
Sampling based approximation of linear functionals in Reproducing Kernel Hilbert Spaces

Gabriele Santin · Toni Karvonen · Bernard Haasdonk

<https://doi.org/10.1007/s10543-021-00870-3>

Abstract In this paper we analyze a greedy procedure to approximate a linear functional defined in a Reproducing Kernel Hilbert Space by nodal values. This procedure computes a quadrature rule which can be applied to general functionals.

For a large class of functionals, that includes integration functionals and other interesting cases, but does not include differentiation, we prove convergence results for the approximation by means of quasi-uniform and greedy points which generalize in various ways several known results. A perturbation analysis of the weights and node computation is also discussed.

Beyond the theoretical investigations, we demonstrate numerically that our algorithm is effective in treating various integration densities, and that it is even very competitive when compared to existing methods for Uncertainty Quantification.

Keywords kernel methods · greedy algorithms · optimal quadrature

Mathematics Subject Classification (2000) 65D05 · 65D15 · 65D30 · 65D32 · 62C10 · 41A05 · 41A25

1 Introduction

Given a strictly positive definite kernel $K : \Omega \times \Omega \rightarrow \mathbb{R}$ on a bounded, measurable set $\Omega \subset \mathbb{R}^d$ and a linear and continuous functional $L \in \mathcal{H}'$ in the dual of the associated reproducing kernel Hilbert space \mathcal{H} , we are interested in the construction

G. Santin
Center for Information and Communication Technology, Fondazione Bruno Kessler, Italy
E-mail: gsantin@fbk.eu

T. Karvonen
The Alan Turing Institute, United Kingdom
E-mail: tkarvonen@turing.ac.uk

B. Haasdonk
Institute for Applied Analysis and Numerical Simulation, University of Stuttgart, Germany
E-mail: haasdonk@mathematik.uni-stuttgart.de

of quadrature-like formulas (or nodal approximants) that approximate $L(f)$ for all $f \in \mathcal{H}$. This means that we look for pairwise distinct centers $X := \{x_i\}_{i=1}^n \subset \Omega$ and weights $W := (w_i)_{i=1}^n \in \mathbb{R}^n$ such that

$$Q_{X,W,L}(f) := \sum_{i=1}^n w_i f(x_i) \approx L(f) \quad \text{for all } f \in \mathcal{H}.$$

We measure the approximation quality of $Q_{X,W,L}$ by means of the worst case error on the unit ball of \mathcal{H} , i.e.,

$$e_{\mathcal{H}}(Q_{X,W,L}) := \sup_{\|f\|_{\mathcal{H}} \leq 1} |Q_{X,W,L}(f) - L(f)|.$$

In terms of this error measure, for a given set X there exist optimal weights $W^* := W^*(X)$, i.e.,

$$W^* := \arg \min_{W \in \mathbb{R}^n} e_{\mathcal{H}}(Q_{X,W,L}),$$

and we use the notation $Q_{X,L} := Q_{X,W^*,L}$ for the weight-optimal quadrature formula.

We will discuss in the following (see Proposition 2.3) how these weights can be computed explicitly, but here we anticipate in particular that, if $v_L \in \mathcal{H}$ is the Riesz representer of L , and if $\Pi_X(v_L)$ is the \mathcal{H} -orthogonal projection of v_L into $V(X) := \text{span}\{K(\cdot, x) : x \in X\} \subset \mathcal{H}$, then it holds that

$$\Pi_X v_L = \sum_{i=1}^n w_i^* K(\cdot, x_i)$$

and $e_{\mathcal{H}}(Q_{X,L}) = \|v_L - \Pi_X v_L\|_{\mathcal{H}}$.

Since it is well known that $\Pi_X v_L$ coincides with the interpolant of v_L on the points X , this means that the optimal weights can be easily computed via standard kernel-based interpolation, and the worst-case error coincides with the \mathcal{H} -norm interpolation error of v_L on X .

Assuming that these optimal weights are used, the question of selecting the centers remains open, and the goal of this paper is to analyze a particularly simple greedy algorithm to do so. The algorithm has been introduced in [37], although we give here a more explicit characterization. It starts from an empty set of centers and iteratively chooses a point among the ones that guarantee the maximal reduction of the worst-case error. This allows to distribute adaptive and possibly non uniform centers tailored to the specific L , and this property is particularly attractive especially in high dimensions.

We show that the algorithm is in fact the f/P -greedy algorithm of [24] known in kernel interpolation, and applied to the Riesz representer v_L of L . In particular we recall how it can be efficiently described and implemented in terms of the Newton basis as in [25, 31].

We then prove two types of error estimates for the approximation of a special class of functionals $L \in \mathcal{H}'$, namely those which are continuous on \mathcal{H} also with

respect to the $L_q(\Omega)$ norm for some $1 \leq q \leq \infty$, i.e., such that there exists $1 \leq q \leq \infty$ and $c_L \geq 0$ with

$$|L(f)| \leq c_L \|f\|_{L_q(\Omega)} \quad \text{for all } f \in \mathcal{H}. \quad (1.1)$$

Observe that this characterization rules out differentiation operators.

First, for certain translational invariant kernels we provide convergence orders for weight-optimal quadrature rules with quasi-uniform sets of centers. These results are a direct consequence of the error rates known for kernel interpolation, and the resulting speed of convergence depends on the input dimension d , the value of q , and the smoothness of the kernels. In particular, smoother kernels lead to faster convergence. For some specific functionals, which are included in our analysis, this result coincides with the ones of [18].

Second, for fairly general kernels we prove convergence with rate $n^{-1/2}$ for the new greedy algorithm, where n is the number of centers. Although various experiments suggest that this rate is far from optimal, it is nevertheless dimension independent and it applies to a wide class of kernels, namely, continuous kernels on bounded domains. Moreover, also for translational invariant kernels this rate is strictly better than the one for quasi-uniform points for a range of values of d , q , and of the smoothness of the kernel. This gap between the two rates is quite small, but it holds for a wider range of kernels for increasing d if q is fixed.

The motivation for the study of the class of functionals (1.1) comes from integration functionals $L(f) := \int_{\Omega} f(x)g_L(x)dx$, where $g_L \in L_p(\Omega)$ for some $1 \leq p \leq \infty$. Indeed, in this case we can take q such that $\frac{1}{p} + \frac{1}{q} = 1$ and $c_L := \|g_L\|_{L_p(\Omega)}$, since

$$|L(f)| \leq \int_{\Omega} |f(x)| |g_L(x)| dx \leq \|g_L\|_{L_p(\Omega)} \|f\|_{L_q(\Omega)}.$$

In this case $Q_{X,L}$ is a quadrature formula in the classical sense. Observe in particular that this class includes also the case of $g_L \in L_p(\Omega) \setminus L_{\infty}(\Omega)$ for some $1 \leq p < \infty$, which is not covered in [18].

However, our analysis comprises other interesting examples that will be discussed in Section 2. In particular, we can consider functionals L which represent any quadrature rule with bounded weights, including Monte Carlo ones. In this case, constructing a quadrature rule $Q_{X,L}$ via the greedy algorithm means to find an approximated quadrature with possibly fewer centers, or a compression of the quadrature L .

Moreover, given the equivalence between the weight-optimal quadrature of L and the interpolation of v_L , and the equivalence between the new greedy algorithm for quadrature and the f/P -greedy algorithm for interpolation, our analysis alternatively applies to the interpolation via f/P -greedy of the class of functions that are Riesz representers of functionals satisfying the condition (1.1). The convergence results will then be also be convergence results in \mathcal{H} for interpolation. These results are potentially very interesting, as we remark that for general functions $v \in \mathcal{H}$, the error $\|v - \Pi_X v\|_{\mathcal{H}}$ can decay arbitrarily slowly even for nicely chosen X (see e.g. [15, Section 8.4.2]). Instead, for this class of functions, rates of convergence are obtained here for interpolation with both quasi-uniform and greedy points.

Several examples of functionals in this class are discussed in Section 2.

On the computational side we remark that both the selection of the greedy points and the computation of the optimal weights can be performed by the sole knowledge of the Riesz representer v_L . As we will recall, this can be computed rather efficiently and explicitly and, when it needs instead to be approximated, we give stability bounds on the resulting perturbed quadrature formula. This easiness of computation is very advantageous if the quadrature rule is then applied to the approximation of $L(f)$ for a given f whose evaluation is expensive. This is the case for example for Uncertainty Quantification, where an integration functional is used to estimate the mean and variance of f . Alternatively, when only few evaluations are available it can be advantageous to leverage the well-known equivalence between kernel interpolation and Gaussian process regression [9, Chapter 17] and view $Q_{X,L}(f)$ as a Gaussian random variable whose standard deviation, equal to the worst case error, attempts to quantify the epistemic uncertainty in the approximation $Q_{X,L}(f) \approx L(f)$. This approach has been especially popular in integration, where it is known as the Bayesian quadrature [5, 23, 30]. The greedy algorithm we study here is sometimes called the sequential Bayesian quadrature in this context [14].

We mention also that other data-based algorithms to approximate linear functionals exists in different settings, and are actively investigated e.g. in the setting of empirical interpolation and reduced order modelling [1, 6, 12, 46].

The paper is structured as follows. In Section 2 we recall some basic facts on kernel spaces, list some properties of linear and continuous functionals on these spaces, and recall the computation and properties of weight-optimal quadrature rules and their connection with interpolation. Section 4 introduces the greedy algorithm and discusses its equivalence with the f/P -greedy algorithm for interpolation. The convergence results for quasi-uniform points and translational invariant kernels are discussed in Section 3, while the ones for greedy points and general kernels are presented in Section 5. Some stability results are shown in Section 6, providing in particular bounds on the worst case error of a quadrature rule obtained from a perturbed Riesz representer. Finally, the greedy method is tested on both synthetic examples and on a benchmark problem in Uncertainty Quantification in Section 7.

2 Kernels and approximation of linear functionals

We recall some basics of kernel theory, and for a more general treatment we refer e.g. to [10, 42]. A strictly positive definite (s.p.d.) kernel on $\Omega \subset \mathbb{R}^d$ is a symmetric function $K: \Omega \times \Omega \rightarrow \mathbb{R}$ such that for all $n \in \mathbb{N}$ and for all sets $X_n := \{x_i\}_{i=1}^n \subset \Omega$ of pairwise distinct points, the kernel matrix $A := A(K, X_n) \in \mathbb{R}^{n \times n}$ defined by $A_{ij} := K(x_i, x_j)$ is positive definite. We assume this condition in the following, and additionally that $\Omega \subset \mathbb{R}^d$ is bounded and Lebesgue measurable, and K is continuous in both variables.

Each s.p.d. kernel is uniquely associated to a reproducing kernel Hilbert space (RKHS) $\mathcal{H} := \mathcal{H}_K(\Omega)$ with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$, which is usually called native space of K on Ω , and which is a Hilbert space of functions $f: \Omega \rightarrow \mathbb{R}$ such that $K(\cdot, x) \in \mathcal{H}$ for all $x \in \Omega$ and $\langle f, K(\cdot, x) \rangle_{\mathcal{H}} = f(x)$ for all $x \in \Omega$ and $f \in \mathcal{H}$.

These two properties mean in particular that the function $v_{\delta_x}(\cdot) := K(\cdot, x)$ is an element of \mathcal{H} for all $x \in \Omega$, and that it is the Riesz representer of the linear functional $\delta_x : \mathcal{H} \rightarrow \mathbb{R}$, $\delta_x(f) := f(x)$, which is thus continuous, i.e., $\delta_x \in \mathcal{H}'$. Actually also the converse holds, i.e., any Hilbert space \mathcal{H} of functions $f : \Omega \rightarrow \mathbb{R}$ where the set $\{\delta_x : x \in \Omega\}$ is contained in \mathcal{H}' is an RKHS, and the corresponding kernel is strictly positive definite whenever the elements of this set are also linearly independent.

Observe that for these particular functionals the Riesz representer v_{δ_x} can be obtained by applying the functional δ_x to one of the two variables of the kernel, i.e., $v_{\delta_x} = K(\cdot, x) = (\delta_x)^y(K(\cdot, y))$, where the upper index denotes the variable with respect to which the functional is applied. This is actually always the case, as we recall in the next proposition.

Proposition 2.1 (Riesz representer [42, Theorem 16.7]) *Let $L \in \mathcal{H}'$ be a linear and continuous functional on \mathcal{H} . Then the Riesz representer $v_L \in \mathcal{H}$ of L is given by $v_L(\cdot) = L^y(K(\cdot, y))$.*

Strictly positive definite kernels allow especially to solve interpolation problems, as stated in the following proposition, which is a collection of various classical results (see e.g. [42])

Proposition 2.2 (Kernel interpolation) *For an s.p.d. kernel K and a set $X := \{x_i\}_{i=1}^n \subset \Omega$ of pairwise distinct points, we denote as $V(X) := \text{span}\{K(\cdot, x) : x \in X\}$ the subspace of \mathcal{H} spanned by the kernel translates at X , and as $\Pi_X : \mathcal{H} \rightarrow V(X)$ the \mathcal{H} -orthogonal projector into $V(X)$.*

Then for each $v \in \mathcal{H}$ the function $\Pi_X v$ interpolates v at X . Moreover, we have

$$\Pi_X v = \sum_{i=1}^n \alpha_i K(\cdot, x_i), \quad (2.1)$$

where $\alpha \in \mathbb{R}^n$ is the unique solution of the linear system $A\alpha = (v(x_1), \dots, v(x_n))^T$ with A the kernel matrix of K on X .

This proposition implies in particular that the interpolant of an arbitrary function in \mathcal{H} can be computed on arbitrary pairwise distinct points. Moreover, this interpolant coincides with the orthogonal projection into $V(X)$, and it is thus an optimal approximant with respect to the \mathcal{H} -norm. As mentioned in the introduction, these properties can be translated to the approximation of a linear functional. Indeed, if $L \in \mathcal{H}'$ and $X := \{x_i\}_{i=1}^n \subset \Omega$ are pairwise distinct points, for a given set of weights $W := (w_i)_{i=1}^n \in \mathbb{R}^n$ we can define

$$Q_{X,W,L}(f) := \sum_{i=1}^n w_i f(x_i), \quad f \in \mathcal{H},$$

and the corresponding worst-case error $e_{\mathcal{H}}(Q_{X,W,L}) := \sup_{\|f\|_{\mathcal{H}} \leq 1} |Q_{X,W,L}(f) - L(f)|$.

Now L has a representer $v_L \in \mathcal{H}$ so that $Lf = \langle f, v_L \rangle_{\mathcal{H}}$ and by the Riesz representation theorem, which provides an isometry between \mathcal{H} and \mathcal{H}' , we have

$$\text{dist}_{\mathcal{H}'}(L, \text{span}\{\delta_{x_j} : x_j \in X\}) = \text{dist}_{\mathcal{H}}(v_L, V(X)),$$

and thus the \mathcal{H}^l -optimal quadrature formula for samples on $X \subset \Omega$ is given by the \mathcal{H} -orthogonal projector Π_X onto $V(X)$, which happens to be the interpolation operator. The weight-optimal quadrature formula is thus

$$Q_{X,L}(f) = \langle f, \Pi_X v_L \rangle_{\mathcal{H}} = L(\Pi_X f).$$

This proves the following proposition.

Proposition 2.3 (Weights-optimal quadrature) *There exist unique weights $W^* = (w_1^*, \dots, w_n^*)^T$ that minimize the worst case error given L and X , i.e.,*

$$W^* := \arg \min_{W \in \mathbb{R}^n} e_{\mathcal{H}}(Q_{X,W,L}), \quad (2.2)$$

and they are the coefficients of the orthogonal projection of v_L into $V(X)$, i.e.,

$$\Pi_X v_L = \sum_{i=1}^n w_i^* K(\cdot, x_i). \quad (2.3)$$

Moreover, $\Pi_X v_L$ is the Riesz representer of the functional $Q_{X,L} := Q_{X,W^*,L} \in \mathcal{H}'$, it holds

$$Q_{X,L}(f) = L(\Pi_X f) \quad \text{for all } f \in \mathcal{H}, \quad (2.4)$$

and

$$e_{\mathcal{H}}(Q_{X,L}) = \|v_L - \Pi_X v_L\|_{\mathcal{H}}. \quad (2.5)$$

Observe that this proposition implies that in practice the quadrature weights can be found just by computing the interpolant of v_L at the points X , and, according to Proposition 2.2, this corresponds to the solution of a linear system. Moreover, (2.4) proves that applying the quadrature formula is equivalent to exactly applying L to the interpolant, or that $Q_{X,L}$ is exact on $V(X)$.

Remark 2.1 (Positive definite kernels) The results of this section can be formulated also for positive definite kernels, i.e., those for which the kernel matrix is required to be only positive semidefinite. However, in this case some complications arise since the kernel matrix can be singular also for pairwise distinct points X , and in particular the elements $K(\cdot, x_i)$ do not need to be linearly independent, and thus they span $V(X)$ without being a basis. Nevertheless, the same results can be derived if more attention is paid, for example by showing that different representations (2.1) can describe a unique function. More importantly, we do not explicitly extend the current presentation to positive definite kernels since it is not clear if the greedy algorithm of the next section, which is the main topic of this paper, can be run without producing singular matrices (and thus an early termination) in the case of (non strictly) positive definite kernels.

Although the construction of $Q_{X,L}$ and the greedy algorithm that we will introduce work for any $L \in \mathcal{H}'$, we recall that our error analysis will apply only to functionals that satisfy (1.1) for some $1 \leq q \leq \infty$ and $c_L \geq 0$. Observe that this definition is well posed, since $\mathcal{H} \subset L_q(\Omega)$ for all $1 \leq q \leq \infty$ and for all $f \in \mathcal{H}$ since Ω is assumed to be bounded and K continuous.

Now that the equivalence between worst-case quadrature with optimal weights and interpolation has been detailed, we conclude this section with a more precise discussion of some relevant examples of functionals satisfying the condition (1.1).

Example 2.1 If $L(f) := \int_{\Omega} f(x)g_L(x)dx$ with $g_L \in L_p(\Omega)$ for some $1 \leq p \leq \infty$, we can take q such that $\frac{1}{p} + \frac{1}{q} = 1$ and we have $c_L := \|g_L\|_{L_p(\Omega)}$ since

$$|L(f)| \leq \int_{\Omega} |f(x)| |g_L(x)| dx \leq \|g_L\|_{L_p(\Omega)} \|f\|_{L_q(\Omega)}.$$

In this case $Q_{X,L}$ is a classical quadrature rule.

Example 2.2 If $L(f) := \sum_{i \in I} \rho_i f(z_i)$ where $I \subset \mathbb{N}$ is a countable index set, $\{\rho_i\}_{i \in I} \subset \mathbb{R}$, $\{z_i\}_{i \in I} \subset \Omega$, and if $\rho := \{\rho_i\}_{i \in I} \in \ell_1(I)$, then we can take $q := \infty$ and $c_L := \|\rho\|_{\ell_1(I)} = \sum_{i \in I} |\rho_i|$ since

$$|L(f)| \leq \sum_{i \in I} |\rho_i| |f(z_i)| \leq \left(\max_{i \in I} |f(z_i)| \right) \|\rho\|_{\ell_1(I)} \leq \|f\|_{L_{\infty}(\Omega)} \|\rho\|_{\ell_1(I)}.$$

This includes for example any quadrature formula with weights $\rho := \{\rho_i\}_{i \in I} \in \ell_1(I)$ and nodes $\{z_i\}_{i \in I} \subset \Omega$. This is the case for example of quadrature rules with positive and bounded weights, and in particular of any Monte Carlo quadrature with $M := |I|$, since $\sum_{i \in I} \rho_i = \sum_{i=1}^M |\Omega|/M = |\Omega|$.

In this case $Q_{X,L}$ can be understood as a compression of the quadrature rule, if $|X| \leq |I|$. Or, even for $|X| = |I|$, $Q_{X,L}$ is weight-optimal and thus can provide a strictly better worst-case error than L .

Considering instead functions which are Riesz representers of functionals satisfying (1.1), we have the following examples.

Example 2.3 Given a number $B > 0$, the functions in the class

$$\mathcal{H}_B := \left\{ v := \sum_{i \in I} \alpha_i K(\cdot, x_i) : \sum_{i \in I} |\alpha_i| \leq B \right\} \subset \mathcal{H}$$

are the Riesz representers of functionals of Example 2.2 with $\|\rho\|_{\ell_1(I)} \leq B$. This set is commonly used to study convergence rates of greedy algorithms (see e.g. [2, 7, 41]), and for this set our results on the greedy algorithm coincide with the rates obtained in [44] for the f/P -greedy algorithm.

Example 2.4 Under the present assumptions (see [40] for a general analysis) the kernel is a Mercer kernel, and it can be proven (see e.g. [42, Chapter 10]) that the operator $T : L_2(\Omega) \rightarrow L_2(\Omega)$ given by

$$T(f) := \int_{\Omega} K(x, y) f(y) dy,$$

is compact and self adjoint. It has a sequence $\{\lambda_j\}_{j \in \mathbb{N}}$ of non increasing and positive eigenvalues and corresponding $L_2(\Omega)$ -orthonormal eigenvectors $\{\varphi_j\}_{j \in \mathbb{N}}$ such that $\{\lambda_j^{1/2} \varphi_j\}_{j \in \mathbb{N}}$ is an \mathcal{H} -orthonormal basis of \mathcal{H} . The image $I_T := T(L_2(\Omega))$ is dense in \mathcal{H} , and every function $v \in I_T$ is the Riesz representer of a functional L which satisfies (1.1). Indeed, if $u \in L_2(\Omega)$ is such that $T(u) = v$, it can be proven that for all $f \in \mathcal{H}$ it holds $\langle f, T(u) \rangle_{\mathcal{H}} = \langle f, u \rangle_{L_2(\Omega)}$, and thus

$$\begin{aligned} |L(f)| &:= |\langle f, v \rangle_{\mathcal{H}}| = |\langle f, T(u) \rangle_{\mathcal{H}}| = |\langle f, u \rangle_{L_2(\Omega)}| \\ &\leq \|u\|_{L_2(\Omega)} \|f\|_{L_2(\Omega)} \quad \text{for all } f \in \mathcal{H}. \end{aligned}$$

It follows that (1.1) holds with $c_L := \|u\|_{L_2(\Omega)}$ and $q := 2$.

These functions are the easiest example of the class analyzed in [36, 38] to study superconvergence phenomena in \mathcal{H} , i.e., functions for which kernel interpolation with quasi-uniform points leads to an improved convergence order.

We conclude this general section by stating in the following proposition an obvious fact that will be useful later.

Proposition 2.4 (Restriction of L) *Assume that $L \in \mathcal{H}'$, and assume that there exists $1 \leq q \leq \infty$ such that L is continuous w.r.t. the L_q -norm on $L_q(\Omega) \cap \mathcal{H}$, with norm bounded by c_L . Then for any \mathcal{H} -closed subspace $V \subset \mathcal{H}$ also the functional $L_V := L \circ \Pi_V$ is continuous on $L_q(\Omega) \cap \mathcal{H}$, with norm $c_{L_V} \leq c_L$.*

Proof Since L is continuous from $L_q(\Omega) \cap \mathcal{H}$ to \mathbb{R} with constant c_L , i.e.,

$$\sup_{f \in \mathcal{H}, f \neq 0} \frac{|L(f)|}{\|f\|_{L_q(\Omega)}} = c_L,$$

then

$$\begin{aligned} c_{L_V} &:= \sup_{f \in \mathcal{H}, f \neq 0} \frac{|L_V(f)|}{\|f\|_{L_q(\Omega)}} = \sup_{f \in \mathcal{H}, f \neq 0} \frac{|L(\Pi_V(f))|}{\|f\|_{L_q(\Omega)}} \\ &= \sup_{f \in V, f \neq 0} \frac{|L(f)|}{\|f\|_{L_q(\Omega)}} \leq \sup_{f \in \mathcal{H}, f \neq 0} \frac{|L(f)|}{\|f\|_{L_q(\Omega)}} = c_L, \end{aligned}$$

which is the desired bound.

3 Convergence for quasi-uniform points and translational invariant kernels

We first analyze convergence rates for quadrature formulas that use quasi-uniform points on spaces generated by translational invariant kernels. In this section we thus assume that $K(x, y) := \phi(x - y)$ for some $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$, and that ϕ has a Fourier transform $\hat{\phi}$ such that there exists $\tau > 0$, and $c, C > 0$ with

$$c \left(1 + \|\omega\|_2^2\right)^{-\tau} \leq \hat{\phi}(\omega) \leq C \left(1 + \|\omega\|_2^2\right)^{-\tau} \quad \text{for all } \omega \in \mathbb{R}^d. \quad (3.1)$$

If additionally Ω has a Lipschitz boundary and it satisfies an interior cone condition, then \mathcal{H} is norm equivalent to the Sobolev space $W_2^\tau(\Omega)$, and in particular there exists a constant $c_E > 0$ such that

$$\|u\|_{W_2^\tau(\Omega)} \leq c_E \|u\|_{\mathcal{H}} \quad \text{for all } u \in \mathcal{H}. \quad (3.2)$$

Observe that this norm equivalence is possible only if $\tau > d/2$, since it implies in particular that $W_2^\tau(\Omega)$ is an RKHS (see e.g. Chapter 10 in [42]). In fact, in the following we only need the continuous embedding of \mathcal{H} into $W_2^\tau(\Omega)$ to hold with a norm inequality (3.2), and for this it is sufficient to assume the validity of the upper bound in (3.1).

For functions in Sobolev spaces we can use the following inequality from [26, Theorem 1.1], that bounds the L_q -norm of a function with scattered zeros. Observe that bounds for general Sobolev norms are also proven in the same Theorem 1.1.

Here and in the following we denote $(x)_+ := \max(x, 0)$ for $x \in \mathbb{R}$ and, for $u : \Omega \rightarrow \mathbb{R}$, $\|u\|_{\ell_\infty(X)}$ denotes the maximum absolute value of u evaluated on $X \subset \Omega$. Moreover, we use the fill distance h_X and the separation distance q_X , defined by

$$h_X := \sup_{x \in \Omega} \min_{y \in X} \|x - y\|_2, \quad q_X := \frac{1}{2} \min_{x_i \neq x_j, x_i, x_j \in X} \|x_i - x_j\|_2,$$

to quantify the distribution of the points X in Ω .

Theorem 3.1 (Norm inequality [26]) *Let $\Omega \subset \mathbb{R}^d$ be bounded and satisfy an interior cone condition. Let $1 \leq q \leq \infty$ and $\tau \in \mathbb{R}$ be such that $\tau > d/2$. Then there exist $c_S > 0$ and $h_0 > 0$ such that if $X \subset \Omega$ is finite and $h_X \leq h_0$, then for all $u \in W_2^\tau(\Omega)$ with $u = 0$ on X it holds*

$$\|u\|_{L_q(\Omega)} \leq c_S h_X^{\tau - d(1/2 - 1/q)_+} |u|_{W_2^\tau(\Omega)}. \quad (3.3)$$

For general point sets, a geometric constraint implies that there exists a constant $c > 0$ depending only on Ω such that $h_X \geq cn^{-1/d}$. If one uses a quasi-uniform sequence $\{X_n\}_{n \in \mathbb{N}} \subset \Omega$, $|X_n| = n$, of points, i.e., such that there exists a constant $c > 0$ with $h_{X_n} \leq cq_{X_n}$ for all $n \in \mathbb{N}$, then it can be proven that there exists a second constant c_Q independent of n , and an index $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$ it holds $h_{X_n} \leq c_Q n^{-1/d}$. We refer for example to Chapter 2 in [24] for explicit estimates of these constants. Moreover, observe that this requirement applies to sequences of scattered points, i.e., the sets $\{X_n\}_{n \in \mathbb{N}}$ do not need to have any grid structure. Using a sequence of quasi-uniform sets allows to rewrite any bound expressed in terms of h_X as a bound involving only the number of points $n = |X|$.

With these tools in hand we can prove the following result. The proof follows a standard procedure used in combination with a sampling inequality to derive an error bound, with the only difference that the L_q -continuity of L will guarantee that the error bound is in the \mathcal{H} -norm. We remark that the same idea has been used in [3, 17, 18], as well as most other work on error estimates for kernel and Bayesian quadrature rules, to derive error bounds for the integration functionals of Example 2.1 with $g_L \in L_\infty(\Omega)$.

Theorem 3.2 (Convergence rates for quasi-uniform points) *Assume that the kernel is translational invariant and that it has a Fourier transform that satisfies the upper bound in (3.1). Under the assumptions of Theorem 3.1, let $L \in \mathcal{H}'$ be a linear functional such that there exist $1 \leq q \leq \infty$ and $c_L \geq 0$ such that*

$$|L(f)| \leq c_L \|f\|_{L_q(\Omega)} \quad \text{for all } f \in \mathcal{H},$$

and let $v_L \in \mathcal{H}$ be its Riesz representer.

Then, if $X \subset \Omega$ is a set of pairwise distinct points with $h_X \leq h_0$, it holds

$$e_{\mathcal{H}}(Q_{X,L}) = \|v_L - \Pi_X v_L\|_{\mathcal{H}} \leq c_{SC} c_L h_X^{\tau-d(1/2-1/q)_+}. \quad (3.4)$$

In particular, if $\{X_n\}_{n \in \mathbb{N}} \subset \Omega$ is a sequence of sets of pairwise distinct points for which there exists $c_Q > 0$ such that $h_{X_n} \leq c_Q n^{-1/d}$ for all $n \in \mathbb{N}$, then for all $n \in \mathbb{N}$ with $n \geq n_0 := (c_Q/h_0)^d$ it holds

$$\|v_L - \Pi_{X_n} v_L\|_{\mathcal{H}} \leq c_U n^{-\frac{\tau}{d} + \left(\frac{1}{2} - \frac{1}{q}\right)_+}, \quad (3.5)$$

where the constant $c_U := c_{LCS} c_E c_Q^{\tau-d(1/2-1/q)_+}$ depends on L only via q and c_L .

Proof For any $f \in \mathcal{H}$, equation (2.4) and the continuity of L on $L_q(\Omega)$ give

$$\begin{aligned} |L(f) - Q_{X,L}(f)| &= |L(f) - L(\Pi_X f)| = |L(f - \Pi_X f)| \\ &\leq c_L \|f - \Pi_X f\|_{L_q(\Omega)}. \end{aligned}$$

Now, since $u := f - \Pi_X f$ vanishes on X , the bound of Theorem 3.1 and the norm inequality (3.2) give

$$\begin{aligned} |L(f) - Q_{X,L}(f)| &\leq c_L \|f - \Pi_X f\|_{L_q(\Omega)} \\ &\leq c_L c_S h_X^{\tau-d(1/2-1/q)_+} \|f - \Pi_X f\|_{W_2^\tau(\Omega)} \\ &\leq c_L c_S h_X^{\tau-d(1/2-1/q)_+} \|f - \Pi_X f\|_{W_2^\tau(\Omega)} \\ &\leq c_L c_S c_E h_X^{\tau-d(1/2-1/q)_+} \|f - \Pi_X f\|_{\mathcal{H}} \\ &\leq c_L c_S c_E h_X^{\tau-d(1/2-1/q)_+} \|f\|_{\mathcal{H}}, \end{aligned}$$

and it follows from (2.5) that

$$\|v_L - \Pi_X v_L\|_{\mathcal{H}} = \sup_{\|f\|_{\mathcal{H}} \leq 1} |L(f) - Q_{X,L}(f)| \leq c_L c_S c_E h_X^{\tau-d(1/2-1/q)_+}.$$

To obtain (3.5) we can just bound h_X from above with $c_Q n^{-1/d}$ and obtain that for all $h_X \leq h_0$, i.e., for all $n \geq (c_Q/h_0)^d$, holds

$$\|v_L - \Pi_{X_n} v_L\|_{\mathcal{H}} \leq c_L c_S c_E c_Q^{\tau-d(1/2-1/q)_+} n^{-\tau/d + (1/2-1/q)_+}.$$

Observe that the proof is just a consequence of the L_q -continuity of L and of the fact that the quadrature rule is an exact application of the functional to the interpolant (see (2.4)). It is clear that similar results can be obtained for any other kernel for which error estimates in the L_q -norm are available for interpolation (and we give an example at the end of this section). Moreover, the rate of convergence just comes from the fact that quasi-uniform points give a good error for interpolation with translational invariant kernels, and in particular the bound makes no distinction between different functionals.

It remains open to investigate if approximation by L -adapted points can achieve a better approximation rate. We expect this to be the case, and also that these rates can be achieved by adaptive quadratures via greedy algorithms. An example supporting this claim is discussed in Section 7.2.

Remark 3.1 (Superconvergence) As a consequence of the theorem, for some class of functions we can derive superconvergence with respect to L_r -norms in the sense of [36, 38], i.e., a rate of convergence of kernel interpolation which is better than the one for generic functions in \mathcal{H} .

Namely, taking $1 \leq r \leq \infty$, for the interpolation of a generic function $u \in \mathcal{H}$ the application of the inequality of Theorem 3.1 gives the standard error estimate

$$\|u - \Pi_X u\|_{L_r(\Omega)} \leq Ch_X^{\tau-d(\frac{1}{2}-\frac{1}{r})_+} \|u - \Pi_X u\|_{\mathcal{H}} \leq Ch_X^{\tau-d(\frac{1}{2}-\frac{1}{r})_+} \|u\|_{\mathcal{H}}. \quad (3.6)$$

On the other hand, for any function v such that the functional $L := \langle v, \cdot \rangle_{\mathcal{H}}$ satisfies the assumptions of Theorem 3.2 for some q , the theorem gives

$$\|v - \Pi_X v\|_{\mathcal{H}} \leq Ch_X^{\tau-d(\frac{1}{2}-\frac{1}{q})_+},$$

and thus (3.6) can be improved to

$$\|v - \Pi_X v\|_{L_r(\Omega)} \leq Ch_X^{2\tau-d\left(\left(\frac{1}{2}-\frac{1}{r}\right)_+ + \left(\frac{1}{2}-\frac{1}{q}\right)_+\right)}.$$

In the case of L_2 -approximation, which is mainly addressed in [36, 38], this means that this class of functions can be approximated with an order of $2\tau - d(1/2 - 1/q)_+$ instead of $\tau - d(1/2 - 1/q)_+$. In particular, for the class of functions of Example 2.4 this result coincides with the one of [36], even if more general function classes are included in our analysis.

Remark 3.2 (Further characterization of the admissible functionals) Observe that in the case of the kernels and native spaces discussed in this section, the condition $L \in \mathcal{H}' \cap (L_q(\Omega))'$ may be equivalent to certain smoothness assumptions on the Riesz representer v_L . Nevertheless, we do not have a satisfactory and complete characterization of this relation, so we leave this point open for future investigations.

Finally, we mention that there are other translational invariant kernels that are comprised in this analysis, and for which there are error statements for interpolation which are completely analogous to the ones of Theorem 3.1 but prove faster decay of

the error. For example, for the Gaussian and Inverse Multiquadric kernels the error estimates of [32] allow to prove exponential rates of convergence for the approximation with sequences of quasi-uniform points. Since the proof is completely analogous to that of Theorem 3.2, we omit it here. Nevertheless, we remark that a similar superconvergence as in Remark 3.1 happens also in this case, since the \mathcal{H} -norm of the error can be bounded by a term decaying (exponentially) with h_X .

4 The greedy algorithm

We can now define and discuss the greedy algorithm. We assume only that K is an s.p.d. kernel on a set Ω and that $L \in \mathcal{H}'$ is a linear functional. We recall that the notation $Q_{X,L}$ denotes the fact that optimal weights are used, and in particular the algorithm is completely determined just by the selection of the set of points.

Definition 4.1 (Greedy algorithm) Let $X_0 := \emptyset$ and $Q_{X_0,L}(f) := 0$ for all $f \in \mathcal{H}$. For any $n \in \mathbb{N}$, the greedy algorithm selects a point

$$x_n \in \arg \min_{x \in \Omega \setminus X_{n-1}} e_{\mathcal{H}}(Q_{X_{n-1} \cup \{x\}, L})$$

with $X_n := X_{n-1} \cup \{x_n\}$.

Observe that thanks to Proposition 2.3 the algorithm is equivalent to the iterative selection of points to interpolate the Riesz representer v_L with a greedy selection rule given by

$$x_n \in \arg \min_{x \in \Omega \setminus X_{n-1}} e_{\mathcal{H}}(Q_{X_{n-1} \cup \{x\}, L}) = \arg \min_{x \in \Omega \setminus X_{n-1}} \|v_L - \Pi_{X_{n-1} \cup \{x\}}(v_L)\|_{\mathcal{H}}, \quad (4.1)$$

i.e., the new point provides the locally \mathcal{H} -optimal update of the interpolant of v_L .

The locally optimal selection rule is known to be the f/P -greedy selection, as shown in [24, 44]. Although well known, we prove this fact and also give the complete definition of the algorithm to stress the fact that it can be efficiently implemented in an iterative way. To this end, we first recall that the interpolation error can be bounded similarly as the worst-case quadrature error. Indeed, in the case of interpolation it is common to define the power function

$$P_{X_n}(x) := \sup_{\|f\|_{\mathcal{H}} \leq 1} |f(x) - (\Pi_{X_n} f)(x)|, \quad (4.2)$$

and, in the language of this paper, it is clear that this is the worst-case error for the weight-optimal quadrature of the functional $L := \delta_x$. In fact it holds also that

$$P_{X_n}(x) = \|K(\cdot, x) - \Pi_{X_n} K(\cdot, x)\|_{\mathcal{H}},$$

since $K(\cdot, x)$ is the Riesz representer of the point-evaluation functional. This implies in particular that the power function is continuous in $x \in \Omega$, it vanishes if and only

if $x \in X_n$, and $P_0(x) = \sqrt{K(x, x)}$ for all $x \in \Omega$. Moreover, if $\{v_i\}_{i=1}^n$ is any \mathcal{H} -orthonormal basis of $V(X_n)$, by the definition of orthogonal projection we clearly have

$$\Pi_{X_n} v = \sum_{i=1}^n \langle v, v_i \rangle_{\mathcal{H}} v_i \quad \text{for all } v \in \mathcal{H}, \quad (4.3)$$

and in particular

$$P_{X_n}(x) = \left\| K(\cdot, x) - \sum_{i=1}^n \langle K(\cdot, x), v_i \rangle_{\mathcal{H}} v_i \right\|_{\mathcal{H}} = \sqrt{K(x, x) - \sum_{i=1}^n v_i(x)^2}. \quad (4.4)$$

In the following we write P_n instead of P_{X_n} to simplify the notation.

Using this power function, we can now recall that the f/P -greedy rule (4.1) selects a new point as

$$x_n \in \arg \max_{x \in \Omega \setminus X_{n-1}} \frac{|v(x) - (\Pi_{X_{n-1}} v)(x)|}{P_{n-1}(x)}, \quad (4.5)$$

and it is thus clear that the name of this selection rule is indeed given by the fact that it selects a point that maximizes the ratio between the function interpolation residual (the “ f ” component) and the power function (the “ P ” component).

Remark 4.1 (P-greedy) We recall that the same greedy algorithm for interpolation, but using the P -greedy selection rule (i.e., select at each iteration one of the points which maximize the power function in (4.2)) has been shown to produce sequences of points with $h_{X_n} \leq cn^{-\frac{1}{d}(1-\varepsilon)}$ for all $\varepsilon > 0$ in [35]. This has been refined to hold also for $\varepsilon = 0$ in [43], i.e., actually the P -greedy algorithm selects sequences of points which satisfy $h_{X_n} \leq cn^{-\frac{1}{d}}$, and it follows that quadrature based on centers selected by the P -greedy algorithm give exactly the approximation order of Theorem 3.2.

To proceed and describe any iterative algorithm that works with nested sequences of interpolation (or quadrature) points, it is convenient to use the following Newton basis, which is an instance of an \mathcal{H} -orthonormal basis.

Proposition 4.1 (Newton basis [25, 31]) *Given a sequence of nested sets $\{X_n\}_{n \in \mathbb{N}}$ of pairwise distinct points in Ω with $X_0 := \emptyset$ and $X_{n+1} := X_n \cup \{x_{n+1}\}$, the Newton basis is a sequence $\{v_k\}_{k \in \mathbb{N}} \subset \mathcal{H}$ such that for all $n \in \mathbb{N}$ the set $\{v_k\}_{k=1}^n$ is an \mathcal{H} -orthonormal basis of $V(X_n)$. This means that $V(X_n) = \text{span}\{v_i\}_{i=1}^n$ and that $\langle v_i, v_j \rangle_{\mathcal{H}} = \delta_{ij}$ for all $i \neq j$ and for all $n \in \mathbb{N}$. Moreover, the Newton basis property $v_i(x_j) = 0$ for all $1 \leq j < i \leq n$ is satisfied.*

The basis can be constructed by a Gram-Schmidt procedure over $\{K(\cdot, x_i)\}_{i \in \mathbb{N}}$, which gives

$$v_1 = \frac{K(\cdot, x_1)}{P_0(x_1)}, \quad v_n = \frac{K(\cdot, x_n) - \sum_{k=1}^{n-1} v_k(x_n) v_k}{P_{n-1}(x_n)}, \quad n \in \mathbb{N}. \quad (4.6)$$

Using this Newton basis it becomes easy to describe the efficient update of the interpolant for an increasing set of points (see [24, 31]), and also to prove the local optimality of the f/P -greedy selection (see [44]).

Proposition 4.2 (Efficient update and f/P -greedy [24, 31, 44]) Let $v \in \mathcal{H}$, $X_n := \{x_i\}_{i=1}^n := X_{n-1} \cup \{x_n\} \subset \Omega$, and let $\{v_k\}_{k=1}^n$ be the Newton basis of $V(X_n)$.

Then the interpolant and the power function can be updated as

$$\Pi_{X_n} v = \Pi_{X_{n-1}} v + \langle v, v_n \rangle_{\mathcal{H}} v_n = \Pi_{X_{n-1}} v + \langle v - \Pi_{X_{n-1}} v, v_n \rangle_{\mathcal{H}} v_n, \quad (4.7)$$

$$P_n(x)^2 = P_{n-1}(x)^2 - v_n(x)^2. \quad (4.8)$$

Moreover, the locally \mathcal{H} -optimal selection rule is given by the f/P -selection rule, i.e.,

$$\arg \min_{x \in \Omega \setminus X_{n-1}} \|v - \Pi_{X_{n-1} \cup \{x\}} v\|_{\mathcal{H}} = \arg \max_{x \in \Omega \setminus X_{n-1}} \frac{|v(x) - (\Pi_{X_{n-1}} v)(x)|}{P_{n-1}(x)}. \quad (4.9)$$

This proposition and Proposition 2.3 clearly show that the greedy algorithm of Definition 4.1 coincides with the f/P -greedy algorithm for the interpolation of v_L . In particular, the update rule (4.7) can be used to incrementally construct the interpolant by computing only a new Newton basis element and the corresponding coefficient each time a point is added. Moreover, this update of the interpolant and the formula (4.8) for the power function allow the efficient update of the selection rule (4.9).

Remark 4.2 (Change of basis) Observe that, once the selection of the points is completed, it is convenient to express the interpolant back from the Newton basis to the standard basis $\{K(\cdot, x_i)\}_{i=1}^n$ to compute the quadrature of a function f . Indeed, for $\Pi_{X_n}(v_L) = \sum_{i=1}^n w_i^* K(\cdot, x_i)$ it holds

$$Q_{X_n, L}(f) = \langle f, \Pi_{X_n} v_L \rangle_{\mathcal{H}} = \sum_{i=1}^n w_i^* \langle K(\cdot, x_i), v \rangle_{\mathcal{H}} = \sum_{i=1}^n w_i^* f(x_i),$$

i.e., the computation of $Q_{X_n, L}(f)$ requires only the knowledge of the weights and the evaluations of f on X_n . Using instead $\Pi_{X_n} v_L = \sum_{i=1}^n \langle v_L, v_i \rangle_{\mathcal{H}} v_i$ we have

$$Q_{X_n, L}(f) = \sum_{i=1}^n \langle v_L, v_i \rangle_{\mathcal{H}} \langle v_i, f \rangle_{\mathcal{H}},$$

where the terms $\langle v_i, f \rangle_{\mathcal{H}}$ are not directly accessible.

Computing this change of basis is trivial, since the matrix of change of basis is lower triangular (see [31]). We remark moreover that an efficient implementation of the whole greedy approximation process to compute X_n and $\{w_i^*\}_{i=1}^n$ (or $\{\alpha_i\}_{i=1}^n$, as in Proposition 2.2) via f/P -greedy is available in Matlab [33] and in Python [34].

Finally, for later use we remark also that it can be proven that for all $X_n := X_{n-1} \cup \{x_n\}$ there is a splitting of the interpolation error in the form

$$\|v - \Pi_{X_n} v\|_{\mathcal{H}}^2 = \|v - \Pi_{X_{n-1}} v\|_{\mathcal{H}}^2 - \left(\frac{v(x) - (\Pi_{X_{n-1}} v)(x)}{P_{n-1}(x)} \right)^2, \quad (4.10)$$

which represents a kind of energy splitting between the interpolant and the residual. This last formula was in fact the starting point for the introduction of the f/P -greedy algorithm in [24], but we followed a different route in the presentation to highlight instead the local optimality of the selection rule.

5 Convergence for greedy points and continuous kernels

For the greedy algorithm of Definition 4.1, or equivalently the f/P -greedy algorithm applied to v_L , we can now derive convergence results. The proof follows the lines of the one of Theorem 3.7 in [7], and in particular it makes use of the following proposition.

Proposition 5.1 (Lemma 3.4 in [7]) *If $\{a_n\}_{n \in \mathbb{N}}$ is a sequence of non-negative numbers such that for a given $A > 0$ it holds $a_1 \leq A$ and $a_{n+1} \leq a_n(1 - a_n/A)$, then $a_n \leq An^{-1}$ for all $n \in \mathbb{N}$.*

The theorem proves that the worst case error converges to zero at least as $n^{-1/2}$. As we will explain when running numerical experiments in Section 7, this has to be understood as a preliminary result, since faster convergence is often observed in practice. Nevertheless, this speed of convergence is independent of the input space dimension and, as we will comment later, it is better than the one that we proved for quasi-uniform points in certain cases.

We remark that convergence of kernel-based integral approximations based instead on P -greedy points has been recently analyzed in [16].

Theorem 5.1 (Convergence rates for greedy quadrature) *For any continuous s.p.d. kernel on a bounded and Lebesgue measurable set $\Omega \subset \mathbb{R}^d$, let $L \in \mathcal{H}'$ be a linear functional such that there exists $1 \leq q \leq \infty$ and $c_L \geq 0$ such that*

$$|L(f)| \leq c_L \|f\|_{L^q(\Omega)} \quad \text{for all } f \in \mathcal{H},$$

and let $v_L \in \mathcal{H}$ be its Riesz representer.

Then, for the sequence of points $\{X_n\}_{n \in \mathbb{N}}$ selected by the greedy algorithm it holds

$$e_{\mathcal{H}}(Q_{X_n, L}) = \|v_L - \Pi_{X_n} v_L\|_{\mathcal{H}} \leq c_G n^{-1/2} \quad \text{for all } n \in \mathbb{N},$$

where $c_G := \max \left\{ \|v_L\|_{\mathcal{H}}, c_L |\Omega|^{1/q} \max_{x \in \Omega} \sqrt{K(x, x)} \right\}$.

Proof For notational simplicity we denote the residual by $r_n := v_L - \Pi_{X_n} v_L$. We assume that $r_n \neq 0$ for any finite n , otherwise we are done. Then equation (4.10) for v_L reads

$$\|r_n\|_{\mathcal{H}}^2 = \|r_{n-1}\|_{\mathcal{H}}^2 - \frac{|r_{n-1}(x_n)|^2}{P_{n-1}(x_n)^2} = \|r_{n-1}\|_{\mathcal{H}}^2 \left(1 - \frac{|r_{n-1}(x_n)|^2}{P_{n-1}(x_n)^2 \|r_{n-1}\|_{\mathcal{H}}^2} \right), \quad (5.1)$$

and by the definition (4.1) of the f/P -greedy selection we have that

$$\frac{|r_{n-1}(x_n)|}{P_{n-1}(x_n)} = \max_{x \in \Omega \setminus X_{n-1}} \frac{|r_{n-1}(x)|}{P_{n-1}(x)}. \quad (5.2)$$

Now let $\bar{x} \in \Omega$ be such that

$$\bar{x} := \arg \max_{x \in \Omega} |r_{n-1}(x)| = \arg \max_{x \in \Omega \setminus X_{n-1}} |r_{n-1}(x)|, \quad (5.3)$$

where the two maxima are equal since $r_{n-1} = 0$ on X_{n-1} and thus $\bar{x} \notin X_{n-1}$. Since $\bar{x} \notin X_{n-1}$ it also holds $P_{n-1}(\bar{x}) \neq 0$, and then using (5.2) and (5.3) we have

$$\frac{|r_{n-1}(x_n)|}{P_{n-1}(x_n)} = \max_{x \in \Omega \setminus X_{n-1}} \frac{|r_{n-1}(x)|}{P_{n-1}(x)} \geq \frac{|r_{n-1}(\bar{x})|}{P_{n-1}(\bar{x})} = \frac{\|r_{n-1}\|_{L^\infty(\Omega)}}{P_{n-1}(\bar{x})} \geq \frac{\|r_{n-1}\|_{L^\infty(\Omega)}}{\|P_{n-1}\|_{L^\infty(\Omega)}}. \quad (5.4)$$

Moreover, since $Id - \Pi_V = \Pi_{V^\perp}$ for any closed subspace $V \subset \mathcal{H}$, and using again the fact that orthogonal projections are self adjoint, we have

$$\begin{aligned} \|r_{n-1}\|_{\mathcal{H}}^2 &= \langle r_{n-1}, r_{n-1} \rangle_{\mathcal{H}} = \langle v_L - \Pi_{X_{n-1}} v_L, r_{n-1} \rangle_{\mathcal{H}} = \langle \Pi_{V(X_{n-1})^\perp} v_L, r_{n-1} \rangle_{\mathcal{H}} \\ &= \langle v_L, \Pi_{V(X_{n-1})^\perp} r_{n-1} \rangle_{\mathcal{H}} = \left(L \circ \Pi_{V(X_{n-1})^\perp} \right) (r_{n-1}). \end{aligned}$$

Using Proposition 2.4 we have that also the functional $L \circ \Pi_{V(X_{n-1})^\perp}$ is continuous with respect to the L_q -norm on $L_q(\Omega) \cap \mathcal{H}$, with norm bounded by c_L . Using this fact, and since the boundedness of Ω implies $\|r_{n-1}\|_{L_q(\Omega)} \leq \|r_{n-1}\|_{L^\infty(\Omega)}$, we can thus control the norm of the residual as

$$\|r_{n-1}\|_{\mathcal{H}}^2 \leq c_L \Pi_{V(X_{n-1})^\perp} \|r_{n-1}\|_{L_q(\Omega)} \leq c_L \|r_{n-1}\|_{L_q(\Omega)} \leq c_L |\Omega|^{1/q} \|r_{n-1}\|_{L^\infty(\Omega)},$$

i.e., $\|r_{n-1}\|_{L^\infty(\Omega)} \geq \|r_{n-1}\|_{\mathcal{H}}^2 / (c_L |\Omega|^{1/q})$. Combining this bound and inequality (5.4) we can continue to obtain

$$\begin{aligned} \frac{|r_{n-1}(x_n)|^2}{P_{n-1}(x_n)^2 \|r_{n-1}\|_{\mathcal{H}}^2} &\geq \frac{\|r_{n-1}\|_{L^\infty(\Omega)}^2}{\|P_{n-1}\|_{L^\infty(\Omega)}^2 \|r_{n-1}\|_{\mathcal{H}}^2} \geq \frac{\|r_{n-1}\|_{\mathcal{H}}^4}{\|P_{n-1}\|_{L^\infty(\Omega)}^2 \|r_{n-1}\|_{\mathcal{H}}^2 c_L^2 |\Omega|^{2/q}} \\ &= \frac{\|r_{n-1}\|_{\mathcal{H}}^2}{\|P_{n-1}\|_{L^\infty(\Omega)}^2 c_L^2 |\Omega|^{2/q}}. \end{aligned}$$

Now we can set $A_{n-1} := \|P_{n-1}\|_{L^\infty(\Omega)}^2 c_L^2 |\Omega|^{2/q}$ and, since P_n is non increasing in n , we have

$$A_{n-1} \leq A_0 = c_L^2 |\Omega|^{2/q} \|P_0\|_{L^\infty(\Omega)}^2 = c_L^2 |\Omega|^{2/q} \max_{x \in \Omega} K(x, x),$$

and thus the last inequality reads

$$\frac{|r_{n-1}(x_n)|^2}{P_{n-1}(x_n)^2 \|r_{n-1}\|_{\mathcal{H}}^2} \geq \frac{\|r_{n-1}\|_{\mathcal{H}}^2}{A_0}.$$

Inserting this result in (5.1) and defining $A := \max \left\{ \|v_L\|_{\mathcal{H}}^2, A_0 \right\}$, we finally have

$$\|r_n\|_{\mathcal{H}}^2 \leq \|r_{n-1}\|_{\mathcal{H}}^2 \left(1 - \frac{\|r_{n-1}\|_{\mathcal{H}}^2}{A_0} \right) \leq \|r_{n-1}\|_{\mathcal{H}}^2 \left(1 - \frac{\|r_{n-1}\|_{\mathcal{H}}^2}{A} \right)$$

and $\|r_0\|_{\mathcal{H}}^2 = \|v_L\|_{\mathcal{H}}^2 \leq A$, and thus the result follows by applying Lemma 5.1 with $a_n := \|r_{n-1}\|_{\mathcal{H}}^2$.

Before discussing some consequences of this result, we show in the following corollary that the same proof idea applies also to other greedy interpolation strategies which are commonly used in kernel interpolation and quadrature.

Corollary 5.1 (Convergence for other selection rules) *Under the same hypotheses of Theorem 5.1, the same convergence result holds if the f/P -greedy selection rule (4.5) is replaced by one of the following ones:*

a) *The f -greedy selection rule, i.e.,*

$$x_n \in \arg \max_{x \in \Omega \setminus X_{n-1}} |v(x) - (\Pi_{X_{n-1}} v)(x)|.$$

b) *The selection rule of Algorithm 4 on p. 93 in [29], i.e.,*

$$x_n \in \arg \max_{x \in \Omega \setminus X_{n-1}} \frac{|v(x) - (\Pi_{X_{n-1}} v)(x)|}{\sqrt{K(x, x)}}.$$

Proof The proof of Theorem 5.1 depends on the f/P -greedy selection rule only via equation (5.4) (and on equation (5.2), which is nevertheless used only to derive the latter). We thus only need to show how to obtain the same bound as in (5.4) with these other selection rules:

a) In this case, by the definition of x_n , it immediately holds that

$$\frac{|r_{n-1}(x_n)|}{P_{n-1}(x_n)} = \frac{\|r_{n-1}\|_{L^\infty(\Omega)}}{P_{n-1}(x_n)} \geq \frac{\|r_{n-1}\|_{L^\infty(\Omega)}}{\|P_{n-1}\|_{L^\infty(\Omega)}}.$$

b) For this selection rule, using the same argument as in the proof of Theorem 5.1 one obtains

$$\frac{|r_{n-1}(x_n)|}{P_{n-1}(x_n)} \geq \frac{\|r_{n-1}\|_{L^\infty(\Omega)}}{\left\| \sqrt{K(x, x)} \right\|_{L^\infty(\Omega)}}$$

in place of (5.4), and then the following of the proof is simplified since $\sqrt{K(x, x)} = P_0(x)$ for all $x \in \Omega$.

Remark 5.1 (Improvement over existing f -greedy convergence results) Observe that existing convergence results for f -greedy are quite weaker. Indeed, the two results of rates of convergence are both from [24], where it is shown (see Corollary 3.3.8) that for $K \in C^2$ and for all $v \in \mathcal{H}$ it holds

$$\min_{1 \leq i \leq n} \|v - \Pi_{X_n} v\|_{L^\infty(\Omega)} \leq Cn^{-1/d},$$

which can be improved to $n^{-2/d}$ if $K \in C^4$ and if it is verified a posteriori that all the points are selected in the interior of Ω . Alternatively, Proposition 3.2.2 in [24] states that for any continuous kernel and for each $v \in \mathcal{H}$ there exists a subsequence $\{m_n\}_{n \in \mathbb{N}} \subset \mathbb{N}$ such that

$$\|v - \Pi_{X_{m_n}} v\|_{L^\infty(\Omega)} \leq Cm_n^{-1/2}.$$

Although Corollary 5.1 applies only to some functions v and not to the entire \mathcal{H} , it gives nevertheless a bound on the error in the \mathcal{H} -norm, which is stronger than the L_∞ -norm, and it works for the entire sequence of interpolants.

Remark 5.2 (Other selection criteria) We remark that the proof does not directly apply instead to the power scaled residual greedy selection (psr-greedy) which has been recently introduced in [8]. In this case the new point is selected as

$$x_n \in \arg \max_{x \in \Omega \setminus X_{n-1}} P_{n-1}(x) |v(x) - (\Pi_{X_{n-1}} v)(x)|.$$

The present approach fails in this case. Indeed, following the same idea, one needs to bound from below the ratio $P_{n-1}(\bar{x})/P_{n-1}(x_{n-1})$, where \bar{x} is defined as in (5.3), and this quantity may be arbitrarily small.

Theorem 5.1 is a first result on the convergence of the algorithm, and it has the advantage of providing rates of convergence that do not depend on the input dimension. Nevertheless, it is far from optimal in the sense that the ideal result that one can aim for in the setting of adaptive algorithms is rather the following: Given a functional L such that there exist a sequence of possibly unknown (or not computable) point sets $\{\bar{X}_n\}_{n \in \mathbb{N}}$, and $t, C > 0$ such that

$$\|v_L - \Pi_{\bar{X}_n} v_L\|_{\mathcal{H}} \leq C n^{-t} \quad \text{for all } n \in \mathbb{N},$$

find a constructive algorithm that selects points $\{X_n\}_{n \in \mathbb{N}}$ such that

$$\|v_L - \Pi_{X_n} v_L\|_{\mathcal{H}} \leq C' n^{-t} \quad \text{for all } n \in \mathbb{N},$$

with a possibly larger constant $C' \geq C$. There is no reason to expect that this can be achieved by greedy algorithms in particular, but this is actually the case for other types of greedy algorithms (see e.g. [4]), and they are furthermore particularly attractive for their easiness of implementation.

In the direction of this optimal expectation, it should also be mentioned that our algorithm is closely related to the greedy algorithms of [41], where a dictionary in a generic Hilbert space H is used to approximate a single function $f \in H$. In this case, the paper [7] proves that for every Hilbert space there exist a dictionary and a function in the class of Example 2.3 such that the rate of $n^{-1/2}$ can not be improved. Nevertheless, in the case of the current paper we are using not a generic dictionary, but rather the particular one that is generated by translates of the reproducing kernel, and thus there is no reason to believe that the result of Theorem 5.1 can not be improved.

Despite being not optimal, we remark that also the rates of convergence proved here are better than the ones of Theorem 3.2 in certain cases. Observe that this of course does not mean, even in this case, that greedy points are better than quasi-uniform points, but only that the estimate is better.

Proposition 5.2 (Comparison of the rates for greedy and quasi-uniform points)
Under the assumptions of Theorem 3.2, let $r := -\tau/d + (1/2 - 1/q)_+$ be the rate of

convergence for quadrature with quasi-uniform points. Then the rates of convergence of Theorem 5.1 are better than the ones of Theorem 3.2 if $q > 2$ and

$$\frac{d}{2} < \tau < d \left(1 - \frac{1}{q}\right),$$

and in this case it holds $-\frac{1}{2} < r < -\frac{1}{q}$, i.e., the improvement is at most of order $n^{-1/2}$.

If L is an integration functional as in Example 2.1, the condition on τ is equivalent to requiring that $d/2 < \tau < d/p$ if $g_L \in L_p(\Omega)$.

Proof Since $\tau > d/2$, the proved rate of convergence of the greedy algorithm is better if

$$-\frac{1}{2} < -\frac{\tau}{d} + \left(\frac{1}{2} - \frac{1}{q}\right)_+, \quad (5.5)$$

and this can happen only if $\left(\frac{1}{2} - \frac{1}{q}\right)_+ > 0$, i.e., $q > 2$. In this case we have $(1/2 - 1/q)_+ = 1/2 - 1/q \in (0, 1/2]$, and then (5.5) guarantees that the greedy estimate is better for all d and τ such that

$$\frac{d}{2} < \tau < \frac{d}{2} + \frac{d}{2} - \frac{d}{q} = d \left(1 - \frac{1}{q}\right).$$

When substituted into r , these two bounds give the values $-1/2$ and $-1/q$.

In the particular case of integration functionals, this is equivalent to require that $d/2 < \tau < d/p$ if $g_L \in L_p(\Omega)$.

We conclude this section with some remarks on the results.

Remark 5.3 (Almost optimal rate when $\tau \approx d/2$) If $\Omega \subset \mathbb{R}^d$ is sufficiently regular, it is known that the rate $n^{-\tau/d}$ is optimal for approximation of $L(f) := \int_{\Omega} f(x)dx$ for functions in the Sobolev space $W_2^{\tau}(\Omega)$ with $\tau > d/2$. That is, there is no sequence of quadrature rules whose worst case error decays faster than this; see [27, Section 1.3.11] and [28, Section 4.2.4]. Theorem 5.1 therefore shows that the greedy quadrature algorithm is almost optimal if τ is close to $d/2$.

Remark 5.4 (Integration on manifolds) We remark that both Theorem 3.2 and Theorem 5.1 remain valid also for certain sets Ω which are not flat subsets of \mathbb{R}^d .

Namely, Theorem 5.1 only assumes that L_p spaces can be defined over Ω , and this is a fairly general assumption. Otherwise, the space \mathcal{H} is treated as a generic Hilbert space, without particular links to the structure of the underlying subset of \mathbb{R}^d .

Theorem 3.2, on the other hand, makes use of the error estimate of Theorem 3.1, which holds for $\Omega \subset \mathbb{R}^d$. Nevertheless, similar results exist for more general sets such as for manifolds embedded in \mathbb{R}^s (see e.g. [11]), and in this case the error rates depend on the dimension d of the manifold. These kind of results can be used in the proof of the theorem, and we analyze an example of this setting in Section 7.3.

6 Perturbations of the Riesz representer

The construction of both a generic weight-optimal quadrature formula, and the greedy selection of the centers, are based on the interpolation of the Riesz representer v_L . The whole algorithm is thus based on the availability of evaluations of v_L , and this can be an unrealistic assumption in some cases. In particular, since $v_L = L^y(K(\cdot, y))$, we need to have an efficient and exact way to compute the functional on the kernel, which is often not possible for example in the case of integration on arbitrary sets Ω .

Instead, we can assume to have an accurate but expensive approximation \tilde{L} of L , which can be evaluated but is too expensive to be practical when the evaluation of the integrand f is expensive. Since the evaluation of the kernel is instead cheap, we can compute the exact Riesz representer of \tilde{L} , i.e., $\tilde{v}_L := \tilde{L}^y K(\cdot, y)$, and use it as an approximation of the exact Riesz representer v_L .

For example, \tilde{L} can be in the form of a high-accuracy quadrature rule $\tilde{L}(f) := \sum_{i=1}^M \rho_i f(z_i)$ where M is very large. In this case computing $\tilde{L}(f)$ for a function f which is expensive to evaluate is not a viable option, while we can construct $\tilde{v}_L := \sum_{i=1}^M \rho_i K(\cdot, z_i)$.

As a first step, the following proposition states that the worst case error of the weight-optimal quadrature is stable with respect to the use of an approximated Riesz representer, provided that also the perturbed functional is continuous. Observe that the result is a consequence of the fact that kernel-based interpolation is a continuous process in Hilbert space, being a projector.

Similar results appear in [5, Appendix B] and [39, Section 2] for numerical integration, although they are less sharp. For example, the former has $\sqrt{n}\varepsilon_L^2$ in the place of ε_L^2 in the upper bound.

Proposition 6.1 (Stability) *Let $L, \tilde{L} \in \mathcal{H}'$ be two functionals with Riesz representers $v_L, \tilde{v}_L \in \mathcal{H}$, where \tilde{L} is a perturbation of L with*

$$\varepsilon_L := \sup_{\|f\|_{\mathcal{H}} \leq 1} |L(f) - \tilde{L}(f)|.$$

Let $X \subset \Omega$ and let $Q_{X,L}$ and $\tilde{Q}_{X,L} := Q_{X,\tilde{L}}$ be the weight-optimal quadrature rules obtained by the interpolation on X of v_L and \tilde{v}_L , respectively.

Then the error obtained by using the wrong Riesz representer \tilde{v}_L to approximate L can be bounded as

$$\sup_{\|f\|_{\mathcal{H}} \leq 1} |L(f) - \tilde{Q}_{X,L}(f)|^2 = e_{\mathcal{H}}(Q_{X,L})^2 + \varepsilon_L^2.$$

Proof By linearity we have that $\delta_L := v_L - \tilde{v}_L \in \mathcal{H}$ is the Riesz representer of the functional $L - \tilde{L} \in \mathcal{H}'$, and by assumption $\|\delta_L\|_{\mathcal{H}} = \varepsilon_L$. We can then rewrite the worst-case error of the statement as required, i.e.,

$$\begin{aligned} \sup_{\|f\|_{\mathcal{H}} \leq 1} |L(f) - \tilde{Q}_{X,L}(f)|^2 &= \|v_L - \Pi_X(\tilde{v}_L)\|_{\mathcal{H}}^2 = \|v_L - \Pi_X v_L + \Pi_X \delta_L\|_{\mathcal{H}}^2 \\ &= \|v_L - \Pi_X v_L\|_{\mathcal{H}}^2 + \|\Pi_X \delta_L\|_{\mathcal{H}}^2 + 2 \langle v_L - \Pi_X v_L, \Pi_X \delta_L \rangle_{\mathcal{H}} \\ &= e_{\mathcal{H}}(Q_{X,L})^2 + \varepsilon_L^2 + 2 \langle v_L - \Pi_X v_L, \Pi_X \delta_L \rangle_{\mathcal{H}} \\ &= e_{\mathcal{H}}(Q_{X,L})^2 + \varepsilon_L^2, \end{aligned}$$

where in the last step we used the fact that $\Pi_X \delta_L \in V(X)$ and $v_L - \Pi_X v_L \in V(X)^\perp$.

We remark that if \tilde{L} is a quadrature rule, then ε_L is merely the worst-case error of this quadrature rule, and can thus be estimated. Moreover, for e.g. ρ_i and \tilde{z}_i the weights and points of an inexpensive (quasi) Monte Carlo rule, ε_L decays with a well-known rate when M increases.

This result is useful to quantify the error introduced by an approximated knowledge of v_L only if X is fixed. If instead the greedy algorithm is used, then also the points X themselves depend crucially on v_L , and thus perturbations of the Riesz representer lead to different sequences of quadrature centers. To deal with this case we have the following result. In this case we need to assume that \tilde{L} , and not L , is continuous on $L_q(\Omega) \cap \mathcal{H}$, since it is the one used to run the greedy algorithm.

Proposition 6.2 (Stability of the greedy quadrature) *Let $L, \tilde{L} \in \mathcal{H}'$ be two functionals with Riesz representers $v_L, \tilde{v}_L \in \mathcal{H}$, where \tilde{L} is a perturbation of L with*

$$\varepsilon_L := \sup_{\|f\|_{\mathcal{H}} \leq 1} |L(f) - \tilde{L}(f)|.$$

Assume furthermore that there exists $1 \leq q \leq \infty$ and $\tilde{c}_L > 0$ such that \tilde{L} is continuous on $L_q(\Omega) \cap \mathcal{H}$ with norm bounded by \tilde{c}_L .

Let $X_n := X_n(\tilde{L}) \subset \Omega$ be the set of points selected after n iterations of the greedy algorithm applied to \tilde{v}_L , and let $\tilde{Q}_{X_n, L} := \tilde{Q}_{X_n, \tilde{L}}$ be the corresponding weight-optimal quadrature rule. Then for n sufficiently large it holds

$$\sup_{\|f\|_{\mathcal{H}} \leq 1} |L(f) - \tilde{Q}_{X_n, L}(f)|^2 \leq 3\varepsilon_L c_G n^{-1/2} + \varepsilon_L^2,$$

where $c_G := \max \left\{ \|\tilde{v}_L\|_{\mathcal{H}}, \tilde{c}_L |\Omega|^{1/q} \max_{x \in \Omega} \sqrt{K(x, x)} \right\}$.

Proof For simplicity of notation we set $X := X_n$, since n is fixed here. Defining δ_L as in the previous proposition we have

$$\begin{aligned} \sup_{\|f\|_{\mathcal{H}} \leq 1} |L(f) - \tilde{Q}_{X, L}(f)|^2 &= \|v_L - \Pi_X \tilde{v}_L\|_{\mathcal{H}}^2 = \|\tilde{v}_L + \delta_L - \Pi_X \tilde{v}_L\|_{\mathcal{H}}^2 \\ &= \|\tilde{v}_L - \Pi_X \tilde{v}_L\|_{\mathcal{H}}^2 + \|\delta_L\|_{\mathcal{H}}^2 + 2\langle \tilde{v}_L - \Pi_X \tilde{v}_L, \delta_L \rangle_{\mathcal{H}} \\ &\leq \|\tilde{v}_L - \Pi_X \tilde{v}_L\|_{\mathcal{H}}^2 + \|\delta_L\|_{\mathcal{H}}^2 + 2\|\tilde{v}_L - \Pi_X \tilde{v}_L\|_{\mathcal{H}} \|\delta_L\|_{\mathcal{H}} \\ &= \|\tilde{v}_L - \Pi_X \tilde{v}_L\|_{\mathcal{H}} (\|\tilde{v}_L - \Pi_X \tilde{v}_L\|_{\mathcal{H}} + 2\|\delta_L\|_{\mathcal{H}}) + \|\delta_L\|_{\mathcal{H}}^2, \end{aligned}$$

and inserting the estimate of Theorem 5.1 for the functional \tilde{L} gives the result, provided that n is large enough to guarantee that $c_G n^{-1/2} \leq \varepsilon_L$.

7 Numerical experiments

In this section we test the greedy algorithm on various integration test problems. We start by an example where greedy points provide the same rate of convergence of the worst case error as quasi-uniform points, but possibly with an arbitrarily better constant, and then we analyze a case where also the rate is strictly better for greedy points. Then, we show how the algorithm performs on a manifold, and finally we test the method on an Uncertainty Quantification benchmark problem.

The experiments are run by executing the f/P -greedy algorithm using the Riesz representer v_L as a target function, where, as it is customary, the maximization is performed on a discrete training set $X_{tr} \subset \Omega$, which is described in each example. All the experiments use the Python implementation [34] of the f/P -greedy algorithm.

In the first three examples the functionals are as in Example 2.1 for some Ω and some function g_L which are defined in each example. The Riesz representers v_L can not be computed exactly, so we use an approximation $v_{\tilde{L}}$ constructed using a quadrature rule given by $\tilde{N} \in \mathbb{N}$ nodes $\{\tilde{x}_i\}_{i=1}^{\tilde{N}} \subset \Omega$ and weights $\{\tilde{w}_i\}_{i=1}^{\tilde{N}} \subset \mathbb{R}$. This approximation gives

$$\tilde{L}(f) := \sum_{i=1}^{\tilde{N}} \tilde{w}_i f(\tilde{x}_i) g_L(\tilde{x}_i), f \in \mathcal{H}, \quad v_{\tilde{L}}(x) = \sum_{i=1}^{\tilde{N}} \tilde{w}_i g_L(\tilde{x}_i) K(x, \tilde{x}_i).$$

The quadrature rule $Q_{X_n, \tilde{L}}$ is constructed by computing its representer, i.e., the interpolant $\Pi_{X_n} v_{\tilde{L}}(x) = \sum_{i=1}^n w_i^* K(x, x_i)$, where X_n are either quasi-uniform or greedy points. In any case, the worst case error is computed as

$$\begin{aligned} e_{\mathcal{H}}(Q_{X_n, \tilde{L}}) &= \|v_{\tilde{L}} - \Pi_{X_n} v_{\tilde{L}}\|_{\mathcal{H}} = \left(\|v_{\tilde{L}}\|_{\mathcal{H}}^2 - \|\Pi_{X_n} v_{\tilde{L}}\|_{\mathcal{H}}^2 \right)^{1/2} \\ &= \left(\sum_{i,j=1}^{\tilde{N}} \tilde{w}_i \tilde{w}_j g_L(\tilde{x}_i) g_L(\tilde{x}_j) K(\tilde{x}_i, \tilde{x}_j) - \sum_{i,j=1}^n w_i^* w_j^* K(x_i, x_j) \right)^{1/2}. \end{aligned}$$

To simplify the interpretation of the errors, we always scale $v_{\tilde{L}}$ to have unit norm. To obtain the worst case error with respect to the exact L , one should consider the perturbation error ε_L as in Proposition 6.1 and Proposition 6.2. Observe that in this step, to reduce the value of ε_L a quadrature rule with a possibly very large number of centers \tilde{N} may be used (e.g., a Monte Carlo quadrature), since the formula has to be applied only once to the kernel to obtain $v_{\tilde{L}}$.

This step requires anyhow some care. Indeed, as long as a formula with a fixed set of nodes is used, it may happen that $\tilde{L} \in \mathcal{H}'$ even if L is unbounded. This is the case for example of the functional of Section 7.2 if $\alpha \geq 2$ and $\{\tilde{x}_i\}_{i=1}^{\tilde{N}} \in \Omega \setminus \{0\}$. It should thus be taken into account that this approach produces in any case an approximation of \tilde{L} , which may be very far away from L if the wrong approximation is chosen.

7.1 Integration with a compactly supported density

We consider the unit square $\Omega := [0, 1]^2 \subset \mathbb{R}^2$ as input domain and a quadratic Matérn kernel, which is defined as $K(x, y) := \phi(\|x - y\|)$ with $\phi(r) := e^{-\gamma r} (3 + 3\gamma r + (\gamma r)^2)$,

with a free parameter which is set to the value $\gamma = 1$ in this experiment. The kernel corresponds to a value of $\tau = 4$ for the decay of the Fourier transform in (3.1) (see e.g. Appendix D in [10]).

We consider the integration functional as in Example 2.1 where g_L is the indicator function of the square $[0.3, 0.5] \times [0.6, 0.8] \subset \Omega$. One would expect that optimal quadrature points for this functional are uniformly distributed inside the support of g_L . As discussed at the beginning of Section 7, we approximate L with a tensor product Gaussian quadrature with $M := 100^2$ points scaled to the square $[0.3, 0.5] \times [0.6, 0.8]$.

The greedy algorithm is run by selecting points from a uniform grid X_{T_r} of 10^4 points in Ω , and it is terminated when 500 points are selected, or when the maximal absolute interpolation error on X_{T_r} is below the tolerance 10^{-12} . In this way, $n = 358$ points are selected, and they are shown in Figure 7.1, left. The points are colored according to the magnitude of the corresponding weight, and they are overlapped to the contour lines of the density g_L .

The greedy algorithm selects almost all points inside the support of the density, which is the behavior one would expect from an optimal algorithm. Nevertheless, some points are selected in the area where $g_L = 0$, even if the corresponding weights are relatively small. The fact that the centers are selected almost only in the support of g_L seems to be an interesting feature of the greedy method, and it is an open issue to prove whether or not it is the case that optimal points (non necessarily greedy) are localized in this way. If this is the case, the selection of a few centers outside of the support may be caused by the use of the approximate Riesz representer. Moreover, most of the weights are positive, and the negative ones are mostly of small magnitude.

In Figure 7.1, right, we show the worst-case error w.r.t. \tilde{L} obtained with these greedy points. Since the points are nested, it is possible to show the decay of the worst case error as a function of the number of centers. As a comparison, we report also the decay of the worst case error for the optimal quadrature rule which uses uniform grids of points of increasing cardinality. The figure suggests that both approximation errors decay as $n^{-\tau/d}$, and in particular, after an initial phase, the convergence of the greedy error is faster than $n^{-1/2}$, and this suggests that indeed the rate of Theorem 5.1 is in general suboptimal. Moreover, although the rate of convergence is the same for the two distributions of points in the asymptotic regime, the greedy error is smaller and the ratio ρ between the two errors is roughly constant as a function of the number of points. This is due to the fact that the uniform points are filling the full Ω , while the greedy ones fill only the support of g_L . We remark that this ratio ρ can be made arbitrarily small by reducing the support of g_L , but no improvements should be expected in the rate of convergence by using greedy methods.

7.2 Integration with a singular density

We consider here $\Omega := [0, 1]^2 \subset \mathbb{R}^2$ and L as in Example 2.1, and set $g_L(x) = \|x - x_c\|^{-\alpha}$ with $\alpha \geq 0$ and $x_c := [0.5, 0.5]^T$.

To obtain the best rate of convergence of the error in Theorem 3.2, one should choose the smallest possible q that satisfies the hypotheses of the theorem for a given

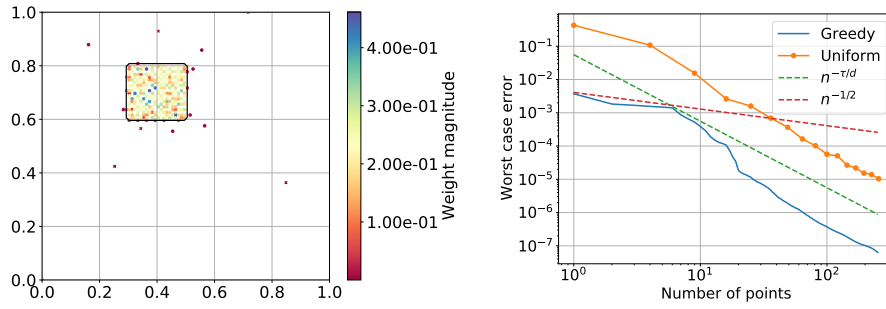


Fig. 7.1 Results of the greedy algorithm in the example of Section 7.1. The left figure shows the contour plot of the density g_L (bold grayscale line), and the position of the points selected by the algorithm, which are colored according to the magnitude of the corresponding positive (circles) or negative (crosses) weight. The right figure shows the decay of the worst-case error as a function of the number of points for the greedy points and for a grid of equally spaced points, and rates of decay scaled to the greedy error.

L . Indeed, if q is sufficiently small the term $(1/2 - 1/q)_+$ vanishes in the exponent. In this case, we can choose the smallest q such that $g_L \in L_p(\Omega)$ with $1/p + 1/q = 1$, and since $g_L \in L_p(\Omega)$ if and only if $p < d/\alpha = 2/\alpha$, it follows that the limiting value is $q = 2/(2 - \alpha)$, which gives $(1/2 - 1/q)_+ = (\alpha/2 - 1/2)_+$. Since the integral is well defined if and only if $p \geq 1$, this implies also that we need to require $1 < d/\alpha$, i.e., $\alpha < d = 2$, to have a bounded functional.

We thus consider values $\alpha \in \{1, 3/2, 2 - 10^{-10}\}$, and run the experiments with the same kernel, and the same setting of the greedy algorithm as in Section 7.1. The integral defining L is approximated by a Monte Carlo approximation \tilde{L} that uses $M := 10^4$ uniformly randomly distributed points on the unit square. The results are reported in Figure 7.2, where the points and weights of the greedy approximation are shown, together with the rate of convergence of the greedy and uniform quadrature as functions of the number of points.

The figure clearly shows that the greedy points are increasingly concentrated around the singularity of g_L as α increases, with weights which are larger. Moreover, the rates of convergence show that quadrature rules with uniform points have a worst case error converging to zero as $n^{-\tau/d + (\alpha/2 - 1/2)_+}$, which is strictly slower than $n^{-\tau/d}$ for $\alpha > 1$. The greedy algorithm, on the other hand, selects non uniform points and it provides a convergence like $n^{-\tau/d}$.

In other words, quasi-uniform points are sub-optimal for too skewed linear functionals.

7.3 Integration on a manifold

A similar experiment as in Section 7.1 is repeated on the sphere $\Omega := \mathbb{S}^2 \subset \mathbb{R}^3$. Assuming Ω is represented in Cartesian coordinates, we consider the integration functional as in Example 2.1 where $g_L(x) := \exp((x - x_c)^T \Sigma (x - x_c))$, $x_c := [0, -1, 0]^T$, and Σ is the diagonal matrix with diagonal $[-5, -5, -3]^T$. The approximation \tilde{L} of

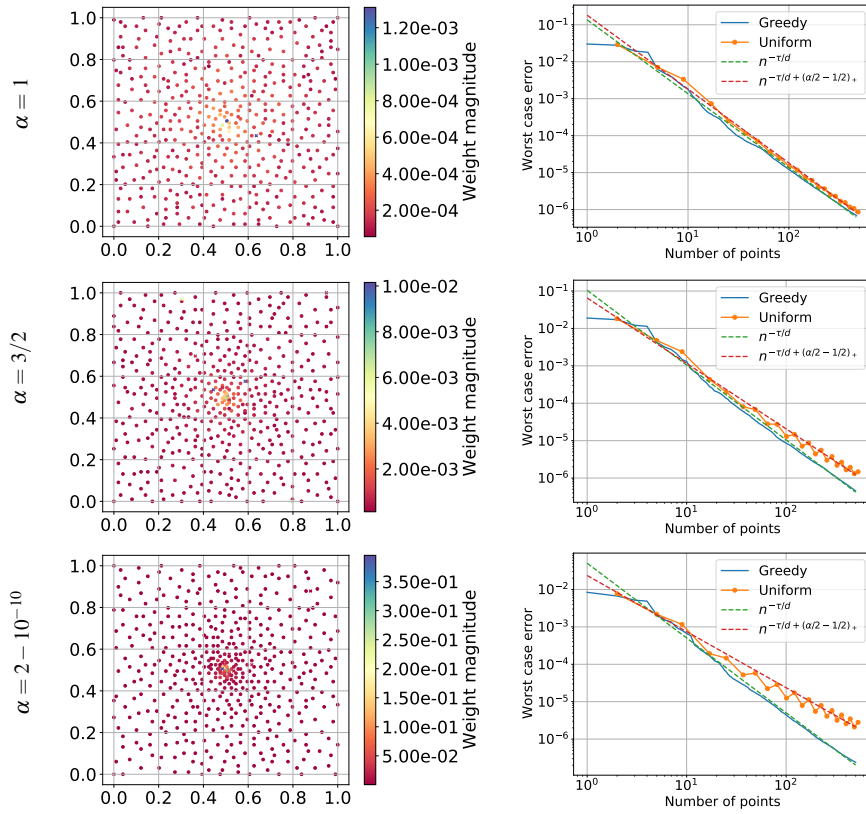


Fig. 7.2 Results of the greedy algorithm in the example of Section 7.2 for values of $\alpha \in \{1, 3/2, 2 - 10^{-10}\}$ (from top to bottom). The left plot shows the position of the points selected by the algorithm, which are colored according to the magnitude of the corresponding positive (circles) or negative (crosses) weight (left column). The right plot shows the decay of the worst-case error as a function of the number of points, both for greedy and equally spaced points. The rates of decay are scaled by a constant.

L is similarly realized with a Monte Carlo approximation with $M := 10^4$ uniformly random points. We use the same setting as in Section 7.1 for the kernel, its parameters, the greedy algorithm and the termination criteria, with the only difference that the points are selected starting from a set of uniformly random points on the sphere.

Moreover, we compare the greedy points with sets of minimal energy points (see e.g. [13]). For each n , these sets $X_n := \{x_i\}_{i=1}^n$ are defined as minima on $(\mathbb{S}^2)^n$ of the Riesz energy

$$E(X_n) := \sum_{i=1}^n \sum_{j=i+1}^n \|x_i - x_j\|^{-2}.$$

We use here the precomputed points from [45], which are found by numerical minimization of this functional.

In this case the density is not constant and does not have compact support, and this is reflected in the fact that the points selected by the greedy algorithm are more spread over the entire Ω (see Figure 7.3, left). Nevertheless, the points are more concentrated in the area where the density g_L is larger, and the weights are accordingly larger. Also in this experiment, a few weights of small magnitude are negative.

Again, the convergence of the worst-case error (see Figure 7.3, right) shows that the greedy algorithm produces quadrature weights with a worst-case error decaying like $n^{-\tau/d}$, if one sets $d = 2$ as the dimension of Ω as an embedded manifold in \mathbb{R}^3 . This is in accordance with known error estimate for kernel interpolation on manifolds (see [11] and Remark 5.4). The same behavior is clearly observed with integration with the minimal energy points.




Fig. 7.3 Results of the greedy algorithm in the example of Section 7.3, represented in spherical coordinates. The left figure shows the contour plot of the density g_L (grayscale lines), and the position of the points selected by the algorithm, which are colored according to the magnitude of the corresponding positive (circles) or negative (crosses) weight. The right figure shows the decay of the worst-case error as a function of the number of points for the greedy points and for sets of minimal energy points, and rates of decay scaled to the greedy error. Observe that d is the dimension of the manifold, and not of the embedding space.

7.4 An Uncertainty Quantification example

As a final example we test the greedy algorithm on the benchmark Uncertainty Quantification (UQ) problem described in [21].

We briefly describe the setting of the problem, and we refer to the cited paper for a thorough discussion. We have a Partial Differential Equation (PDE) modelling a two-phase flow in a porous medium, which depends on three input parameters (the injection rate, the relative permeability degree, and the reservoir porosity) and represents the saturation of some carbon dioxide which is injected into a one-dimensional aquifer over a time interval $[0, T]$.

The aquifer is discretized into 250 equal sized cells, and for a fixed value of the parameter triple $\theta := [\theta_1, \theta_2, \theta_3]^T \in \mathbb{R}^3$, a time-dependent numerical solution of the PDE can be computed by the Finite Volume (FV) method. We denote as $s(\theta) \in$

\mathbb{R}^{250} the nodal values of the numerical solution at the final time T computed with parameters θ , and thus the FV discretization defines a map $\theta \in \mathbb{R}^3 \mapsto s(\theta) \in \mathbb{R}^{250}$.

Instead than just computing the solution for a fixed value, in this case one is interested to quantify the effect on the solution $s(\theta)$ of an uncertain knowledge of the input parameters. In particular, each parameter triple $\theta := [\theta_1, \theta_2, \theta_3]^T$ is assumed to represent a sample from three random variables with known distributions, and thus the solution $s(\theta)$ itself is a random variable with values in \mathbb{R}^{250} . The goal of the benchmark problem is to estimate the mean $\mu_s \in \mathbb{R}^{250}$ and the standard deviation $\sigma_s \in \mathbb{R}^{250}$ of this random vector using as few solutions of the PDE as possible.

The benchmark includes also a dataset [22] which contains a spatial discretization of the input parameters into $N := 10000$ points, i.e., a set $X := \{\theta_i\}_{i=1}^N \subset \mathbb{R}^3$ representing independent samples of the parameters drawn accordingly to the respective distributions, and an implementation of the FV solver to compute the values $Y := \{y_i := s(\theta_i)\}_{i=1}^N \subset \mathbb{R}^{250}$. Moreover, the mean and standard deviation vectors computed by the various methods analyzed in [21] are also available for comparison. As a reference solution, the paper uses the integral computed by a Monte Carlo approximation which uses the full set of nodes X .

In this case, we run the greedy algorithm with the same Matérn kernel and $\gamma := 1/2$. Following [21], the points are selected from X itself, since this discretization incorporates information of the distribution of the parameters and this is assumed to be a known information. The resulting quadrature rule is applied to each of the 250 entries of the solution vector. Namely, if we assume that the points selected by the greedy algorithm are the subset $X_n := \{\theta_{i_1}, \dots, \theta_{i_n}\} \subset X$, we obtain an approximated mean vector $\tilde{\mu}_s \in \mathbb{R}^{250}$ as

$$(\tilde{\mu}_s)_j := \mathcal{Q}_{X_n}(s_j) = \sum_{k=1}^n w_k^* (y_{i_k})_j,$$

where $(y_{i_k})_j$ is the j -th component of the i_k -th output vector. Similarly, the approximated standard deviation is the vector $\tilde{\sigma}_s \in \mathbb{R}^{250}$ with

$$(\tilde{\sigma}_s)_j := \left(\sum_{k=1}^n w_k^* (y_{i_k})_j^2 - (\tilde{\mu}_s)_j^2 \right)^{1/2},$$

where we used the fact that the variance is the difference between the mean of the square and the square of the mean.

Observe that this process is actually extracting a compressed quadrature rule from the reference Monte Carlo one.

For these approximated vectors, we can compute the ℓ_2 errors with respect to the reference mean and standard deviation provided by the Monte Carlo quadrature. Following [21], these are computed as

$$E_\mu := \frac{1}{250} \|\mu_s - \tilde{\mu}_s\|, \quad E_\sigma := \frac{1}{250} \|\sigma_s - \tilde{\sigma}_s\|.$$

The values of E_μ, E_σ are reported in Figure 7.4 for increasing values of the number of quadrature points. We also report the same errors obtained in the cited paper using

some state-of-the-art methods, namely arbitrary polynomial chaos expansion (aPC), spatially adaptive sparse grids (aSG), P -greedy kernel interpolation (PGreedy), and Hybrid stochastic Galerkin (HSG). It is clear that the present method yields comparable results, and it even outperforms them if sufficiently many centers are used.

We remark that in the plot (and in the paper) for each of the methods two given parameter values are tested. Also in the case of the algorithm of this paper, a quite high sensitivity to the parameter γ of the kernel was observed, even if we report only the results for a representative value.




Fig. 7.4 Errors E_μ (approximation of the mean, left plot) and E_σ (approximation of the standard deviation, right plot) as functions of the number of points obtained by the greedy algorithm in the example of Section 7.4. The results are compared with the ones obtained with arbitrary polynomial chaos expansion (aPC), spatially adaptive sparse grids (aSG), P -greedy kernel interpolation (PGreedy), and Hybrid stochastic Galerkin (HSG) from [21].

8 Future work

The numerical experiments strongly suggest that the theoretical rates obtained for the greedy points are suboptimal, and this behavior would deserve additional investigation. This is particularly evident in cases when the target function or functional has some localized features. In this case, greedy methods seem to be able to localize the interpolant (or quadrature formula), giving a clear advantage over quasi-uniform points. An understanding of this effect beyond numerical evidence is still not available.

In more general terms, it is not clear how to characterize functions or functionals for which optimal quadrature or interpolation points are localized. The experiment of Section 7.1 seems to suggest that this is the case for integration functionals with densities with localized support, but a proof of this fact is missing.

Future work should focus also on the study of the sign of the weights, and on a stability analysis of the greedy quadrature formulas.

Moreover, the sensitivity of the results on the parameter of the kernel (e.g. as observed in Section 7.4) calls for efficient techniques for the estimation of γ , which in this kind of application need especially to rely on few or no additional samples of the expensive function to be integrated. In this setting, a viable and promising option is to use Bayesian methods based on the maximization of a likelihood of the model

give the data. Recently, investigations of these techniques from a deterministic point of view have been published (see [19, 20]). It is anyhow not yet clear how these can be coupled with greedy, or generally adaptive methods.

Acknowledgements We thank Tizian Wenzel for several comments on an early version of this manuscript. Moreover, we thank the two anonymous referees for a number of useful and constructive remarks that helped improving this paper.

References

1. Antil, H., Field, S.E., Herrmann, F., Nocketto, R.H., Tiglio, M.: Two-step greedy algorithm for reduced order quadratures. *Journal of Scientific Computing* **57**(3), 604–637 (2013)
2. Barron, A.R., Cohen, A., Dahmen, W., DeVore, R.A.: Approximation and learning by greedy algorithms. *The Annals of Statistics* **36**(1), 64–94 (2008)
3. Bezhav, A.Yu.: Cubature formulae on scattered meshes. *Russian Journal of Numerical Analysis and Mathematical Modelling* **6**(2), 95–106 (1991)
4. Binev, P., Cohen, A., Dahmen, W., DeVore, R., Petrova, G., Wojtaszczyk, P.: Convergence rates for greedy algorithms in reduced basis methods. *SIAM Journal on Mathematical Analysis* **43**(3), 1457–1472 (2011)
5. Briol, F.X., Oates, C.J., Girolami, M., Osborne, M.A., Sejdinovic, D.: Probabilistic integration: A role in statistical computation? (with discussion and rejoinder). *Statistical Science* **34**(1), 1–22 (2019)
6. Brown, D., Smith, R.J.E., Freise, A.: Fast simulation of Gaussian-mode scattering for precision interferometry. *Journal of Optics* **18**(2), 025,604 (2016). DOI 10.1088/2040-8978/18/2/025604
7. DeVore, R.A., Temlyakov, V.N.: Some remarks on greedy algorithms. *Advances in Computational Mathematics* **5**(2-3), 173–187 (1996)
8. Dutta, S., Farthing, M.W., Perracchione, E., Savant, G., Putti, M.: A greedy non-intrusive reduced order model for shallow water equations (2020)
9. Fasshauer, G., McCourt, M.: Kernel-Based Approximation Methods Using MATLAB, *Interdisciplinary Mathematical Sciences*, vol. 19. World Scientific Publishing (2015)
10. Fasshauer, G.E.: Meshfree Approximation Methods with MATLAB, *Interdisciplinary Mathematical Sciences*, vol. 6. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ (2007)
11. Fuselier, E., Wright, G.B.: Scattered data interpolation on embedded submanifolds with restricted positive definite kernels: Sobolev error estimates. *SIAM Journal on Numerical Analysis* **50**(3), 1753–1776 (2012)
12. Gaß, M., Glau, K.: Parametric integration by magic point empirical interpolation. *IMA Journal of Numerical Analysis* **39**(1), 315–341 (2017)
13. Hardin, D., Saff, E.: Discretizing manifolds via minimum energy points. *Notices of the AMS* **51**(10), 1186–1194 (2004)
14. Huszár, F., Duvenaud, D.: Optimally-weighted herding is Bayesian quadrature. In: 28th Conference on Uncertainty in Artificial Intelligence, pp. 377–385 (2012)
15. Iske, A.: Approximation Theory and Algorithms for Data Analysis, *Texts in Applied Mathematics*, vol. 68. Springer, Cham (2018)
16. Kanagawa, M., Hennig, P.: Convergence guarantees for adaptive Bayesian quadrature methods. In: *Advances in Neural Information Processing Systems*, vol. 32, pp. 6234–6245 (2019)
17. Kanagawa, M., Sriperumbudur, B.K., Fukumizu, K.: Convergence guarantees for kernel-based quadrature rules in misspecified settings. In: *Advances in Neural Information Processing Systems*, vol. 29, pp. 3288–3296 (2016)
18. Kanagawa, M., Sriperumbudur, B.K., Fukumizu, K.: Convergence analysis of deterministic kernel-based quadrature rules in misspecified settings. *Foundations of Computational Mathematics* **20**, 155–194 (2019)
19. Karvonen, T., Tronarp, F., Särkkä, S.: Asymptotics of maximum likelihood parameter estimates for gaussian processes: The ornstein-uhlenbeck prior. In: 2019 IEEE 29th International Workshop on Machine Learning for Signal Processing (MLSP), pp. 1–6 (2019)
20. Karvonen, T., Wynne, G., Tronarp, F., Oates, C., Särkkä, S.: Maximum likelihood estimation and uncertainty quantification for Gaussian process approximation of deterministic functions. *SIAM/ASA Journal on Uncertainty Quantification* **8**(3), 926–958 (2020)

21. Köppel, M., Franzelin, F., Kröker, I., Oladyskhin, S., Santin, G., Wittwar, D., Barth, A., Haasdonk, B., Nowak, W., Pflüger, D., Rohde, C.: Comparison of data-driven uncertainty quantification methods for a carbon dioxide storage benchmark scenario. *Computational Geosciences* **23**(2), 339–354 (2019)
22. Köppel, M., Franzelin, F., Kröker, I., Oladyskhin, S., Wittwar, D., Santin, G., Barth, A., Haasdonk, B., Nowak, W., Pflüger, D., Rohde, C.: Datasets and executables of data-driven uncertainty quantification benchmark in carbon dioxide storage (2017)
23. Larkin, F.M.: Gaussian measure in Hilbert space and applications in numerical analysis. *Rocky Mountain Journal of Mathematics* **2**(3), 379–422 (1972)
24. Müller, S.: Komplexität und Stabilität von kernbasierten Rekonstruktionsmethoden (Complexity and Stability of Kernel-based Reconstructions). Ph.D. thesis, Fakultät für Mathematik und Informatik, Georg-August-Universität Göttingen (2009). URL <https://ediss.uni-goettingen.de/handle/11858/00-1735-0000-0006-B3BA-E>
25. Müller, S., Schaback, R.: A Newton basis for kernel spaces. *Journal of Approximation Theory* **161**(2), 645–655 (2009)
26. Narcowich, F.J., Ward, J.D., Wendland, H.: Sobolev bounds on functions with scattered zeros, with applications to radial basis function surface fitting. *Mathematics of Computation* **74**(250), 743–763 (2005)
27. Novak, E.: Deterministic and Stochastic Error Bounds in Numerical Analysis, *Lecture Notes in Mathematics*, vol. 1349. Springer-Verlag (1988)
28. Novak, E., Woźniakowski, H.: Tractability of Multivariate Problems. Volume I: Linear Information. European Mathematical Society (2008)
29. Oettershagen, J.: Construction of Optimal Cubature Algorithms with Applications to Econometrics and Uncertainty Quantification. Ph.D. thesis, Institut für Numerische Simulation, Universität Bonn (2017)
30. O’Hagan, A.: Bayes–Hermite quadrature. *Journal of Statistical Planning and Inference* **29**(3), 245–260 (1991)
31. Pazouki, M., Schaback, R.: Bases for kernel-based spaces. *Journal of Computational and Applied Mathematics* **236**(4), 575–588 (2011)
32. Rieger, C., Zwicknagl, B.: Sampling inequalities for infinitely smooth functions, with applications to interpolation and machine learning. *Adv. Comput. Math.* **32**(1), 103–129 (2008)
33. Santin, G.: VKOGA, Matlab implementation. <https://gitlab.mathematik.uni-stuttgart.de/pub/ians-anm/vkoga> (2019)
34. Santin, G.: VKOGA, Python implementation. <https://github.com/GabrieleSantin/VKOGA> (2020)
35. Santin, G., Haasdonk, B.: Convergence rate of the data-independent P-greedy algorithm in kernel-based approximation. *Dolomites Research Notes on Approximation* **10**, 68–78 (2017)
36. Schaback, R.: Improved error bounds for scattered data interpolation by radial basis functions. *Mathematics of Computation* **68**(225), 201–216 (1999)
37. Schaback, R.: Greedy sparse linear approximations of functionals from nodal data. *Numerical Algorithms* **67**(3), 531–547 (2014)
38. Schaback, R.: Superconvergence of kernel-based interpolation. *Journal of Approximation Theory* **235**, 1–19 (2018)
39. Sommariva, A., Vianello, M.: Numerical cubature on scattered data by radial basis functions. *Computing* **76**(3–4), 295–310 (2006)
40. Steinwart, I., Scovel, C.: Mercer’s theorem on general domains: On the interaction between measures, kernels, and RKHSs. *Constructive Approximation* **35**(3), 363–417 (2012)
41. Temlyakov, V.N.: Greedy approximation. *Acta Numerica* **17**, 235–409 (2008)
42. Wendland, H.: Scattered Data Approximation, *Cambridge Monographs on Applied and Computational Mathematics*, vol. 17. Cambridge University Press, Cambridge (2005)
43. Wenzel, T., Santin, G., Haasdonk, B.: A novel class of stabilized greedy kernel approximation algorithms: Convergence, stability & uniform point distribution. ArXiv 1911.04352 (2019). Accepted for publication on *Journal of Approximation Theory*
44. Wirtz, D., Haasdonk, B.: A vectorial kernel orthogonal greedy algorithm. *Dolomites Research Notes on Approximation* **6**, 83–100 (2013)
45. Wright, G.B.: SpherePts. <https://github.com/gradywright/spherepts/> (2020)
46. Yano, M.: Discontinuous Galerkin reduced basis empirical quadrature procedure for model reduction of parametrized nonlinear conservation laws. *Advances in Computational Mathematics* (2019)