smallest element is $(n+1, 1)$ and

$$w_t = -\gamma_1 e_1 - \dots - \gamma_{n-1} e_{n-1} + \gamma_n e_n + \alpha_{n+1,1} \tilde{e}_{n+1,1},$$

with $\gamma_1, \dots, \gamma_{n-1}$ like above and $\gamma_n = 1/n$ (second case). In particular $w_t \neq 0$.

- (3) Let $N(n) := \lceil 1 + \log_2(n \ln(n+1)) \rceil$ then $n-1 \geq N(n)$ for all $n \geq 7$. If n is the smallest index of an a_n which has not been chosen yet and if this $n \geq 7$ then for any i with $n-1 \geq i \geq N(n)$

$$\langle e_i, w_t \rangle \leq -\frac{1}{\ln(n+1)}.$$

- (4) For each $n \geq 1$ there exists a step $t \geq 1$ with $x_t = a_n$.
- (α) (1) is saying that no point on the line from 0 to an $a_n, b_n, c_{n,m}$ or $d_{n,m}$ is chosen that differs from $a_n, b_n, c_{n,m}$ and $d_{n,m}$. To see this first observe that only points $\phi(x)$ will be chosen at any stage t for which $\langle w_t, \phi(x) \rangle > 0$. By assumption $w_t \neq 0$. If there exists an e_n with $\langle e_n, w_t \rangle \neq 0$ then either $\langle a_n, w_t \rangle$ or $\langle b_n, w_t \rangle$ is strictly positive. Also, if there is an $\tilde{e}_{n,m}, (n, m) \neq (2, 1)$, such that $\langle \tilde{e}_{n,m}, w_t \rangle > 0$ then $\langle d_{n,m}, w_t \rangle$ is strictly positive. Similarly, if $\langle \tilde{e}_{2,1}, w_t \rangle < 0$ then $\langle d_{2,1}, w_t \rangle > 0$. Assuming that none of these cases apply we have that either $\langle \tilde{e}_{2,1}, w_t \rangle > 0$ or there is an $\tilde{e}_{n,m}, (n, m) \neq (2, 1)$, with $\langle \tilde{e}_{n,m}, w_t \rangle < 0$. In the first case $\langle c_{1,1}, w_t \rangle > 0$. In the latter case and with $\langle \tilde{e}_{2,1}, w_t \rangle = 0$ let $(n', m') := \min\{(n, m) : \langle \tilde{e}_{n,m}, w_t \rangle < 0\}$ where the minimum is taken wrt. the lexicographic ordering. We have $\langle c_{n',m'-1}, w_t \rangle = \alpha_{n',m'-1} \langle \tilde{e}_{n',m'-1}, w_t \rangle - \alpha_{n',m'} \langle \tilde{e}_{n',m'}, w_t \rangle > -\alpha_{n',m'} \langle \tilde{e}_{n',m'}, w_t \rangle > 0$, if $m' > 1$, and if $m' = 1$ then

$$\langle c_{n'-1, N_{n'-1}}, w_t \rangle = \alpha_{n'-1, N_{n'-1}} \langle \tilde{e}_{n'-1, N_{n'-1}}, w_t \rangle - \alpha_{n',1} \langle \tilde{e}_{n',1}, w_t \rangle > -\alpha_{n',1} \langle \tilde{e}_{n',1}, w_t \rangle > 0.$$

If the chosen $\phi(x)$ is on the line from 0 to an a_n then $a_n = \xi \phi(x)$ with $\xi \geq 1$ and $0 < \langle \phi(x), w_t \rangle \leq \xi \langle \phi(x), w_t \rangle = \langle a_n, w_t \rangle$ and $\phi(x) = a_n$. The same argument applies to $b_n, c_{n,m}$ and $d_{n,m}$.

- (β) We prove by induction over $t \geq 2$ that (2) holds. We start with the *induction basis*. $w_1 = c_{1,1} = e_1 + \alpha_{2,1} \tilde{e}_{2,1}$ and we have $\langle w_1, b_n \rangle \leq 0, \langle w_1, d_n \rangle \leq 0$ for all n , $\langle w_1, a_n \rangle = 0$ for all $n \geq 2$, $\langle w_1, c_{n,i} \rangle = 0$ if either $n > 2$ or $(n = 2$ and $i > 2)$. Furthermore,

$$\begin{aligned} \langle w_1, c_{1,1} \rangle &= \|c_{1,1}\|^2 = 1 + \alpha_{2,1}^2 = 1 + \frac{\langle e_2, a_2 \rangle}{2} = 1 + \frac{C}{8} \left\lceil \frac{4}{\ln(3)} \right\rceil + \frac{1}{4} \\ &\leq 1 + \frac{C}{4} \left\lceil \frac{2}{\ln(2)} \right\rceil + \frac{1}{4} \\ &< 1 + \frac{C}{2} \left\lceil \frac{2}{\ln(2)} \right\rceil = \langle w_1, a_1 \rangle \end{aligned}$$

since $C \geq 1, \lceil 2/\ln(2) \rceil = 3$ and hence

$$\frac{C}{4} \left\lceil \frac{2}{\ln(2)} \right\rceil > \frac{1}{4}.$$

Also, $\langle w_1, c_{2,1} \rangle = \alpha_{2,1}^2 < \langle w_1, c_{1,1} \rangle < \langle w_1, a_1 \rangle$ and $x_1 = a_1$. Therefore,

$$w_2 = w_1 - a_1 = e_1 + \alpha_{2,1} \tilde{e}_{2,1} - (a'_1 + 1)e_1 = -a'_1 e_1 + \alpha_{2,1} \tilde{e}_{2,1}$$

and w_2 has the promised form.

- (γ) Next, we address the *induction step*.

- (i) Assuming w_t has the given form in step t we can observe that

$$\langle x_t, w_t \rangle \geq \frac{\langle a_n, e_n \rangle}{n} > 0$$

since in the first case

$$\langle c_{n,i}, w_t \rangle = \beta_n \gamma_n + \alpha_{n,i}^2 = -(i-1)\beta_n^2 + (\alpha_{n,1}^2 + (i-1)\beta_n^2) = \frac{\langle e_n, a_n \rangle}{n},$$

in case that $i > 1$ or, for $i = 1$,

$$\langle c_{n,i}, w_t \rangle = \alpha_{n,1}^2 = \frac{\langle e_n, a_n \rangle}{n}.$$

In the second case,

$$\langle a_n, w_t \rangle \geq \gamma_n \langle e_n, a_n \rangle = \frac{\langle e_n, a_n \rangle}{n}.$$

- (ii) For the b_j (i) implies that, first, no $j \geq n$ will have been chosen in t since for these $\langle b_j, w_t \rangle \leq 0$ holds in both cases. Also, if for a j , $1 \leq j \leq n - 1$, $\gamma_j < 2^j \langle a_n, e_n \rangle / n$ then $\langle b_j, w_t \rangle = \gamma_j 2^{-j} < \langle a_n, e_n \rangle / n$ and $b_j \neq x_t$. On the other hand, if $\gamma_j \geq 2^j \langle a_n, e_n \rangle / n$ and $x_t = b_j$ then the coefficient changes by -2^{-j} , i.e. the new coefficient is

$$\gamma_j - 2^{-j} \geq 2^j \frac{\langle a_n, e_n \rangle}{n} - 2^{-j}.$$

The coefficient is also always non-negative since γ_j is a multiple of 2^{-j} and b_j will not be selected if $\gamma_j = 0$. In total, all cases are consistent with our induction hypothesis and we are safe against any application of b_j .

- (iii) In terms of a_j , we can directly observe that $\langle a_j, w_t \rangle = 0$ if $j > n$ and $\langle a_j, w_t \rangle = -\gamma_j \langle e_j, a_j \rangle \leq 0$ if $j < n$. So only a_n might have been chosen at time t . However, in the first case we have that

$$\langle a_n, w_t \rangle = \gamma_n \langle a_n, e_n \rangle = -(i - 1) \beta_n \langle a_n, e_n \rangle \leq \frac{N_n - 1}{n N_n} \langle a_n, e_n \rangle < \frac{\langle a_n, e_n \rangle}{n}$$

and $x_t \neq a_n$. In the second case, if $\gamma_n = 1/n$, then

$$\langle a_n, w_t \rangle = \frac{\langle a_n, e_n \rangle}{n}.$$

Thus a_n might be chosen, and, in case it is, then the new coefficient is $\gamma_n - a'_n - 1/n = -a'_n$ which is consistent with the induction hypothesis.

- (iv) Turning to the $c_{n',i'}$ elements we can observe that for $(n, i) > (2, 1)$ we have $\langle w_t, c_{1,1} \rangle \leq 0$. For $(n, i) = (2, 1)$ we are in the first case since $N_2 = 4$ and

$$\begin{aligned} \langle w_t, c_{1,1} \rangle &= \langle -\gamma_1 e_1 + \gamma_2 e_2 + \alpha_{2,1} \tilde{e}_{2,1}, e_1 + \alpha_{2,1} \tilde{e}_{2,1} \rangle \\ &= -\gamma_1 + (1/2) \langle e_2, a_2 \rangle \\ &\leq -(1/2) \langle e_2, a_2 \rangle + 1/2 + (1/2) \langle e_2, a_2 \rangle. \end{aligned}$$

But, $1/2 < (1/2) \langle a_2, e_2 \rangle$ and $c_{1,1}$ will never be chosen. For the remaining $c_{n',i'}$ elements we have in the first case in step t that no $c_{n',i'}$ will be chosen if $n' \neq n$ or $i' \neq i$ since in these cases

$$(*) \quad \langle w_t, c_{n',i'} \rangle \begin{cases} = 0 & \text{if } n' > n, \\ = \gamma_n \beta_n = -(i - 1) \beta_n^2 \leq 0 & \text{if } n' = n \text{ and } (i' > i \text{ or } i' < i - 1), \\ = \gamma_n \beta_n - \alpha_{n,i}^2 < 0 & \text{if } n' = n \text{ and } i' + 1 = i, \\ \leq 0 & \text{if } 2 \leq n' = n - 1 \\ & \text{and } i' = N_{n'} \text{ and } i = 1, \quad (*) \\ = -\gamma_{n'} \beta_{n'} < \langle w_t, b_{n'} \rangle & \text{if } n' < n \text{ and } (*) \text{ does not apply.} \end{cases}$$

(*) follows for $n \geq 3$ from

$$\begin{aligned} \langle w_t, c_{n-1, N_{n-1}} \rangle &= \langle w_t, \beta_{n-1} e_{n-1} + \alpha_{n-1, N_{n-1}} \tilde{e}_{n-1, N_{n-1}} - \alpha_{n,1} \tilde{e}_{n,1} \rangle \\ &= -\gamma_{n-1} \beta_{n-1} - \alpha_{n,1}^2 = \frac{\gamma_{n-1}}{(n-1)N_{n-1}} - \alpha_{n,1}^2 \\ &\leq \gamma_{n-1} 2^{-n} - \frac{\langle e_n, a_n \rangle}{n} \leq 2^{-n} a'_{n-1} - \frac{a'_n}{n} \\ &= C \left(2^{-n} \left\lceil \frac{2^{n-1}}{\ln(n)} \right\rceil 2^{-(n-1)} - \frac{1}{n} \left\lceil \frac{2^n}{\ln(n+1)} \right\rceil 2^{-n} \right) \\ &\leq C n^{-1} 2^{-2n-1} \left(n \left\lceil \frac{2^n}{\ln(n^2)} \right\rceil - 2^{n-1} \left\lceil \frac{2^n}{\ln(n+1)} \right\rceil \right) \leq 0. \end{aligned}$$

If now $x_t = c_{n,i}$ then, in case $i < N_n$,

$$\begin{aligned} w_{t+1} &= w_t - c_{n,i} \\ &= - \sum_{i=1}^{n-1} \gamma_i e_i - (i-1) \beta_n e_n + \alpha_{n,i} \tilde{e}_{n,i} - \beta_n e_n - \alpha_{n,i} \tilde{e}_{n,i} + \alpha_{n,i+1} \tilde{e}_{n,i+1} \\ &= - \sum_{i=1}^{n-1} \gamma_i e_i - i \beta_n e_n + \alpha_{n,i+1} \tilde{e}_{n,i+1} \end{aligned}$$

which has the desired form. Similarly, in case that $i = N_n$

$$\begin{aligned} w_{t+1} &= w_t - c_{n, N_n} \\ &= - \sum_{i=1}^{n-1} \gamma_i e_i - (N_n - 1) \beta_n e_n + \alpha_{n, N_n} \tilde{e}_{n, N_n} - \beta_n e_n - \alpha_{n, N_n} \tilde{e}_{n, N_n} + \alpha_{n+1,1} \tilde{e}_{n+1,1} \\ &= - \sum_{i=1}^{n-1} \gamma_i e_i - N_n \beta_n e_n + \alpha_{n+1,1} \tilde{e}_{n+1,1} \end{aligned}$$

which has the form of the second case since $-N_n \beta_n = 1/n$.

In the second case no element $c_{n',i'}$ will be chosen since

$$\langle w_t, c_{n',i'} \rangle = \begin{cases} 0 & \text{if } n' > n + 1 \text{ or } (n' = n + 1 \text{ and } i' > 1), \\ \alpha_{n+1,1}^2 < \langle e_n, a_n \rangle / n & \text{if } n' = n + 1 \text{ and } i' = 1, \\ \beta_n / n - \alpha_{n+1,1}^2 < 0 & \text{if } n' = n \text{ and } i' = N_n, \\ \beta_n / n < 0 & \text{if } n' = n \text{ and } i' < N_n, \\ -\gamma_{n'} \beta_{n'} < \langle w_t, b_{n'} \rangle \text{ or } = 0 & \text{if } n' < n. \end{cases}$$

(v) We turn to the d elements:

- First case: if $(n, i) = (2, 1)$ then $\langle w_t, d_2 \rangle = -(1/2) \alpha_{2,1}^2 < 0$ and otherwise $\langle w_t, d_2 \rangle = 0$. In either way d_2 will not be chosen. For any other n' we have that $\langle w_t, d_{n'} \rangle = 0$ if $(n, i) \neq (n', 1)$. Otherwise

$$\langle w_t, d_{n'} \rangle = (1/2) \alpha_{n,1}^2 = \frac{1}{2} \frac{\langle e_n, a_n \rangle}{n}$$

and $\langle w_t, d_{n'} \rangle < \langle a_n, e_n \rangle / n$ and $d_{n'}$ will never be chosen.

- Second case: $\langle w_t, d_2 \rangle \leq 0$ and all $d_{n'}$ with $n' \neq n + 1$ are zero. Finally

$$\langle w_t, d_{n+1} \rangle = (1/2)\alpha_{n+1,1}^2 = \frac{\langle e_{n+1}, a_{n+1} \rangle}{n + 1} < \frac{\langle e_n, a_n \rangle}{n}$$

since the sequence $\langle e_n, a_n \rangle$ is non-increasing.

So in step $t + 1$ the element w_{t+1} will have the right form, and, certainly, $w_{t+1} \neq 0$.

- (δ) Next we prove (3). Since the smallest index n for which a_n has not been chosen in rounds 1 to t is assumed to be larger than 7 we can assume that $t \geq 2$.

For each m , $\langle a_m, e_m \rangle / m \geq a'_m / m \geq 1 / (m \ln(m + 1))$. Hence, from (2) we conclude for all i with $n - 1 \geq i \geq N(n)$

$$\begin{aligned} -\langle e_i, w_{t+1} \rangle = \gamma_i &\geq \min \left\{ a'_i, \max \left\{ 2^i \frac{\langle a_n, e_n \rangle}{n} - 2^{-i}, 0 \right\} \right\} \\ &\geq \min \left\{ \frac{1}{\ln(i + 1)}, \max \left\{ \frac{2^i}{n \ln(n + 1)} - 2^{-i}, 0 \right\} \right\}. \end{aligned}$$

Using the assumption $i \geq N(n)$ we observe that

$$\begin{aligned} i &\geq 1 + \log_2(n \ln(n + 1)) = \log_2(2n \ln(n + 1)) \\ &\implies 2^i \geq 2n \ln(n + 1) \\ &\implies \frac{2^i}{n \ln(n + 1)} - 1 \geq 1 \end{aligned}$$

and since $i > 1$

$$\frac{2^i}{n \ln(n + 1)} - 2^{-i} \geq \frac{2^i}{n \ln(n + 1)} - 1 \geq \frac{1}{\ln(i + 1)} \geq \frac{1}{\ln(n + 1)}.$$

So, for $n - 1 \geq i \geq N(n)$,

$$-\langle e_i, w_{t+1} \rangle \geq \frac{1}{\ln(n + 1)}.$$

- (ϵ) (4) also follows from (2). First, if a_n has been chosen in any step t then for all $t' > t$ (2) tells us that $\langle a_n, w_{t'} \rangle \leq 0$. Consequently, $x_{t'} \neq a_n$ and a_n will be chosen at most ones. Also, if a_n is the element with minimal index which has not been chosen yet in time t then either $\langle w_t, a_n \rangle = \langle w_t, a_m \rangle = 0$ or $\langle w_t, a_n \rangle > \langle w_t, a_m \rangle$ for all $m > n$. In the first case no element $a_{n'}$ will be chosen and in the second case, if an $a_{n'}$ will be chosen it will be the one with the smallest index in the set of elements which have not been chosen yet. So the elements a_n will be chosen in order and no element will be skipped.

Let us now assume that $\{a_m : a_m \neq x_t \text{ for all } t \geq 1\}$ is not empty and let a_n be the element with the smallest index in this set.

The argument in (γ) shows us that no $c_{m,j}$ with $m > n$ will be chosen. Also, if $c_{m,j}$, $m \leq n$, has been chosen in any step t then for all $t' > t$ we again infer from (2) and the argument in (γ) that $c_{m,j}$ will not be chosen in t' and, hence, each $c_{m,j}$ is not chosen more than ones. Also none of the $d_{n,i}$ elements will be ever chosen.

So the only way that an a_n is never chosen is that infinite many b_m elements are selected. Yet, no b_m with $m \geq n$ will be chosen since the inner product with the weight vector is less or equal to zero. Also each b_m , $m < n$ can only be chosen finite many times before the weight vector in direction e_m becomes 0 and the inner product with b_m becomes 0 too (at which point it will certainly not be chosen any more). So only finite many applications of b_m 's are possible with a contradiction that a_n will not be chosen.

(e): *Unboundedness.* — d(3) and d(4) allow us now to show that the sequence $\{\|w_t\|\}_{t \geq 1}$ is unbounded. Assume that at stage t the element n is the smallest index such that $\langle a_n, w_t \rangle$ is positive.

For $n \geq 7$ we know from d.3 that $|\langle e_i, w_t \rangle| \geq 1/\ln(n+1)$ for all i , $N(n) \leq i < n$. Hence, for $n \geq 7$,

$$\begin{aligned} \|w_t\|^2 &= \sum_{i=1}^{\infty} |\langle e_i, w_t \rangle|^2 \geq \sum_{i=N(n)}^{n-1} \frac{1}{(\ln(n+1))^2} \geq \frac{n-1-N(n)}{(\ln(n+1))^2} \\ &= \frac{n-1 - \lceil 1 + \log_2(n \ln(n+1)) \rceil}{(\ln(n+1))^2} \geq \frac{n-3}{(\ln(n+1))^2} - \frac{\log_2(n \ln(n+1))}{(\ln(n+1))^2}. \end{aligned}$$

Furthermore,

$$\frac{\log_2(n \ln(n+1))}{(\ln(n+1))^2} = \frac{\ln(n) + \ln(\ln(n+1))}{\ln(2)(\ln(n+1))^2} \leq \frac{2}{\ln(2)},$$

since $\ln(x) \leq x$ for all $x > 0$ and $\ln(n+1) > 0$. Hence,

$$\|w_t\|^2 \geq \frac{n-3}{(\ln(n+1))^2} - \frac{2}{\ln(2)}.$$

The right side goes to infinity in n and, since for every n there is a t at which a_n is chosen due to d.4, the norm of w_t crosses any boundary at one time t . \square

COROLLARY 6.3. — *There exists a continuous kernel on $[0, 1]$, a Borel probability measure on $[0, 1]$ which assigns positive measure to open subsets of $[0, 1]$ and an initialization for which the algorithm does not converge with a $1/t$ rate to \mathbf{m} .*

Proof. — We consider the Hilbert space $(\mathcal{H}, \langle \cdot, \cdot \rangle)$ from Proposition 6.2 with the corresponding feature map $\phi : [0, 1] \rightarrow \mathcal{H}$. Define the continuous kernel function $k(x, y) := \langle \phi(x), \phi(y) \rangle$ on $[0, 1]$ and let the corresponding RKHS be $(\mathcal{K}, (\cdot, \cdot))$. The geometry of the two spaces is closely related. We have for scalars a_i, b_j and $x_i, y_j \in [0, 1]$, $i = 1, \dots, n, j = 1, \dots, m$, that

$$\left\langle \sum_{i=1}^n a_i \phi(x_i), \sum_{j=1}^m b_j \phi(y_j) \right\rangle = \sum_{i=1}^n \sum_{j=1}^m k(x_i, y_j) = \left(\sum_{i=1}^n a_i k(x_i, \cdot), \sum_{j=1}^m b_j k(y_j, \cdot) \right).$$

Furthermore, we know that the Bochner-integral $\mathbf{m} \in \mathcal{H}$ lies in $\text{cch } \phi[\mathcal{X}]$ which equals the closure of $\text{ch } \phi[\mathcal{X}]$ [Sim11][Theorem 5.2, p. 71] and there exists a sequence $\{n_i\}_{i \in \mathbb{N}}$, $n_i \in \mathbb{N}$, elements $x_{ij} \in [0, 1]$, and non-negative weights a_{ij} with $\sum_{j=1}^{n_i} a_{ij} = 1$ such that the sequence $\{s_i = \sum_{j=1}^{n_i} a_{ij} \phi(x_{ij})\}_{i \in \mathbb{N}}$ converges to \mathbf{m} in norm, i.e. $\|\mathbf{m} - s_i\| \rightarrow 0$

for $i \rightarrow \infty$. The corresponding sequence $\{\bar{s}_i = \sum_{j=1}^{n_i} a_{ij}k(x_{ij}, \cdot)\}_{i \in \mathbb{N}}$ is a Cauchy sequence in \mathcal{K} since

$$\|\bar{s}_i - \bar{s}_j\|_{\mathcal{K}} = \|s_i - s_j\|_{\mathcal{H}}$$

and has a limit $\mathbf{n} \in \mathcal{K}$ because \mathcal{K} is complete. In particular, for any $x \in \mathcal{X}$

$$\begin{aligned} |(\mathbf{n}, k(x, \cdot)) - \langle \mathbf{m}, \phi(x) \rangle| &= |(\mathbf{n} - \bar{s}_i, k(x, \cdot)) + (\bar{s}_i, k(x, \cdot)) - \langle s_i, \phi(x) \rangle + \langle s_i - \mathbf{m}, \phi(x) \rangle| \\ &\leq k(x, x)\|\mathbf{n} - \bar{s}_i\|_{\mathcal{K}} + \|\phi(x)\|_{\mathcal{H}}\|\mathbf{m} - s_i\|_{\mathcal{H}} \rightarrow 0 \text{ (in } i) \end{aligned}$$

and $(\mathbf{n}, k(x, \cdot)) = \langle \mathbf{m}, \phi(x) \rangle$ for every $x \in X$. Furthermore, for arbitrary l points $x_1, \dots, x_l \in X$ and scalars a_1, \dots, a_l it holds that

$$\left(\mathbf{n}, \sum_{i=1}^l a_i k(x_i, \cdot) \right) = \left\langle \mathbf{m}, \sum_{i=1}^l a_i \phi(x_i) \right\rangle$$

and

$$\left| \|\mathbf{n}\|_{\mathcal{K}} - \|\mathbf{m}\|_{\mathcal{H}} \right| \leq \left| \|\mathbf{n}\|_{\mathcal{K}} - \|\bar{s}_i\|_{\mathcal{K}} \right| + \left| \|\bar{s}_i\|_{\mathcal{K}} - \|s_i\|_{\mathcal{H}} \right| + \left| \|s_i\|_{\mathcal{H}} - \|\mathbf{m}\|_{\mathcal{H}} \right|$$

which also goes to 0 in i and therefore $\|\mathbf{n}\|_{\mathcal{K}} = \|\mathbf{m}\|_{\mathcal{H}}$.

The function $k(x, \cdot) : X \rightarrow \mathcal{H}$ is continuous and therefore Bochner-integrable with respect to P . Denote the Bochner integral with $\mathbf{n}' = \int k(x, \cdot) dP$. For any $x \in X$

$$\mathbf{n}(x) = (\mathbf{n}, k(x, \cdot)) = \langle \mathbf{m}, \phi(x) \rangle = E\langle \phi(\cdot), \phi(x) \rangle = Ek(x, \cdot) = (\mathbf{n}', k(x, \cdot)) = \mathbf{n}'(x)$$

and $\mathbf{n} = \mathbf{n}'$.

Now if the algorithm is applied in $(\mathcal{K}, (\cdot, \cdot))$ with the initialization $k(x_0, \cdot)$, where x_0 is the element in \mathcal{X} that maps to the initialization $\phi(x_0)$ that we use in Proposition 6.2, then sequences of elements x_t and of weights w_t are generated. The weights w_t are of the form $k(x_0, \cdot) + \sum_{i=1}^t k(x_i, \cdot) - t\mathbf{n}$. The sequence x_1, x_2, \dots also maximizes the objective in $(\mathcal{H}, \langle \cdot, \cdot \rangle)$. This can be seen by an induction over the weights $\tilde{w}_t \in \mathcal{H}$ that are generated by the algorithm. The induction step is the following.

$$\begin{aligned} \max_{x \in [0,1]} \langle \tilde{w}_t, \phi(x) \rangle &= \max_{x \in [0,1]} \left(\langle \phi(x_0), \phi(x) \rangle + \sum_{i=1}^t \langle \phi(x_i), \phi(x) \rangle - t \langle \mathbf{m}, \phi(x) \rangle \right) \\ &= \max_{x \in [0,1]} \left((k(x_0, \cdot), k(x, \cdot)) + \sum_{i=1}^t (k(x_i, \cdot), k(x, \cdot)) - t(\mathbf{n}, k(x, \cdot)) \right) \\ &= (w_t, k(x_{t+1}, \cdot)) \\ &= (k(x_0, \cdot), k(x_{t+1}, \cdot)) + \sum_{i=1}^t (k(x_i, \cdot), k(x_{t+1}, \cdot)) - t(\mathbf{n}, k(x_{t+1}, \cdot)) \\ &= \left\langle \phi(x_0) + \sum_{i=1}^t \phi(x_i) - t\mathbf{m}, \phi(x_{t+1}) \right\rangle = \langle \tilde{w}_t, \phi(x_{t+1}) \rangle. \end{aligned}$$

From Proposition 6.2 we can now infer that the sequence $\{\|\tilde{w}_t\|_{\mathcal{H}}\}_{t \in \mathbb{N}} = \{\|w_t\|_{\mathcal{K}}\}_{t \in \mathbb{N}}$ is unbounded and the algorithm does not converge with the fast rate in \mathcal{K} . \square

List of Notation

| | | |
|---|---|----|
| \mathbf{m} | Mean embedding | 4 |
| \mathbf{m}_n | Mean embedding of the empirical measure | 4 |
| width | Width of the projection of a convex set C onto a function h | 4 |
| β_l | Lebesgue measure (volume) of the l -dimensional unit ball | 9 |
| $\mathcal{H} \odot \mathcal{H}$ | Symmetric tensor product | 10 |
| $\underline{g} \otimes h$ | Tensor product of functions as a linear operator | 21 |
| \mathfrak{C} | Covariance operator | 21 |
| T_k | Integral operator associated to the kernel k | 22 |
| \mathfrak{C}_c | Centered covariance operator | 22 |
| $\mathcal{H}_1 \oplus \mathcal{H}_2$ | Direct sum of \mathcal{H}_1 and \mathcal{H}_2 | 22 |
| ψ_γ | A continuous function that approximates a step function | 26 |
| \mathcal{H}_S | Restriction to the support S of the data: $\mathcal{H}_S = \{h _S : h \in \mathcal{H}\}$ | 32 |
| $\tilde{\mathfrak{C}}_c^S$ | Centered covariance operator corresponding to \mathcal{H}_S | 33 |
| $K(f, t)$ | K -functional | 36 |
| \mathcal{H}^+ | Extension of \mathcal{H} which contains the constant functions | 40 |
| \mathcal{H}^- | \mathcal{H} without the constant functions | 41 |
| \mathfrak{C}_n | An empirical cov. operator that is represented as a function in $\mathcal{H} \odot \mathcal{H}$ | 64 |
| $C_{\odot, n}$ | Empirical convex set which contains \mathfrak{C}_n | 64 |
| \mathfrak{C} | A covariance operator that is represented as a function in $\mathcal{H} \odot \mathcal{H}$ | 64 |
| C_\odot | Convex set which contains \mathfrak{C} | 64 |
| $\mathbf{m}_{y, n}$ | Empirical mean embedding in the regression setting | 64 |
| $\mathbf{m}_{y, n}^\otimes$ | Empirical mean embedding in the regression setting; tensor representation | 64 |
| $\widehat{h}(y, x)$ | Trivial extension of $h(x)$ to the product space | 65 |
| $\widehat{\mathcal{H}} \odot \widehat{\mathcal{H}}$ | Extension of $\mathcal{H} \odot \mathcal{H}$ corresponding to $h \mapsto \widehat{h}$ | 65 |
| κ_y | Kernel function of $\widehat{\mathcal{H}} \odot \widehat{\mathcal{H}}$ | 65 |
| \mathbf{m}_y | Mean embedding in the regression setting | 70 |
| $l \times k$ | Product kernel for regression where $l = f_0 \otimes f_0$ | 74 |
| \check{h} | Extension of h to the product space: $\check{h} = \langle 1, \cdot \rangle_{\mathbb{R}} \otimes h(\cdot)$ | 78 |
| \mathbf{m}_y^\otimes | Mean embedding in the regression setting; tensor representation | 78 |
| C_y^\otimes | Convex set for regression; tensor representation | 78 |
| $C_{y, n}^\otimes$ | Empirical convex set for regression; tensor representation | 78 |
| $\tilde{\mathfrak{C}}_c^\otimes$ | Centered covariance operator for regression; tensor representation | 79 |
| $\widehat{Y}^{(n)}$ | Capped response variable | 80 |
| $\widehat{\mathbf{m}}_{y, n}^\otimes$ | Empirical mean embedding corresponding to $\widehat{Y}^{(n)}$ | 80 |
| $\widehat{\mathfrak{C}}_c^{\otimes, (n)}$ | Covariance operator corresponding to $\widehat{Y}^{(n)}$ | 83 |
| κ_\oplus | Kernel function used for simultaneous estimation | 88 |

BIBLIOGRAPHY

- [AHPV05] Pankaj K. Agarwal, Sariel Har-Peled, and Kasturi R. Varadarajan, *Geometric approximation via coresets*, Combinatorial and computational geometry, Mathematical Sciences Research Institute Publications, vol. 52, Cambridge University Press, 2005, pp. 1–30. ↑485
- [Aro50] Nachman Aronszajn, *Theory of reproducing Kernels*, Trans. Am. Math. Soc. **68** (1950), 337–404. ↑490, 557, 561, 562
- [BBDG19] François-Xavier Briol, Alessandro Barp, Andrew B. Duncan, and Mark Girolami, *Statistical inference for generative models with maximum mean discrepancy*, 2019, <https://arxiv.org/abs/1906.05944>. ↑529, 530
- [BBZ07] Gilles Blanchard, Olivier Bousquet, and Laurent Zwald, *Statistical properties of kernel principal component analysis*, Mach. Learn. **66** (2007), no. 2-3, 259–294. ↑488, 587
- [BCC⁺22] Gábor Braun, Alejandro Carderera, Cyrille W. Combettes, Hamed Hassani, Amin Karbasi, Aryan Mokhtari, and Sebastian Pokutta, *Conditional gradient methods*, 2022, <https://arxiv.org/abs/2211.14103v1>. ↑569
- [Bee93] Gerald Beer, *Topologies on closed and closed convex sets*, Mathematics and its Applications (Dordrecht), vol. 268, Kluwer Academic Publishers, 1993. ↑494
- [BHPI02] Mihai Bădoiu, Sariel Har-Peled, and Piotr Indyk, *Approximate clustering via coresets*, Proceedings of the thirty-fourth annual ACM symposium on theory of computing (STOC 2002), Montreal, ACM Press, 2002, pp. 250–257. ↑485
- [BLJO12] Francis Bach, Simon Lacoste-Julien, and Guillaume Obozinski, *On the equivalence between herding and conditional gradient algorithms*, ICML 2012 International Conference on Machine Learning, 2012, <https://hal.science/hal-00681128>. ↑483, 484, 520
- [BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart, *Concentration inequalities: A nonasymptotic theory of independence*, Oxford University Press, 2013. ↑550, 551
- [BM03] Peter L. Bartlett and Shahar Mendelson, *Rademacher and Gaussian complexities: risk bounds and structural results*, J. Mach. Learn. Res. **3** (2003), no. 3, 463–482. ↑496, 498
- [BT04] Amir Beck and Marc Teboulle, *A conditional gradient method with linear rate of convergence for solving convex linear systems*, Math. Methods Oper. Res. **59** (2004), no. 2, 235–247. ↑575
- [CM09] Stephen L. Campbell and Carl D. Meyer, *Generalized inverses of linear transformations*, Classics in Applied Mathematics, vol. 56, Society for Industrial and Applied Mathematics, 2009. ↑518
- [CWS10] Yutian Chen, Max Welling, and Alexander Smola, *Supersamples from kernel-herding*, Proceedings of the Conference on Uncertainty in Artificial Intelligence, 2010, <https://arxiv.org/abs/1203.3472>. ↑569
- [DF93] Andreas Defant and Klaus Floret, *Tensor norms and operator ideals*, North-Holland Mathematics Studies, vol. 176, North-Holland, 1993. ↑490, 492
- [DM21] Raaz Dwivedi and Lester Mackey, *Kernel thinning*, 2021, <https://arxiv.org/abs/2105.05842>. ↑486
- [DU77] Joseph Diestel and J. Jerry jun. Uhl, *Vector measures*, Mathematical Surveys, vol. 15, American Mathematical Society, 1977. ↑490, 505
- [Dud14] Richard M. Dudley, *Uniform central limit theorems*, 2nd ed., Cambridge Studies in Advanced Mathematics, vol. 142, Cambridge University Press, 2014. ↑488, 489, 495, 551
- [Eng89] Ryszard Engelking, *General topology*, revised and completed ed., Sigma Series in Pure Mathematics, vol. 6, Heldermann Verlag, 1989. ↑505, 553, 559
- [Fre01] David H. Fremlin, *Measure theory*, Torres Fremlin, 2001. ↑488, 489, 502, 522, 525, 526, 530, 531, 551, 552, 553

- [FW56] Marguerite Frank and Philip Wolfe, *An algorithm for quadratic programming*, Nav. Res. Logist. Q. **3** (1956), no. 1-2, 95–110. ↑485
- [FW95] Sally Floyd and Manfred K. Warmuth, *Sample compression, learnability, and the Vapnik–Chervonenkis dimension*, Mach. Learn. **21** (1995), 269–304. ↑486
- [GBR⁺12] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola, *A Kernel two-sample test*, J. Mach. Learn. Res. **13** (2012), 723–773. ↑488, 584
- [GN16] Evarist Giné and Richard Nickl, *Mathematical foundations of infinite-dimensional statistical models*, Cambridge Series in Statistical and Probabilistic Mathematics, vol. 40, Cambridge University Press, 2016. ↑489, 495, 496, 497, 543, 547
- [HCB16] Jonathan H. Huggins, Trevor Campbell, and Tamara Broderick, *Coresets for scalable bayesian logistic regression*, Proceedings of the 30th International Conference on Neural Information Processing Systems, Advances in Neural Information Processing Systems, vol. 29, Curran Associates, Inc., 2016, pp. 4080–4088. ↑485
- [HS14] Nick Harvey and Samira Samadi, *Near-optimal herding*, Proceedings of The 27th Conference on Learning Theory, Proceedings of Machine Learning Research, vol. 35, PMLR, 2014, pp. 1165–1182. ↑486
- [Kol18] Vladimir Koltchinskii, *Asymptotic efficiency in high-dimensional covariance estimation*, Proceedings of the international congress of mathematicians 2018, ICM 2018. Volume IV. Invited lectures, World Scientific; Sociedade Brasileira de Matemática, 2018, pp. 2903–2923. ↑587
- [LLW23] Jian Li, Yong Liu, and Weiping Wang, *Optimal convergence rates for agnostic Nyström kernel learning*, Proceedings of the 40th International Conference on Machine Learning, Proceedings of Machine Learning Research, vol. 202, PMLR, 2023, pp. 19811–19836. ↑487, 585, 586
- [LT91] Michel Ledoux and Michel Talagrand, *Probability in Banach spaces. Isoperimetry and processes*, Classics in Mathematics, Springer, 1991. ↑496
- [LW86] Nick Littlestone and Manfred K. Warmuth, *Relating data compression and learnability*, Tech. report, University of California at Santa Cruz, 1986. ↑486
- [MFSS17] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf, *Kernel mean embedding of distributions: A review and beyond*, Found. Trends Mach. Learn. **10** (2017), no. 1-2, 1–144. ↑584
- [Mur90] Gerard J. Murphy, *C*-algebras and operator theory*, Academic Press Inc., 1990. ↑490
- [Pin94] Iosif Pinelis, *Optimum bounds for the distributions of martingales in Banach spaces*, Ann. Probab. **22** (1994), no. 4, 1679–1706. ↑541
- [Pin99] ———, *Correction : optimum bounds for the distributions of martingales in Banach spaces*, Ann. Probab. **27** (1999), no. 4, 2119–2119. ↑541
- [PR16] Vern I. Paulsen and Mrinal Raghupathi, *An introduction to the theory of reproducing kernel Hilbert spaces*, Cambridge Studies in Advanced Mathematics, vol. 152, Cambridge University Press, 2016. ↑488, 491, 510, 525, 532, 534, 536, 544, 545, 560, 565
- [PT20] Jeff M. Phillips and Wai Ming Tai, *Near-optimal coresets of kernel density estimates*, Discrete Comput. Geom. **63** (2020), no. 4, 867–887. ↑483, 484
- [RCR15] Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco, *Less is more: Nyström computational regularization*, Advances in Neural Information Processing Systems, vol. 28, Curran Associates, Inc., 2015, pp. 1657–1665. ↑487
- [Roc97] R. Tyrrell Rockafellar, *Convex analysis*, 2nd ed., Princeton Mathematical Series, vol. 28, Princeton University Press, 1997. ↑519
- [RS72] Michael Reed and Barry Simon, *Methods of modern mathematical physics. 1: Functional analysis*, Academic Press Inc., 1972. ↑492, 561, 564

- [SC08] Ingo Steinwart and Andreas Christmann, *Support vector machines*, Information Science and Statistics, Springer, 2008. ↑488, 507, 521
- [Sim11] Barry Simon, *Convexity: an analytic viewpoint*, Cambridge Tracts in Mathematics, vol. 187, Cambridge University Press, 2011. ↑600
- [SSM97] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller, *Kernel principal component analysis*, Artificial Neural Networks — ICANN'97, Lecture Notes in Computer Science, vol. 1327, Springer, 1997, pp. 583–588. ↑488, 587
- [SSRR20] Nicholas Sterge, Bharath Sriperumbudur, Lorenzo Rosasco, and Alessandro Rudi, *Gain with no pain: efficiency of kernel-PCA by Nyström sampling*, Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research, vol. 108, PMLR, 2020, pp. 3642–3652. ↑488, 587, 588
- [Tem11] Vladimir N. Temlyakov, *Greedy approximation*, Cambridge Monographs on Applied and Computational Mathematics, vol. 20, Cambridge University Press, 2011. ↑485
- [Wer02] Dirk Werner, *Funktionalanalysis*, 4th ed., Springer-Lehrbuch, Springer, 2002. ↑505
- [WS00] Christopher K. I. Williams and Matthias Seeger, *Using the Nyström method to speed up kernel machines*, Proceedings of the 14th International Conference on Neural Information Processing Systems, MIT Press, 2000, p. 661–667. ↑487

Manuscript received on 27th August 2024,
revised on 12th December 2025,
accepted on 2nd March 2026.

Recommended by Editors S. Gouëzel and É. Gassiat.
Published under license CC BY 4.0.



eISSN: 2644-9463

This journal is a member of Centre Mersenne.



Steffen GRÜNEWÄLDER
University of York
Department of Mathematics
York (United Kingdom)
gruNEW@gmail.com