M. Griebel, P. Oswald

Convergence analysis of online algorithms for vector-valued kernel regression

Convergence analysis of online algorithms for vector-valued kernel regression

Michael Griebel · Peter Oswald

Abstract We consider the problem of approximating the regression function from noisy vector-valued data by an online learning algorithm using an appropriate reproducing kernel Hilbert space (RKHS) as prior. In an online algorithm, i.i.d. samples become available one by one by a random process and are successively processed to build approximations to the regression function. We are interested in the asymptotic performance of such online approximation algorithms and show that the expected squared error in the RKHS norm can be bounded by $C^2(m+1)^{-s/(2+s)}$, where m is the current number of processed data, the parameter $0 < s \leq 1$ expresses an additional smoothness assumption on the regression function, and the constant C depends on the variance of the input noise, the smoothness of the regression function, and other parameters of the algorithm.

1 Introduction

In this paper, we consider the problem of learning the regression function from noisy vector-valued data using an appropriate RKHS as a prior. For relevant background on the theory of kernel methods, see [4, 5, 12, 13, 15] and especially [2, 3, 11] in the vector-valued case. Our focus is on obtaining estimates for the expectation of the squared error norm in the RKHS H of approximations to the regression function which are built in an incremental way by so-called online algorithms. The setting we use is as follows: Let be given $N \leq \infty$ samples $(\omega_m, y_m) \in \Omega \times Y$, $m = 0, \ldots, N-1$, of an input-output process $\omega \to y$, which are i.i.d. with respect to a (generally unknown) probability measure μ defined on $\Omega \times Y$. For simplicity, let Ω be a compact metric space, Y a separable Hilbert space, and μ a Borel measure. What we are looking for is a regression function $f_\mu : \Omega \to Y$ which, in some sense, optimally represents the underlying input-output process. Algorithms for least-squares regression aim to find approximations to the solution

$$f_\mu(\omega) = \mathbb{E}(y|\omega) \in L^2_\rho(\Omega, Y)$$

M. Griebel
Institute for Numerical Simulation, Universität Bonn, Friedrich-Hirzebruch-Allee 7, 53115 Bonn, Germany, and Fraunhofer Institute for Algorithms and Scientific Computing (SCAI), Schloss Birlinghoven, 53754 Sankt Augustin, Germany
Corresponding author, tel.: +49-228-7369829, E-mail: griebel@ins.uni-bonn.de

P. Oswald
Institute for Numerical Simulation, Universität Bonn, Friedrich-Hirzebruch-Allee 7, 53115 Bonn, Germany
E-mail: agp.oswald@gmail.com

of the minimization problem

$$\mathbb{E}(\|f(\omega) - y\|_Y^2) = \int_{\Omega \times Y} \|f(\omega) - y\|_Y^2 \, d\mu(\omega, y) \longmapsto \min \tag{1}$$

for $f \in L_\rho^2(\Omega, Y)$ from the samples $(\omega_m, y_m), m = 0, \ldots, N-1$, where $\rho(\omega)$ is the marginal probability measure generated by $\mu(\omega, y)$ on Ω.[1] For the minimization problem (1) to be meaningful, one needs

$$\mathbb{E}(\|y\|_Y^2) = \int_{\Omega \times Y} \|y\|_Y^2 \, d\mu(\omega, y) = \int_\Omega \mathbb{E}(\|y\|_Y^2 | \omega) \, d\rho(\omega) < \infty.$$

Since solving the discretized least-squares problem

$$\frac{1}{N} \sum_{m=0}^{N-1} \|f(\omega_m) - y_m\|_Y^2 \longmapsto \min \tag{2}$$

for $f \in L_\rho^2(\Omega, Y)$ is an ill-posed problem that makes no sense without further regularization, it is customary to add a prior assumption $f \in H$, where $H \subset L_\rho^2(\Omega, Y)$ is a set of functions $f : \Omega \to Y$ such that point evaluations $\omega \to f(\omega)$ are well-defined. Staying with the Hilbert space setting, candidates for H are vector-valued RKHS

$$H = \{f_v(\omega) = R_\omega^* v : \|f_v\|_H := \|v\|_V, \ v \in V\} \tag{3}$$

introduced by means of a feature map $\mathbf{R} := \{R_\omega\}_{\omega \in \Omega}$, i.e. by a family of bounded linear operators $R_\omega : Y \to V$ mapping into a separable Hilbert space V called feature space. Detailed definitions and necessary assumptions on the feature map are given in the next section. Obviously, the RKHS H and the feature space V are isometric and can be identified.

The online algorithms considered in this paper build, starting from an initial guess $u^{(0)} \in V$, a sequence of successive approximations $u^{(m)} \in V$, where $u^{(m+1)}$ is a linear combination of the previous approximation $u^{(m)}$ and a term that includes the residual $y_m - R_{\omega_m}^* u^{(m)}$ with respect to the currently processed sample (ω_m, y_m). More precisely, the update formula in the feature space V can be written in the form

$$u^{(m+1)}(\omega) = \alpha_m(u^{(m)} + \mu_m R_{\omega_m}(y_m - R_{\omega_m}^* u^{(m)})), \qquad m = 0, 1, \ldots, N-1, \tag{4}$$

By setting $f^{(m)} := f_{u^{(m)}}$, the corresponding iteration in the RKHS H has the equivalent form

$$f^{(m+1)}(\omega) = \alpha_m(f^{(m)}(\omega) + \mu_m K(\omega, \omega_m)(y_m - f^{(m)}(\omega_m))), \qquad m = 0, 1, \ldots, N-1, \tag{5}$$

where $K(\omega, \theta) : Y \to Y, \omega, \theta \in \Omega$, is the operator kernel of the RKHS determined by the feature map. The parameters α_m, μ_m are given by

$$\alpha_m = \frac{m+1}{m+2}, \qquad \mu_m = \frac{A}{(m+1)^t}, \qquad m = 0, 1, \ldots, \tag{6}$$

where the constants $1/2 < t < 1$ and $A > 0$ will be properly fixed later. Following tradition, α_m and μ_m are called regularization and step-size parameters, respectively.

Our main result, namely Theorem 1, concerns a sharp asymptotic estimate for the expected squared error

$$\mathbb{E}(\|u - u^{(m)}\|_V^2) = \mathbb{E}(\|f_u - f^{(m)}\|_H^2), \qquad m \to \infty,$$

of these iterations in the V and H norms, respectively. Here $u \in V$ is the unique minimizer of the minimization problem

$$J(v) := \mathbb{E}(\|f_v - y\|_Y^2) = \int_{\Omega \times Y} \|R_\omega^* v - y\|_Y^2 \, d\mu(\omega, y) \longmapsto \min. \tag{7}$$

---

[1] The symbol $\mathbb{E}$ denotes expectations of random variables with respect to the underlying probability space, which may vary from formula to formula but should be clear from the context.

These estimates hold under standard assumptions on the feature map, the parameters in (6) and the smoothness s of $u \in V$ measured in a scale of smoothness spaces $V_{P_\rho}^s \subset V$ associated with the underlying covariance operator $P_\rho = \mathbb{E}(R_\omega R_\omega^*)$. Our approach is an extension of earlier work [8] on Schwarz iterative methods in the noiseless case, where $y_m = f_u(\omega_m) = R_{\omega_m}^* u$ was assumed.

The remainder of this paper is organized as follows: In Section 2 we introduce vector-valued RKHS, define P_ρ and the associated scale of smoothness spaces $V_{P_\rho}^s$, and discuss properties of the minimization problem (7). This sets the stage for the analysis of our online learning algorithms in V and allows us to formulate our main convergence result, namely Theorem 1. In Section 3 we review related results from the literature and compare them to our new result. In Section 4 we then provide the detailed proof of Theorem 1. In Section 5 we give further remarks on Theorem 1, show the divergence in expectation

$$\|\mathbb{E}(u^{(m)})\|_V \to \infty, \qquad m \to \infty,$$

of the online algorithm (4)-(6) if (7) does not possess a unique minimizer $u \in V$ and consider a simple special case of "learning" an element u of a Hilbert space V from noisy measurements of its coefficients with respect to a complete orthonormal system (CONS) in V.

2 Setting and main result

Let us first introduce our approach to vector-valued RKHS $H \subset L_\rho^2(\Omega, Y)$, where Ω is a compact metric space with Borel probability measure ρ and Y is a separable Hilbert space. Note that such a RKHS H can be implicitly introduced by a family $\mathbf{R} = \{R_\omega\}_{\omega \in \Omega}$ of bounded linear operators $R_\omega : Y \to V$, where V is another separable Hilbert space (we will tactitly assume that V is infinite-dimensional). More precisely, we introduce the notation

$$f_v(\omega) := R_\omega^* v, \quad \omega \in \Omega, \quad v \in V,$$

and set

$$H := \{f_v : v \in V, \quad \|f_v\|_H := \|v\|_V\}.$$

Under the simplifying conditions

$$\|R_\omega^* v\|_Y = 0 \quad \rho - \text{a.e.} \quad \implies \quad v = 0 \qquad \forall\, v \in V, \tag{8}$$

and

$$f_v(\omega) \in C(\Omega, Y) \qquad \forall\, v \in V, \tag{9}$$

we then have $H \subset C(\Omega, Y) \subset L_\rho^2(\Omega, Y)$ and $f_v = 0$ in $L_\rho^2(\Omega, Y)$ implies $v = 0$. By definition, H and V are isometric, which allows us to easily switch between H and V in the sequel.

The condition (9) is the continuity of the operator family $\mathbf{R}$ in the strong operator topology and ensures Bochner integrability of functions from Ω into Y and V, respectively, which appear in the formulas below. Because of the assumed compactness of Ω, it also implies

$$\|R_\omega\|_{Y \to V}^2 \le \Lambda < \infty, \qquad \omega \in \Omega, \tag{10}$$

with some $\Lambda < \infty$. Another consequence is that the operator kernel

$$K(\omega, \theta) := R_\omega^* R_\theta : Y \to Y, \qquad \omega, \theta \in \Omega,$$

associated with the vector-valued RKHS H is a Mercer kernel. The condition (10) is equivalent to the uniform boundedness

$$\|K(\omega, \theta)\|_Y \le \Lambda, \qquad \omega, \theta \in \Omega, \tag{11}$$

of the operator kernel K. Furthermore, (10) is equivalent to the uniform boundedness of the operator family

$$P_\omega := R_\omega R_\omega^* : V \to V, \qquad \omega \in \Omega,$$

in V, i.e.,

$$\|P_\omega\|_V \le \Lambda, \qquad \omega \in \Omega. \tag{12}$$

For fixed V and $\mathbf{R}$ satisfying the above properties, instead of solving the minimization problem (1) on $L_\rho^2(\Omega, Y)$, one now searches for the minimizer $u \in V$ of (7). The solution u of this quadratic minimization problem on V, if it exists, must satisfy the necessary condition

$$\mathbb{E}((R_\omega^* u - y, R_\omega^* w)_Y) = \mathbb{E}((P_\omega u - R_\omega y, w)_V) = 0 \qquad \forall\, w \in V.$$

This condition corresponds to the linear operator equation

$$P_\rho u = \mathbb{E}(R_\omega y), \qquad P_\rho := \mathbb{E}(P_\omega) = \mathbb{E}(R_\omega R_\omega^*), \tag{13}$$

in V. Note that in general we cannot expect that $f_u = f_\mu$, since the closure W of H in $L_\rho^2(\Omega, Y)$ does not necessarily coincide with $L^2(\Omega, Y)$.

The operator $P_\rho : V \to V$ defined in (13), which plays the role of a covariance operator, is bounded and symmetric positive definite. Indeed,

$$\begin{aligned}
(P_\rho v, w)_V &:= \int_{\Omega \times Y} (P_\omega v, w)_V \, d\mu(\omega, y) = \int_\Omega (P_\omega v, w)_V \, d\rho(\omega) \\
&= \int_\Omega (R_\omega^* v, R_\omega^* w)_Y \, d\rho(\omega) = (f_v, f_w)_{L^2(\Omega, Y)} = (v, P_\rho w)_V, \quad v, w \in V,
\end{aligned}$$

and setting $w = v$ we also have $\ker(P_\rho) = \{0\}$ since

$$(P_\rho v, v) = \|f_v\|_{L_\rho^2(\Omega, V)} > 0, \qquad 0 \ne v \in V,$$

by (8). The boundedness of $P_\rho : V \to V$, together with the estimate

$$\|P_\rho\|_V \le \Lambda,$$

follows from (10). The spectrum of P_ρ is thus contained in $[0, \Lambda]$. We also assume that P_ρ is compact. A sufficient condition, which is often satisfied in applications, is the trace class property for P_ρ which holds in particular if the operators R_ω, $\omega \in \Omega$, have uniformly bounded finite rank.

Note that the assumptions on Ω and $\mathbf{R}$ can be weakened, see for example [2], and that the compactness of P_ρ is used only as a technical simplification. In particular, the compactness assumption allows us to define the scale of smoothness spaces $V_{P_\rho}^s$, $s \in \mathbb{R}$, generated by P_ρ using the complete orthonormal system (CONS) $\Psi := \{\psi_k\}$ of eigenvectors of P_ρ and associated eigenvalues $\lambda_1 \ge \lambda_2 \ge \ldots > 0$ with limit 0 in V as follows: $V_{P_\rho}^s$ is the completion of $\mathrm{span}(\Psi)$ with respect to the norm

$$\Big\| \sum_k c_k \psi_k \Big\|_{V_{P_\rho}^s} = \left(\sum_k \lambda_k^{-s} c_k^2 \right)^{1/2},$$

which is well defined on $\mathrm{span}(\Psi)$ for any s. These spaces will appear in the investigation below.

For our convergence analysis in V, we now assume that

$$\mathbb{E}(R_\omega y) \in \mathrm{ran}(P_\rho). \tag{14}$$

In other words, we assume that (7) has a unique solution $u \in V$ for which (13) holds. For general y with $\mathbb{E}(\|y\|_Y^2) < \infty$, only $\mathbb{E}(R_\omega y) \in \mathrm{ran}(P_\rho^{1/2})$ can be established. It can also be proved that the functional $J(v)$ in (7) does not have a global minimum if (14) is violated (the infimum of $J(v)$ is then only approached if $v \to \infty$ in V in a certain way). It will turn out that the condition (14) is necessary for the convergence in V of the online method (4)-(6), see Subsection 5.3.

For a given prior RKHS H induced by the operator family $\mathbf{R}$ with associated space V and for given samples (ω_m, y_m), $m = 0, \ldots, N-1$, with finite N, the standard regularization of the ill-posed problem (2) is to find the minimizer $u_N \in V$ of the minimization problem

$$J_N(v) := \frac{1}{N} \sum_{m=0}^{N-1} \|f_v(\omega_m) - y_m\|_Y^2 + \kappa_N \|v\|_V^2 \longmapsto \min \qquad (15)$$

on V, where $\kappa_N > 0$ is a suitable regularization parameter. Using the representer theorem for Mercer kernels [11], this problem leads to a linear system with a typically dense and ill-conditioned $N \times N$ matrix. There is a huge body of literature, especially in the scalar-valued case $Y = \mathbb{R}$, devoted to setting up, analyzing and solving these problems for fixed N. Alternatively, one can restrict the minimization in (7) to bounded closed subsets in V which, under our assumptions on P_ρ, are compact subsets of $L_\rho^2(\Omega, Y)$, see [5].

We focus here on the online learning algorithm (4)-(6) for finding approximations to the minimizer $u \in V$ of (7) and are interested in the analysis of their asymptotic performance. For this purpose, we define the noise term

$$\varepsilon_\omega := y - f_u(\omega) = y - R_\omega^* u, \qquad \omega \in \Omega,$$

which is a Y-valued random variable on $\Omega \times Y$ (to keep the notation short, the dependence of ε_ω on y is not explicitly shown). By (14) we have $\mathbb{E}(\varepsilon_\omega | \omega) = f_\mu(\omega) - f_u(\omega)$ for any $\omega \in \Omega$. Also, the noise variance

$$\sigma_H^2 := \mathbb{E}(\|\varepsilon_\omega\|_Y^2) = \mathbb{E}(\|y - f_\mu\|_Y^2) + \|f_\mu - f_u\|_{L^2(\Omega, Y)}^2 \qquad (16)$$

with respect to $f_u \in H$ is finite, since $\mathbb{E}(\|y\|_Y^2) < \infty$ was assumed in the first place. The value of σ_H depends on both the average size of the noise $y - f_\mu(\omega)$ on Ω measured in the Y norm and the $L_\rho^2(\Omega, Y)$ distance of f_μ from W.

The online algorithm (4)-(6) is a particular instance of a randomized Schwarz approximation method associated with $\mathbf{R}$. Its noiseless version, where $y_m = R_{\omega_m}^* u$, was studied in [8] under the assumption $u \in V_{P_\rho}^s$, $0 \le s \le 1$, where α_m was as in (6) but μ_m was determined by a steepest descent rule. Our goal in this paper is to derive convergence results for the expected squared error $\mathbb{E}(\|u - u^{(m)}\|_V^2)$, $m = 1, 2, \ldots$. As expected, such estimates again require additional smoothness assumptions on u in the form $u \in V_{P_\rho}^s$ with $0 < s \le 1$. However, in contrast to the noiseless case [8], they include a dependence on the noise variance σ_H^2 in addition to the dependence on the initial error and the smoothness of u. The prize of convergence is a certain decay of the step-sizes $\mu_m \to 0$, as assumed in (6), which is typical for stochastic approximation algorithms. Our main result is as follows:

Theorem 1 *Let Y, V be separable Hilbert spaces, Ω be a compact metric space, μ be a Borel probability measure on $\Omega \times Y$, and ρ be the marginal Borel probability measure on Ω induced by μ. Assume that*

$$\mathbb{E}(\|y\|_Y^2) = \int_{\Omega \times Y} \|y\|_Y^2 \, d\mu < \infty.$$

For the operator family $\mathbf{R} = \{R_\omega\}_{\omega \in \Omega}$, we need the conditions (8) and (10). We also assume that the operator $P_\rho = \mathbb{E}(R_\omega R_\omega^*)$ is compact. Finally, we assume (14) and that $u \in V_{P_\rho}^s$ for some $0 < s \le 1$.

Consider the online learning algorithm (4), where $u^{(0)} \in V$ is arbitrary, the parameters α_m, μ_m are given by (6) with $t = t_s := (1+s)/(2+s)$ and $A = 1/(2\Lambda)$, and the random samples (ω_m, y_m), $m = 0, 1, \ldots, N \le \infty$, are i.i.d. with respect to μ. Then the expected squared error in V satisfies

$$\mathbb{E}(\|u - u^{(m)}\|_V^2) = \mathbb{E}(\|f_u - f^{(m)}\|_H^2) \le C^2 (m+1)^{-s/(2+s)}, \qquad m = 1, 2, \ldots, N+1, \qquad (17)$$

where $f^{(m)} = f_{u^{(m)}}$, the noise variance σ_H^2 is defined in (16), and

$$C^2 = 2\|u - u^{(0)}\|_V^2 + 2\|u\|_V^2 + 8\Lambda^s \|u\|_{V_{P_\rho}^s}^2 + \sigma_H^2/\Lambda.$$

In this generality, Theorem 1 has not yet appeared in the literature, at least to our knowledge. Its proof is given in Section 4. For the parameter range $0 < s \leq 1$, the exponent $-s/(2+s)$ in the right-hand side of (17) is best possible under the general conditions stated in Theorem 1. Estimates of the form (17) also hold for arbitrary values $1/2 < t < 1$ and $0 < A \leq 1/(2\Lambda)$ in (6), albeit with non-optimal exponents depending on t and different constants C varying with t and A. Without the condition (14), which ensures the existence of the minimizer $u \in V$ in (7), the online method (4)-(6) diverges in expectation. Note that convergence estimates for with respect to the weaker $L_\rho^2(\Omega, Y)$ norm are also of great interest, but cannot be obtained within our framework. We will comment on these issues in the concluding Section 5.

There is a huge amount of literature devoted to the convergence theory of various versions of the algorithm (4)-(5), especially for the scalar-valued case $Y = \mathbb{R}$. In particular, it is often considered in the so-called finite horizon case, where $N < \infty$ is fixed and the step-sizes μ_m are chosen in dependence on N so that expectations such as $\mathbb{E}(\|u - u^{(N)}\|_V^2)$ or $\mathbb{E}(\|f_u - f_{u^{(N)}}\|_{L_\rho^2(\Omega, Y)}^2)$, respectively, are optimized for the final approximation $u^{(N)}$. A brief discussion of known results is given in the next section.

3 Results related to Theorem 1

Given the large number of publications on convergence rates for learning algorithms, we will present only a selection of results focusing on the RKHS setting and online algorithms similar to (4)-(5). The results we cite are often stated and proved for the scalar-valued case $Y = \mathbb{R}$, although some authors claim that their methods extend to the case of an arbitrary separable Hilbert space Y with minor modifications. One of the first papers on the vector-valued case is [1], where the authors provide upper bounds in probability for the $L_\rho^2(\Omega, Y)$ error of f_{u_N} if $N \to \infty$ and $\kappa_N \to 0$, where u_N is the solution of (15). These bounds depend in a specific way on the smoothness of $u \in V_{P_\rho}^s$, $0 \leq s \leq 1$, and on the spectral properties of P_ρ. Note that in [1] and in many other papers stronger assumptions on the compactness of P_ρ are made than we do, and that bounds in probability do not automatically imply bounds in expectation. Moreover, the error measured in the $L_\rho^2(\Omega, Y)$ norm is with respect to f_{u_N} and not with respect to approximations such as $f_{u^{(m)}}$, $m \leq N$, which are produced by a special algorithm comparable to (5).

In [14], the authors provide estimates in probability for an algorithm similar to (4)-(5) for the scalar-valued case $Y = \mathbb{R}$. They cover both, convergence in $L_\rho^2(\Omega, \mathbb{R})$ and H norms. There, the main additional assumption needed for the application of certain results from martingale theory is that, for some constant $M_\rho < \infty$, the random variable y satisfies

$$|y| \leq M_\rho$$

a.e. on the support of ρ. If $u^{(0)} = 0$ (as assumed in [14]), then this assumption implies bounds for $\|u - u^{(0)}\|_V = \|u\|_V$ and σ with constants that depend on M_ρ. Up to the specification of constants and using the notation of this paper, the convergence result for the H norm stated in [14, Theorem B] is as follows: Consider the online algorithm (4) with starting value $u^{(0)} = 0$ and parameters

$$\alpha_m = \frac{m + m_0 - 1}{m + m_0}, \qquad \alpha_m \mu_m = \frac{A}{(m + m_0)^{(s+1)/(s+2)}}, \qquad m = 0, 1, \ldots,$$

for some (sufficiently large) m_0 and suitable A. Then, if $u \in V_{P_\rho}^s$ for some $0 < s \leq 2$, we have

$$\mathbb{P}\left(\|u - u^{(m)}\|_V^2 \leq \frac{C}{(m + m_0)^{s/(s+2)}}\right) \geq 1 - \delta, \qquad 0 < \delta < 1, \quad m = 0, 1, \ldots,$$

for some constant $C = C(M_\rho, \|u\|_{V_{P_\rho}^s}, m_0, s, \Lambda, \log(2/\delta)) < \infty$. Here $V = H$ is an RKHS of functions $u : \Omega \to \mathbb{R}$ generated by some scalar-valued Mercer kernel $K : \Omega \times \Omega \to \mathbb{R}$ and

$\Lambda = \max_{\omega \in \Omega} K(\omega, \omega)$. The associated maps R_ω are given by $R_\omega y = yK(\omega, \cdot)$, $y \in \mathbb{R}$. Consequently, $R_\omega^* u = u(\omega)$, $\omega \in \Omega$, corresponds to function evaluation. So, for $0 < s \le 1$, we get the same rate as in our Theorem 1 which was, however, based on the expectation of the squared error in $V = H$ in the more general vector-valued case. What our rather elementary method does not provide is a result for the case $1 < s \le 2$ and for $L_\rho^2(\Omega, Y)$ convergence. For the latter situation, [14, Theorem C] gives the better estimate

$$\mathbb{P}\left(\|u - u^{(m)}\|_{L_\rho^2(\Omega, \mathbb{R})}^2 \le \frac{\bar{C}}{(m + m_0)^{(s+1)/(s+2)}} \right) \ge 1 - \delta, \quad 0 < \delta < 1, \quad m = 0, 1, \ldots,$$

under the same assumptions, but with a different constant

$$\bar{C} = \bar{C}(M_\rho, \|u\|_{V_{P_\rho}^s}, m_0, s, \Lambda, \log(2/\delta)) < \infty.$$

This is almost matching the lower estimates for kernel learning derived in [1] for classes of instances where the spectrum of P_ρ has a prescribed decay of the form $\lambda_k \asymp k^{-b}$ for some $b > 1$. Note that for Mercer kernels, the operator

$$P_\rho : u \in H \quad \longmapsto \quad (P_\rho u)(\cdot) = \int_\Omega K(\cdot, \theta) u(\theta) \, d\rho(\theta)$$

is trace class, while in our Theorem 1 no stronger decay of eigenvalues is assumed.

Estimates in expectation close to our result have also been obtained for slightly different settings. For example, in [16] both, the so-called *regularized* ($\alpha_m < 1$) and the *unregularized* online algorithm ($\alpha_m = 1$) were analyzed in the scalar-valued case $Y = \mathbb{R}$ under assumptions similar to ours for $L_\rho^2(\Omega, \mathbb{R})$ and $V = H$ convergence. We quote only the result for convergence in the RKHS $V = H$. It concerns the so-called *finite horizon* case of the unregularized online algorithm (4) with $\alpha_m = 1$, where one fixes $N < \infty$, chooses a constant step-size $\mu_m = \mu$, $m = 0, \ldots, N-1$, which depends on N, stops the iteration at $m = N$ and asks for a good estimate of the expectation of $\mathbb{E}(\|u - u^{(N)}\|_V^2)$ for the final iterate only. Up to the specification of constants, Theorem 6 in [16] states that, under the condition $u \in V_{P_\rho}^s$, $s > 0$, one can obtain the bound

$$\mathbb{E}(\|u - u^{(N)}\|_V^2) = O(N^{-s/(s+2)}), \qquad N \to \infty,$$

if one sets $\mu_N = cN^{-(s+1)/(s+2)}$ with a properly adjusted value of c. Note that $s > 0$ is arbitrary with the exponent approaching -1 if the smoothness parameter s tends to ∞, while our result gives no improvement for $s > 1$. The drawback of the finite horizon case is that the estimate concerns only a fixed iterate $u^{(N)}$ with an N to be decided beforehand. In a sense, this can be seen as building an approximation to the solution u_N of (15) with $\kappa_N = \mu_N$ from a single pass over the N i.i.d. samples (ω_m, y_m), $m = 0, \ldots, N-1$.

In recent years, attention has shifted to obtaining refined rates when P_ρ possesses faster eigenvalue decay, usually expressed by the property that P_ρ^β is trace class for some $\beta < 1$ or by the slightly weaker assumption

$$\lambda_k = O(k^{-1/\beta}), \qquad k \to \infty, \tag{18}$$

on the eigenvalues of the covariance operator P_ρ. Bounds that require knowledge of $\beta < 1$ are sometimes called capacity dependent, so our bounds in Theorem 1 as well as the cited results from [14, 16] are capacity independent. Capacity dependent convergence rates for the expected squared error for the online algorithm (4)-(5) have been obtained, among others, in [6, 7, 9, 10], again in the scalar-valued case $Y = \mathbb{R}$ and with various parameter settings, including unregularized and finite horizon versions. In [7], rates for $\mathbb{E}(\|u - \bar{u}^{(m)}\|_{L_\rho^2(\Omega, \mathbb{R})}^2)$ have been established, where

$$\bar{u}^{(m)} = \frac{1}{m+1} \sum_{k=0}^m u^{(k)}, \qquad m = 0, 1, \ldots, \tag{19}$$

is the sequence of averages associated with the sequence $u^{(m)}$, $m = 0, 1, \ldots$, obtained by the unregularized iteration (4) with $\alpha_m = 1$ and $u^{(0)} = 0$. That averaging has a similar effect as regularization with $\alpha_m = (m+1)/(m+2)$ in (4) considered in Theorem 1 can be guessed if one observes that

$$\bar{u}^{(m+1)} = \frac{m+1}{m+2} \bar{u}^{(m)} + \frac{1}{m+2} u^{(m+1)},$$

where $u^{(m+1)} = u^{(m)} + \mu_m(y_m - R_{\omega_m} u^{(m)})$, and compares with the regularized iteration (4) with α_m given in (6). To illustrate the influence of β, we formulate the following bound, which is a consequence of [7, Corollary 3]: Under an additional technical assumption on the noise term ε_ω, if the condition (18) holds for some $0 < \beta < 1$ and $u \in V_{P_\rho}^s$, $s > -1$, then for suitable choices for the learning rates μ_m, we have

$$\mathbb{E}(\|u - \bar{u}^{(m)}\|_{L_\rho^2(\Omega, \mathbb{R})}^2) = \begin{cases} O((m+1)^{-(s+1)}), & -1 < s < -\beta, \\ O((m+1)^{-(s+1)/(s+1+\beta)}), & -\beta < s < 1 - \beta, \\ O((m+1)^{-(1-\beta/2)}), & 1 - \beta < s. \end{cases}$$

Thus, stronger eigenvalue decay generally implies stronger asymptotic error decay in the $L_\rho^2(\Omega, \mathbb{R})$ norm. In [6, Section 6], similar rates are obtained in the finite horizon setting for both, the above averaged iterates $\bar{u}^{(N)}$ and for $u^{(N)}$ produced by a two-step extension of the one-step iteration (4).

In addition to $L_\rho^2(\Omega, \mathbb{R})$ convergence results, the paper [10] also provides a capacity dependent convergence estimate in the RKHS norm for the unregularized algorithm (4)-(5) with parameters $\alpha_m = 1$ and $\mu_m = c(m+1)^{-1/2}$. Under the boundedness assumption $|y| \leq M_\rho$, Theorem 2 in [10] implies that

$$\mathbb{E}(\|u - u^{(m)}\|_V^2) = O((m+1)^{-\min(s, 1-\beta)/2} \log^2(m+1)), \qquad m = 1, 2, \ldots,$$

if $u \in V_{P_\rho}^s$ for some $s > 0$, P_ρ^β is trace class for some $0 < \beta < 1$, and c is properly adjusted.

Finally, the scalar-valued least-squares regression problem with $Y = \mathbb{R}$ and RKHS prior space H can also be cast as linear regression problem in $V = H$. This was done in [6, 9]. More abstractly, given a μ-distributed random variable $(\xi_\omega, y) \in V \times \mathbb{R}$ on $\Omega \times \mathbb{R}$, the task is to find approximations to the minimizer $u \in V$ of the problem

$$\mathbb{E}(|(\xi_\omega, v) - y|^2) \longmapsto \min, \qquad v \in V, \tag{20}$$

from i.i.d. samples (ξ_{ω_i}, y_i). If $V = H$ is the RKHS that regularizes the scalar-valued least-squares regression problem on $\Omega \times \mathbb{R}$, then the canonical choice is $\xi_\omega = K(\omega, \cdot)$. In [9], for the iteration

$$u^{(m+1)} = u^{(m)} + \mu_m(y_m - (\xi_{\omega_m}, u^{(m)}))\xi_{\omega_m}, \qquad m = 0, 1, \ldots,$$

weak convergence in V is studied by deriving estimates for quantities such as $\mathbb{E}((v, u - u^{(m)})^2)$ and $\mathbb{E}((\xi_{\omega'}, u - u^{(m)})^2)$ under some simplifying assumptions on the noise and the normalization $\|\xi_\omega\| = 1$. Note that this iteration is nothing but the unregularized iteration (4) with $\alpha_m = 1$ since $(\xi_{\omega_m}, u^{(m)}) = u^{(m)}(\omega_m)$ in this case. In the learning application, the assumption $\|\xi_\omega\| = 1$ means $K(\omega, \omega) = 1$. Moreover, in this case

$$\mathbb{E}((\xi_{\omega'}, u - u^{(m)})^2) = \mathbb{E}(\|u - u^{(m)}\|_{L_\rho^2(\Omega, \mathbb{R})}),$$

since the expectation on the left, in addition to the i.i.d. samples (ξ_{ω_k}, y_k), $k = 0, \ldots, m - 1$, is also taken with respect to the independently ρ-distributed random variable $\xi_{\omega'}$. This is related to learning rates in the $L_\rho^2(\Omega, \mathbb{R})$ norm. The estimates for $\mathbb{E}((\xi_{\omega'}, u - u^{(m)})^2)$ given in [9] concern both, the finite horizon and the online setting and again depend on the parameters $s \geq 0$ (smoothness of u) and $0 < \beta \leq 1$ (capacity assumption on P_ρ). For the estimates of $\mathbb{E}((v, u - u^{(m)})^2)$, the smoothness $s' \geq 0$ of the fixed element $v \in V_{P_\rho}^{s'}$ is traded against the smoothness $s \geq 0$ of $u \in V_{P_\rho}^s$. See [9] for details.

4 Proof of Theorem 1

In this subsection we will use the notation and assumptions outlined above, with the only change that the scalar product in V is simply denoted by $(\cdot, \cdot)$ and the associated norm $\|\cdot\|_V$ is accordingly denoted by $\|\cdot\|$. We also set $e^{(m)} := u - u^{(m)}$. We will prove an estimate of the form

$$\mathbb{E}(\|e^{(m)}\|^2) = \mathrm{O}((m+1)^{-s/(2+s)}), \qquad m \to \infty, \tag{21}$$

under the assumption $u \in V_{P_\rho}^s$, $0 < s \leq 1$, if the parameters A and t in (6) are chosen accordingly. The precise statement and the dependence of the constant in (21) on the initial error, the noise variance and the smoothness assumption are given in the formulation of Theorem 1.

From (4) and $y_m = R_{\omega_m}^* u + \varepsilon_{\omega_m}$ we derive the error representation

$$e^{(m+1)} = \underbrace{\alpha_m(e^{(m)} - \mu_m P_{\omega_m} e^{(m)}) + \bar{\alpha}_m u}_{\bar{e}^{(m+1)}:=} - \alpha_m \mu_m R_{\omega_m} \varepsilon_{\omega_m},$$

where $\bar{\alpha}_m := 1 - \alpha_m = (m+2)^{-1}$, compare also 6. The first term $\bar{e}^{(m+1)}$ corresponds to the noiseless case considered in [8], while the remainder term is the noise contribution. Thus,

$$\|e^{(m+1)}\|^2 = \|\bar{e}^{(m+1)}\|^2 - 2\alpha_m \mu_m (R_{\omega_m} \varepsilon_{\omega_m}, \bar{e}^{(m+1)}) + \alpha_m^2 \mu_m^2 \|R_{\omega_m} \varepsilon_{\omega_m}\|^2. \tag{22}$$

We now estimate the conditional expectation with respect to the given $u^{(m)}$, separately for the three terms in (22). Here and in the following we denote this conditional expectation by $\mathbb{E}'$. For the third term, by (10) and the definition (16) of the noise variance σ_H^2, we have

$$\mathbb{E}'(\|R_{\omega_m} \varepsilon_{\omega_m}\|^2) \leq \Lambda \mathbb{E}(\|\varepsilon_\omega\|_Y^2) = \Lambda \sigma_H^2. \tag{23}$$

For the second term, we need

$$\mathbb{E}((R_\omega \varepsilon_\omega, w)) = \mathbb{E}((y - R_\omega^* u, R_\omega^* w)_Y) = 0 \qquad \forall\, w \in V.$$

This follows directly from the fact that $u \in V$ is the minimizer of the problem (7). Thus, by setting $w = \alpha_m e^{(m)} + \bar{\alpha}_m u$, we get

$$\begin{aligned}
\mathbb{E}'(-2\alpha_m \mu_m &(R_{\omega_m} \varepsilon_{\omega_m}, \bar{e}^{(m+1)})) \\
&= 2\alpha_m \mu_m (\alpha_m \mu_m \mathbb{E}'((R_{\omega_m} \varepsilon_{\omega_m}, P_{\omega_m} e^{(m)})) - \mathbb{E}'((R_{\omega_m} \varepsilon_{\omega_m}, w))) \\
&= 2\alpha_m^2 \mu_m^2 \mathbb{E}'((R_{\omega_m} \varepsilon_{\omega_m}, P_{\omega_m} e^{(m)})) \\
&\leq \alpha_m^2 \mu_m^2 (\mathbb{E}'(\|R_{\omega_m} \varepsilon_{\omega_m}\|^2) + \mathbb{E}'(\|P_{\omega_m} e^{(m)}\|^2)).
\end{aligned}$$

Here, the first term is estimated by (23). For the second term, we substitute the upper bound

$$\mathbb{E}'(\|P_{\omega_m} e^{(m)}\|^2) \leq \Lambda \mathbb{E}'((P_{\omega_m} e^{(m)}, e^{(m)})) = \Lambda(P_\rho e^{(m)}, e^{(m)}), \tag{24}$$

which follows from (10) and the definition of P_ρ. Together this gives

$$\mathbb{E}'(-2\alpha_m \mu_m (R_{\omega_m} \varepsilon_{\omega_m}, \bar{e}^{(m+1)})) \leq \Lambda \alpha_m^2 \mu_m^2 (\sigma_H^2 + (P_\rho e^{(m)}, e^{(m)})) \tag{25}$$

for the second term in (22).

To estimate the first term $\mathbb{E}'(\|\bar{e}^{(m+1)}\|^2)$, we modify the arguments from [8], where the case $\varepsilon_m = 0$ was treated. We use the error decomposition

$$\begin{aligned}
\|\bar{e}^{(m+1)}\|^2 = \bar{\alpha}_m^2 \|u\|^2 &+ 2\alpha_m \bar{\alpha}_m (u, e^{(m)} - \mu_m P_{\omega_m} e^{(m)} \\
&+ \alpha_m^2 (\|e^{(m)}\|^2 - 2\mu_m (e^{(m)}, P_{\omega_m} e^{(m)})) + \mu_m^2 \|P_{\omega_m} e^{(m)}\|^2).
\end{aligned}$$

After taking conditional expectations, we arrive with the definition of P_ρ and (24) at

$$\begin{aligned}
\mathbb{E}'(\|\bar{e}^{(m+1)}\|^2) &= \bar{\alpha}_m^2\|u\|^2 + 2\alpha_m\bar{\alpha}_m(u, e^{(m)} - \mu_m P_\rho e^{(m)}) \\
&\quad + \alpha_m^2(\|e^{(m)}\|^2 - 2\mu_m(e^{(m)}, P_\rho e^{(m)}) + \mu_m^2\mathbb{E}'(\|P_{\omega_m}e^{(m)}\|^2)) \\
&\leq \bar{\alpha}_m^2\|u\|^2 + 2\alpha_m\bar{\alpha}_m(u, e^{(m)} - \mu_m P_\rho e^{(m)}) \\
&\quad + \alpha_m^2(\|e^{(m)}\|^2 - \mu_m(2 - \Lambda\mu_m)(e^{(m)}, P_\rho e^{(m)})).
\end{aligned}$$

Next, to estimate the term $(u, e^{(m)} - \mu_m P_\rho e^{(m)})$, we take an arbitrary $h = P_\rho^{1/2}v \in V_{P_\rho}^1$, where $v \in V = V_{P_\rho}^0$ and $\|h\|_{V_{P_\rho}^1} = \|v\|$. This gives us

$$\begin{aligned}
&2\alpha_m\bar{\alpha}_m(u, e^{(m)} - \mu_m P_\rho e^{(m)}) \\
&= 2\alpha_m\bar{\alpha}_m((u - h, (I - \mu_m P_\rho)e^{(m)}) + (h, (I - \mu_m P_\rho)e^{(m)})) \\
&\leq 2\alpha_m\bar{\alpha}_m\|u - h\|\|(I - \mu_m P_\rho)e^{(m)}\| + 2(\bar{\alpha}_m\mu_m^{-1/2}(I - \mu_m P_\rho)v, \alpha_m\mu_m^{1/2}e^{(m)}) \\
&\leq 2\alpha_m\bar{\alpha}_m\|u - h\|\|e^{(m)}\| + \bar{\alpha}_m^2\mu_m^{-1}\|(I - \mu_m P_\rho)v\|^2 + \alpha_m^2\mu_m\|P_\rho^{1/2}e^{(m)}\|^2 \\
&\leq 2\alpha_m\bar{\alpha}_m\|u - h\|\|e^{(m)}\| + \bar{\alpha}_m^2\mu_m^{-1}\|h\|_{V_{P_\rho}^1}^2 + \alpha_m^2\mu_m(P_\rho e^{(m)}, e^{(m)}).
\end{aligned}$$

Here we have silently used that $\|(I - \mu_m P_\rho)e^{(m)}\| \leq \|e^{(m)}\|$ and similarly

$$\|(I - \mu_m P_\rho)v\| \leq \|v\| = \|h\|_{V_{P_\rho}^1},$$

which holds since $0 < \mu_m \leq A \leq (2\Lambda)^{-1}$ according to (6) and the restriction on A. Substitution into the previous inequality yields

$$\begin{aligned}
\mathbb{E}'(\|\bar{e}^{(m+1)}\|^2) &\leq \bar{\alpha}_m^2(\|u\|^2 + \mu_m^{-1}\|h\|_{V_{P_\rho}^1}^2) + 2\alpha_m\bar{\alpha}_m\|u - h\|\|e^{(m)}\| \\
&\quad + \alpha_m^2(\|e^{(m)}\|^2 - \mu_m(1 - \Lambda\mu_m)(e^{(m)}, P_\rho e^{(m)})).
\end{aligned}$$

Now, combining this estimate for the conditional expectation of the first term in (22) with the bounds (23) and (25) for the third and second terms, respectively, we get

$$\begin{aligned}
\mathbb{E}'(\|\bar{e}^{(m+1)}\|^2) &\leq \alpha_m^2(\|e^{(m)}\|^2 + 2\Lambda\sigma_H^2\mu_m^2) \\
&\quad + 2\alpha_m\bar{\alpha}_m\|u - h\|\|e^{(m)}\| + \bar{\alpha}_m^2(\|u\|^2 + \mu_m^{-1}\|h\|_{V_{P_\rho}^1}^2).
\end{aligned} \tag{26}$$

Here the term involving $(e^{(m)}, P_\rho e^{(m)}) \geq 0$ has been omitted, since its forefactor $-\mu_m(1 - 2\Lambda\mu_m)$ is non-positive due to the restriction on A in (6).

For given

$$u = \sum_k c_k\psi_k \in V_{P_\rho}^s, \qquad 0 < s \leq 1,$$

in (26) we choose

$$h = \sum_{k:\,\lambda_k(m+1)^b \geq B} c_k\psi_k$$

with some fixed constants $b, B > 0$ specified below. This gives

$$\|h\|_{V_{P_\rho}^1}^2 = \sum_{k:\,\lambda_k(m+1)^b \geq B} \lambda_k^{-(1-s)}(\lambda_k^{-s}c_k^2) \leq B^{-(1-s)}(m+1)^{(1-s)b}\|u\|_{V_{P_\rho}^s}^2$$

and

$$\|u - h\|^2 = \sum_{k:\,\lambda_k(m+1)^b < B} \lambda_k^s(\lambda_k^{-s}c_k^2) \leq B^s(m+1)^{-bs}\|u\|_{V_{P_\rho}^s}^2.$$

After substitution into (26), we get

$$\mathbb{E}'(\|\bar{e}^{(m+1)}\|^2) \leq \alpha_m^2(\|e^{(m)}\|^2 + 2\Lambda\sigma_H^2\mu_m^2) + 2\alpha_m\bar{\alpha}_m B^{s/2}(m+1)^{-bs/2}\|u\|_{V_{P_\rho}^s}\|e^{(m)}\| \tag{27}$$

$$+ \bar{\alpha}_m^2(\|u\|^2 + \mu_m^{-1}B^{-(1-s)}(m+1)^{(1-s)b}\|u\|_{V_{P_\rho}^s}^2).$$

Obviously, if $s = 1$, we can set $h = u$ which greatly simplifies the considerations below and leads to a more accurate final estimate, see Section 5.1.

Next, we switch to full expectations in (27) by using the independence assumption for the sampling process and taking into account that

$$\varepsilon_m := \mathbb{E}(\|e^{(m)}\|^2)^{1/2} \geq \mathbb{E}(\|e^{(m)}\|).$$

Together with (6) and $\alpha_m = (m+1)\bar{\alpha}_m$, this gives

$$\varepsilon_{m+1}^2 \leq \alpha_m^2(\varepsilon_m^2 + 2A^2\Lambda\sigma_H^2(m+1)^{-2t}) + 2\alpha_m\bar{\alpha}_m B^{s/2}(m+1)^{-bs/2}\|u\|_{V_{P_\rho}^s}\varepsilon_m$$

$$+ \bar{\alpha}_m^2(\|u\|^2 + A^{-1}B^{-(1-s)}(m+1)^{(1-s)b+t}\|u\|_{V_{P_\rho}^s}^2)$$

$$\leq \alpha_m^2\varepsilon_m^2 + \bar{\alpha}_m^2(2A^2\Lambda\sigma_H^2(m+1)^{2-2t} + 2B^{s/2}(m+1)^{-bs/2+1}\|u\|_{V_{P_\rho}^s}\varepsilon_m\|e^{(m)}\|$$

$$+ \|u\|^2 + A^{-1}B^{-(1-s)}(m+1)^{(1-s)b+t}\|u\|_{V_{P_\rho}^s}^2).$$

In a final step, we assume for a moment that

$$\varepsilon_k \leq C(k+1)^{-r}, \qquad k = 0,\ldots,m, \tag{28}$$

holds for some constants $C, r > 0$. Next, we set

$$a := \max(2 - 2t, -bs/2 + 1 - r, (1-s)b + t)$$

and

$$D := 2A^2\Lambda\sigma_H^2 + 2CB^{s/2}\|u\|_{V_{P_\rho}^s} + \|u\|^2 + A^{-1}B^{-(1-s)}\|u\|_{V_{P_\rho}^s}^2.$$

Since $1/2 < t < 1$ is assumed in (6), we have $a > 0$. Then, for $k = 0, 1, \ldots, m$, the estimate for ε_{k+1} simplifies to

$$\varepsilon_{k+1}^2 \leq \alpha_k^2\varepsilon_k^2 + D\bar{\alpha}_k^2(k+1)^a$$

or, since $\alpha_k^2\bar{\alpha}_{k-1}^2 = \bar{\alpha}_k^2$, to

$$d_{k+1} := \bar{\alpha}_k^{-2}\varepsilon_{k+1}^2 \leq \alpha_k^2\bar{\alpha}_k^{-2}\varepsilon_k^2 + D(k+1)^a = d_k + D(k+1)^a.$$

By recursion we get

$$d_{m+1} \leq d_0 + D\sum_{k=0}^{m}(k+1)^a = \varepsilon_0^2 + D\sum_{k=0}^{m}(k+1)^a$$

and finally

$$\varepsilon_{m+1}^2 \leq (m+2)^{-2}(\|e^{(0)}\|^2 + D(m+2)^{a+1}) < (\|e^{(0)}\|^2 + D)(m+2)^{a-1},$$

since we have $a > 0$ and

$$\sum_{k=0}^{m}(k+1)^a \leq \int_1^{m+2} x^a\,dx < (m+2)^{a+1}. \tag{29}$$

So (28) holds by induction for all m if we ensure that

$$1 - a \geq 2r, \qquad \|e^{(0)}\|^2 + D \leq C^2. \tag{30}$$

To complete the proof of Theorem 1, it remains to maximize r for given $0 < s \le 1$. For this purpose, it is intuitively clear to require

$$a = 1 - 2r = 2 - 2t = -bs/2 + 1 - r = (1-s)b + t.$$

This system of equations has the unique solution

$$t = \frac{1+s}{2+s}, \quad b = \frac{1}{2+s}, \quad 2r = \frac{s}{2+s}, \quad a = \frac{2}{2+s}.$$

Furthermore, the appropriate value for C in (28) must satisfy

$$\|e^{(0)}\|^2 + \|u\|^2 + 2A^2 \Lambda \sigma_H^2 + 2CB^{s/2}\|u\|_{V_{P_\rho}^s} + A^{-1}B^{-(1-s)}\|u\|_{V_{P_\rho}^s}^2 \le C^2.$$

With such choices for t and C, the condition (30) is guaranteed, and (28) gives the desired bound

$$\varepsilon_m^2 \le C^2(m+1)^{-s/(s+2)}, \qquad m = 1, 2, \ldots, N-1.$$

By choosing concrete values for $0 < A \le (2\Lambda)^{-1}$ and $B > 0$, the constant C^2 can be made more explicit. For example, substituting the upper bound

$$2CB^{s/2}\|u\|_{V_{P_\rho}^s} \le \frac{C^2}{2} + 2B^s\|u\|_{V_{P_\rho}^s}^2$$

and rearranging the terms shows that

$$C^2 = 2\left(\|e^{(0)}\|^2 + \|u\|^2 + B^s(2 + (AB)^{-1})\|u\|_{V_{P_\rho}^s}^2 + 2A^2\Lambda\sigma_H^2\right)$$

is appropriate. In particular, setting for simplicity A to its maximal value $A = (2\Lambda)^{-1}$ and taking $B = \Lambda$ gives a more explicit dependence of C^2 on the assumptions on $\|e^{(0)}\|^2$, the noise variance σ_H^2, and the smoothness of u, namely

$$C^2 = 2\|e^{(0)}\|^2 + 2\|u\|^2 + 8\Lambda^s\|u\|_{V_{P_\rho}^s} + \sigma_H^2/\Lambda. \tag{31}$$

This is the constant shown in the formulation of Theorem 1. Obviously, varying A and B changes the trade-off between initial error, noise variance, and smoothness assumptions in the convergence estimate (28). Note also that B is not part of the algorithm and can be adjusted to any value. This concludes the proof of Theorem 1.

5 Further remarks

5.1 Comments on Theorem 1

In the special case $s = 1$, the proof of Theorem 2 is simplified as follows: In (26) we can set $h = u$, and (27) is therefore simplified to

$$\mathbb{E}'(\|\bar{e}^{(m+1)}\|^2) \le \alpha_m^2(\|e^{(m)}\|^2 + 2\Lambda\sigma_H^2\mu_m^2) + \bar{\alpha}_m^2(\|u\|^2 + \mu_m^{-1}\|u\|_{V_{P_\rho}^1}^2). \tag{32}$$

So with $\mu_m = A(m+1)^{-t}$ we directly get a recursion for

$$d_m := \bar{\alpha}_{m-1}^{-2}\varepsilon_m^2 = (m+1)^2\mathbb{E}(\|e^{(m)}\|^2)$$

in the form

$$d_{m+1} \le d_m + (2A^2\Lambda\sigma_H^2(m+1)^{2-2t} + \|u\|^2 + A^{-1}(m+1)^t\|u\|_{V_{P_\rho}^1}^2).$$

Taking into account (29), for $1/2 < t < 1$ we finally arrive at

$$\mathbb{E}(\|e^{(m)}\|^2) \leq \frac{\|e^{(0)}\|^2}{(m+1)^2} + \frac{2A^2\Lambda\sigma_H^2}{(m+1)^{2t-1}} + \frac{\|u\|^2}{m+1} + \frac{A^{-1}\|u\|_{V_{P_\rho}^1}^2}{(m+1)^{1-t}}, \qquad (33)$$

$m = 1, 2, \ldots$. This estimate shows more clearly the guaranteed error decay with respect to the initial error $\|e^{(0)}\|^2$, the noise variance σ_H^2, and the norms $\|u\|^2$ and $\|u\|_{V_{P_\rho}^1}^2$ of the solution u, respectively, in dependence on t. The asymptotically dominant term is of the form

$$O((m+1)^{-\min(2t-1, 1-t)})$$

and is minimized when $t = 2/3$. For this value of t and with $A = (2\Lambda)^{-1}$ we get

$$\mathbb{E}(\|e^{(m)}\|^2) \leq \frac{\|e^{(0)}\|^2}{(m+1)^2} + \frac{\|u\|^2}{m+1} + \frac{2\Lambda\|u\|_{V_{P_\rho}^1}^2 + (2\Lambda)^{-1}\sigma_H^2}{(m+1)^{1/3}}. \qquad (34)$$

Without further assumptions one cannot expect a better error decay rate, see Section 3 and Subsection 5.4.

Another comment concerns the finite horizon setting, which is often treated instead of a true online method. Here one fixes a finite N, chooses a constant learning rate $\mu_m = \mu$ for $m = 0, \ldots, N-1$ in dependence on N, and only asks for a best possible bound for $\mathbb{E}(\|e^{(N)}\|^2)$. Our approach easily provides results for this case as well. We demonstrate this only for $s = 1$. For fixed $\mu_m = \mu$, the error recursion for the quantities d_m now takes the form

$$d_{m+1} \leq d_m + (2A^2\Lambda\sigma_H^2(m+1)^2\mu^2 + \|u\|^2 + \mu^{-1}\|u\|_{V_{P_\rho}^1}^2), \quad m = 0, \ldots, N-1,$$

and gives

$$\mathbb{E}(\|e^{(N)}\|^2) \leq \frac{\|e^{(0)}\|^2}{(N+1)^2} + 2\Lambda\sigma_H^2\mu^2(N+1) + \frac{\|u\|^2 + \mu^{-1}\|u\|_{V_{P_\rho}^1}^2}{N+1}.$$

Setting $\mu = (2\Lambda)^{-1}(N+1)^{-2/3}$ results in a final estimate for the finite horizon case similar to (34), but only for $m = N$.

There are obvious drawbacks of the whole setting in which Theorem 1 is formulated. First, the assumptions are at most qualitative: Since μ, and thus ρ, is usually not at our disposal, we cannot verify the assumption $u \in V_{P_\rho}^s$, nor assess the value of σ_H^2. Moreover, although the restriction to learning rates μ_m of the form (6) may not cause problems in view of the results obtained, the choice of optimal values for t and A is by no means obvious. A rule for the adaptive choice of μ_m that does not require knowledge of the values of s and the size of the norms of u, but leads to the same quantitative error decay as guaranteed by Theorem 1, would be desirable.

5.2 Difficulties with convergence in $L_\rho^2(\Omega, Y)$

Our result for the vector-valued case concerned convergence in V, which is isometric to the RKHS H generated by $\mathbf{R}$. What we did not succeed in is extending our methods to obtain better asymptotic convergence rates of $f_{u^{(m)}} \to f_u$ in the $L_\rho^2(\Omega, Y)$ norm. Under the assumption (14) about the existence of the minimizer u in (7), error estimates in the $L_\rho^2(\Omega, Y)$ norm require the study of $\mathbb{E}(\|P_\rho^{1/2} e^{(m)}\|^2) = \mathbb{E}((P_\rho e^{(m)}, e^{(m)}))$ instead of $\mathbb{E}(\|e^{(m)}\|^2)$. If, in analogy to (22), one examines the error decomposition

$$\|P_\rho^{1/2} e^{(m+1)}\|^2 \leq \|P_\rho^{1/2} \bar{e}^{(m+1)}\|^2 - 2\alpha_m\mu_m(P_\rho R_{\omega_m}\varepsilon_{\omega_m}, \bar{e}^{(m+1)}) + \alpha_m^2\mu_m^2\|P_\rho^{1/2} R_{\omega_m}\varepsilon_{\omega_m}\|^2,$$

then difficulties mostly arise from the first term in the right-hand side. In fact, we have

$$\|P_\rho^{1/2}\bar{e}^{(m+1)}\|^2 = \bar{\alpha}_m^2\|P_\rho^{1/2}u\|^2 + 2\alpha_m\bar{\alpha}_m(P_\rho u, e^{(m)} - \mu_m P_{\omega_m}e^{(m)})$$
$$+ \alpha_m^2(\|P_\rho^{1/2}e^{(m)}\|^2 - 2\mu_m(P_\rho e^{(m)}, P_{\omega_m}e^{(m)})) + \mu_m^2\|P_\rho^{1/2}P_{\omega_m}e^{(m)}\|^2).$$

After taking conditional expectations $\mathbb{E}'(\|P_\rho^{1/2}\bar{e}^{(m+1)}\|^2)$, we get a negative term

$$-2\alpha_m^2\mu_m\|P_\rho e^{(m)}\|^2$$

on the right-hand side which must compensate for positive contributions from terms such as

$$\mathbb{E}'(\|P_\rho^{1/2}P_{\omega_m}e^{(m)}\|^2).$$

Since in general P_ρ does not commute with the operators P_ω, this strategy does not work without additional assumptions.

5.3 A divergence result

If the crucial condition (14) in Theorem 1 does not hold, i.e., if

$$g := \mathbb{E}(R_\omega y) = \mathbb{E}(R_\omega f_\mu(\omega)) \notin \mathrm{ran}(P_\rho). \tag{35}$$

then the sequence $u^{(m)}$ obtained from the online algorithm (4)-(6) diverges in expectation to ∞ in V for any choice of the parameters $1/2 < t < 1$ and $0 < A \le (2\Lambda)^{-1}$. This negative result is equivalent to proving

$$\|\mathbb{E}(u^{(m)})\|^2 \to \infty, \qquad m \to \infty, \tag{36}$$

and shows that (14) is essential in Theorem 1. As before, norm and scalar product in V are denoted by $\|\cdot\|$ and $(\cdot, \cdot)$, respectively.

To establish (36), we expand g and $u^{(0)}$ with respect to the CONS Ψ, and derive an inhomogeneous linear recursion for the coefficients of the expected error trajectory $U^{(m)} := \mathbb{E}(u^{(m)})$. To do this, set

$$u^{(0)} = \sum_{i\in\mathbb{N}} c_i\psi_i, \qquad g = \sum_{i\in\mathbb{N}} g_i\psi_i,$$

where $c_i = (u^{(0)}, \psi_i)$, $g_i = (g, \psi_i)$, and denote $c_i^{(m)} := (U^{(m)}, \psi_i)$, $i \in \mathbb{N}$. Obviously, $u^{(0)} = U^{(0)}$ and thus $c_i = c_i^{(0)}$. From (4) we have

$$U^{(m+1)} = \alpha_m(U^{(m)} + \mu_m(g - P_\rho U^{(m)})) = \alpha_m(I - \mu_m P_\rho)U^{(m)} + \alpha_m\mu_m g, \quad m \ge 0$$

which, by iteration using the properties of the regularization parameters $\alpha_m = (m+1)/(m+2)$, immediately yields

$$U^{(m)} = \frac{1}{(m+1)}\left(\prod_{l=0}^{m-1}(I - \mu_l P_\rho)u^{(0)} + \sum_{k=1}^{m} k\mu_{k-1}\prod_{l=k}^{m-1}(I - \mu_l P_\rho)g\right),$$

m=1,2,.... By linearity of the expectation operator and the fact that Ψ consists of the eigenvectors of P_ρ, we get

$$c_i^{(m)} = \frac{1}{(m+1)}\left(\underbrace{\prod_{l=0}^{m-1}(1 - \mu_l\lambda_i)}_{\varepsilon_m(\lambda_i)}c_i + \underbrace{(\sum_{k=1}^{m} k\mu_{k-1}\prod_{l=k}^{m-1}(1 - \mu_l\lambda_i))}_{d_m(\lambda_i):=}g_i\right), \tag{37}$$

$m = 1, 2, \ldots$, separately for each $i \in \mathbb{N}$. Under the restrictions on the parameters μ_m from (6), we have $1/2 \leq 1 - \mu_l \lambda_i \leq 1$. This shows that we have $0 < \varepsilon_m(\lambda_i) \leq 1$ for the factor in front of c_i, which implies that the contribution of the initial guess $u^{(0)}$ can be neglected for $m \to \infty$.

Next, we focus on lower bounds for the factor $d_m(\lambda_i)$ in front of g_i in (37). For our purposes it is sufficient to show that, for some $m_0 \geq 1$ depending on t,

$$d_m(\lambda) \geq c_0 (m+1) \lambda^{-1}, \qquad c_1 (m+1)^{t-1} \leq \lambda \leq \Lambda, \qquad m \geq m_0, \tag{38}$$

with constants c_0, c_1 independent of λ and m. In fact, with (38) at hand, we have

$$\|U^{(m)}\|^2 = (m+1)^{-2} \sum_{i \in \mathbb{N}} (\varepsilon_m(\lambda_i) c_i + d_m(\lambda_i) g_i)^2$$

$$\geq (m+1)^{-2} \left(\frac{1}{2} \sum_{i \in \mathbb{N}} d_m(\lambda_i)^2 g_i^2 - \sum_{i \in \mathbb{N}} \varepsilon_m(\lambda_i)^2 c_i^2 \right)$$

$$\geq \frac{c_0^2}{2} \left(\sum_{\lambda_i \geq c_0 (m+1)^{t-1}} \lambda_i^{-2} g_i^2 \right) - \frac{\|u^{(0)}\|^2}{(m+1)^2},$$

where the elementary inequality $(a+b)^2 \geq b^2/2 - a^2$, $a, b \in \mathbb{R}$, was used in the first step. But if $g \notin \text{ran}(P_\rho)$, then $\sum_{i \in \mathbb{N}} \lambda_i^{-2} g_i^2 = \infty$ and the above lower bound also tends to infinity. This proves (36).

It remains to show (38). Set

$$\Pi_k^{m-1}(\lambda) := \prod_{l=k}^{m-1} (1 - \mu_k \lambda), \quad k = 0, \ldots, m-1, \qquad \Pi_m^{m-1}(\lambda) = 1.$$

Since $1/2 \leq 1 - \mu_m \lambda \leq 1$, we have $\log(1 - \mu_m \lambda) \geq -\log(2)\mu_m \lambda$ for all $m \geq 0$. This gives

$$\Pi_k^{m-1}(\lambda) \geq 2^{-\lambda \sum_{l=k}^{m-1} \mu_l}, \qquad k = 1, \ldots, m-1.$$

But for $0 < t < 1$ we have

$$\sum_{l=k}^{m-1} \mu_l = A \sum_{l=k+1}^{m} l^{-t} \leq \frac{1}{2\Lambda} \int_{k+1/2}^{m+1/2} x^{-t} \, dx$$

$$= \frac{(m+\frac{1}{2})^{1-t} - (k+\frac{1}{2})^{1-t}}{2\Lambda(1-t)} \leq C(m+1)^{1-t} - (k+1)^{1-t}), \quad k = 1, \ldots, m-1,$$

where C depends on t and Λ. So,

$$\Pi_k^{m-1}(\lambda) \geq 2^{-C\lambda((m+1)^{1-t} - (k+1)^{1-t})}, \qquad k = 1, \ldots, m,$$

and, consequently,

$$d_m(\lambda) \geq \frac{1}{2} \sum_{k \in I_m(\lambda)} k \mu_{k-1}, \qquad I_m(\lambda) = \{ k \leq m : C\lambda((m+1)^{1-t} - (k+1)^{1-t}) \leq 1 \}. \tag{39}$$

Obviously, the cardinality of $I_m(\lambda)$ is equal to $|I_m(\lambda)| = m - k_0$, where k_0 is the largest $k \geq 0$ not in $I_m(\lambda)$, i.e.,

$$(k_0 + 1)^{1-t} < (m+1)^{1-t} - (C\lambda)^{-1} \leq (k_0 + 2)^{1-t}.$$

If

$$(2/C)(m+1)^{t-1} \leq \lambda \leq \Lambda,$$

then $(k_0 + 2)^{1-t} \geq (m+1)^{1-t}/2$, which implies that, for sufficiently large $m \geq m_0$, there exists such a k_0 that satisfies $k_0 + 1 \geq c_2(m+1)$ with some constant $c_2 > 0$, where m_0 and c_2 depend on t. Furthermore, by definition of k_0, we have

$$\frac{1}{C\lambda} < (m+1)^{1-t} - (k_0 + 1)^{1-t} \leq (1-t)\frac{m - k_0}{(k_0 + 1)^t}.$$

Substituting this, $k \geq k_0 + 1 \geq c_2(m+1)$, and $\mu_{k-1} = Ak^{-t}$ into (39) we get

$$
\begin{aligned}
d_m(\lambda) &\geq \frac{A}{2}(k_0+1)^{1-t}(m-k_0) \geq \frac{A(k_0+1)}{2(1-t)C\lambda} \\
&\geq \frac{Ac_2}{2(1-t)C}\frac{m+1}{\lambda}, \qquad (2/C)(m+1)^{t-1} \leq \lambda \leq \Lambda, \quad m \geq m_0.
\end{aligned}
$$

This proves (38) if we set

$$
c_0 = \frac{Ac_2}{2(1-t)C}, \qquad c_1 = \frac{2}{C},
$$

and finishes the argument for (36).

5.4 A special case

Now consider the special "learning" problem of recovering an unknown element $u \in V$ from noisy measurements of its coefficients with respect to a CONS $\Psi = \{\psi_i\}_{i \in \mathbb{N}}$ in V by the online method considered in this paper. To do this, we assume that we are given μ-distributed random samples (i_m, y_m), where $i_m \in \mathbb{N}$ and

$$
y_m = (u, \psi_{i_m}) + \varepsilon_m, \qquad m = 0, 1, \ldots \tag{40}
$$

are the noisy samples of the coefficients (u, ψ_i). Starting from $u^{(0)} = 0$, we want to approximate u by the iterates $u^{(m)}$ obtained from the online algorithm

$$
u^{(m+1)} = \alpha_m u^{(m)} + \alpha_m \mu_m (y_m - (u^{(m)}, \psi_{i_m}))\psi_{i_m}, \qquad m = 0, 1, \ldots, \tag{41}
$$

where the coefficients α_m and μ_m are given by (6) with $\Lambda = 1$. This is a special instance of (4) if we set $\Omega = \mathbb{N}$, $Y = \mathbb{R}$ and define $R_i : \mathbb{R} \to V$ and $R_i^* : V \to \mathbb{R}$ by $R_i y = y\psi_i$ and $R_i^* v = (v, \psi_i)$, respectively. To simplify things further, let i_m be i.i.d. samples from $\mathbb{N}$ with respect to a discrete probability measure ρ on $\mathbb{N}$, and let ε_m be i.i.d. random noise with zero mean and finite variance $\sigma^2 < \infty$ that is independent of i_m. The corresponding operator P_ρ is given by

$$
P_\rho v = \sum_{i \in \mathbb{N}} \rho_i (v, \psi_i)\psi_i,
$$

its eigenvalues $\lambda_i = \rho_i$ are given by ρ, and it is trace class (w.l.o.g., we assume $\rho_1 \geq \rho_2 \geq \ldots$). The spaces $V_{P_\rho}^s$, $-\infty < s < \infty$, can now be identified as sets of formal orthogonal series

$$
V_{P_\rho}^s := \left\{ u \sim \sum_{i \in \mathbb{N}} c_i \psi_i : \quad \|u\|_{V_{P_\rho}^s}^2 = \sum_{i \in \mathbb{N}} \rho_i^{-s} c_i^2 \right\}.
$$

Obviously, $V_{P_\rho}^s \subset V = V_{P_\rho}^0$ for $s > 0$. Since functions $f : \mathbb{N} \to \mathbb{R}$ can be identified with formal series with respect to Ψ by

$$
u \sim \sum_{i \in \mathbb{N}} c_i \psi_i \quad \leftrightarrow \quad f_u : f_u(i) = c_i,
$$

we have $\|f_u\|_{L_\rho^2(\mathbb{N}, \mathbb{R})} = \|u\|_{V_{P_\rho}^{-1}}$ and we can silently identify $L_\rho^2(\mathbb{N}, \mathbb{R})$ with $V_{P_\rho}^{-1}$. Under the assumptions made, the underlying minimization problem (7) on V is

$$
\mathbb{E}(|f_v - y|^2) = \|f_v - f_u\|_{L_\rho^2(\mathbb{N}, \mathbb{R})}^2 + \sigma^2 \quad \longmapsto \quad \min,
$$

and, as expected, has u as its unique solution. This example also shows that sometimes it is more important to consider convergence in V than convergence in the sense of $L_\rho^2(\Omega, Y)$.

The simplicity of this example allows a rather comprehensive convergence theory with respect to the scale of $V_{P_\rho}^s$ spaces. We state the following results without detailed proof.

Theorem 2 *Let $-1 \leq \bar{s} \leq 0 \leq s$, and $\bar{s} < s \leq \bar{s}+2$. Then, for the sampling process described above, the online algorithm (41) converges for $u \in V_{P_\rho}^s$ in the $V_{P_\rho}^{\bar{s}}$ norm with the bound*

$$\mathbb{E}(\|e^{(m)}\|_{V_{P_\rho}^{\bar{s}}}^2) \leq C(m+1)^{-\min(\frac{s-\bar{s}}{s+2},\frac{2}{\bar{s}+4})}(A^{\bar{s}-s}\|u\|_{V_{P_\rho}^s}^2 + A^{2+\bar{s}}\sigma^2), \quad m = 1,2,\ldots, \quad (42)$$

if the parameters t and A in (6) satisfy

$$t = t_{s,\bar{s}} := \max((s+1)/(s+2),(\bar{s}+3)/(\bar{s}+4)), \qquad 0 < A \leq 1/2.$$

Setting $\bar{s} = 0$, one concludes from (42) that the convergence estimate for the online algorithm (4), which holds by Theorem 1 for $0 < s \leq 1$ in the general case, is indeed matched. For $\bar{s} = -1$, which corresponds to $L_\rho^2(\mathbb{N},\mathbb{R})$ convergence, the rate is better and in line with known lower bounds.

The estimate (42) for the online algorithm (41) is best possible in the sense that, under the conditions of Theorem 2, the exponent in (42) cannot be increased without additional assumptions on ρ. In particular, there is no further improvement for $s > \bar{s}+2$, i.e., the estimate actually saturates at $s = \bar{s}+2$. This can be seen from the following result.

Theorem 3 *Let $-1 \leq \bar{s} \leq 0 \leq s$, $\bar{s} < s$ and $\sigma > 0$. For the online algorithm (41) we have*

$$\sup_\rho \sup_{u: \|u\|_{V_{P_\rho}^s}=1} (m+1)^{\min((s-\bar{s})/(2+s),2/(\bar{s}+4))}\mathbb{E}(\|e^{(m)}\|_{V_{P_\rho}^{\bar{s}}}^2) \geq c > 0, \quad (43)$$

$m = 1,2,\ldots,$, where c depends on $\bar{s}$, s, σ, and the parameters t and A in (6), but is independent of m.

The proofs of these statements are elementary but rather tedious and will be given elsewhere. Let us just note that the simplicity of this example allows us to reduce the considerations to explicit linear recursions for expectations associated with the decomposition coefficients $c_i^{(m)} := (e^{(m)}, \psi_i)$ of the errors $e^{(m)} = u - u^{(m)}$ with respect to Ψ for each $i \in \mathbb{N}$ separately. This is because

$$\mathbb{E}(\|e^{(m)}\|_{V_{P_\rho}^{\bar{s}}}^2) = \sum_i \rho_i^{-\bar{s}}\mathbb{E}((c_i^{(m)})^2), \qquad \|u\|_{V_{P_\rho}^s}^2 = \|e^{(0)}\|_{V_{P_\rho}^s}^2 = \sum_i \rho_i^{-s}c_i^2 \quad (44)$$

and

$$c_i^{(m+1)} = \bar{\alpha}_m c_i + \alpha_m c_i^{(m)} + \alpha_m \begin{cases} \mu_m(y_{i_m} - (u^{(m)}, \psi_{i_m})), & i_m = i \\ 0, & i_m \neq i \end{cases}$$

$$= \bar{\alpha}_m c_i + \alpha_m(c_i^{(m)} - \delta_{i,i_m}\mu_m(c_i^{(m)} + \varepsilon_m))$$

for $m = 0,1,\ldots,$ where $\delta_{i,i_m} = 1$ with probability ρ_i, and $\delta_{i,i_m} = 0$ with probability $1 - \rho_i$. So if we denote $\varepsilon_{m,i} := \mathbb{E}((c_i^{(m)})^2)$ and $\bar{\varepsilon}_{m,i} := \mathbb{E}(c_i^{(m)})$ and use the independence assumption, we get a system of linear recursions

$$\varepsilon_{m+1,i} = \alpha_m^2(1 - \rho_i\mu_m(2 - \mu_m))\varepsilon_{m,i} + 2\alpha_m\bar{\alpha}_m(1 - \rho_i\mu_m)c_i\bar{\varepsilon}_{m,i} + \bar{\alpha}_m^2 c_i^2 + \rho_i\alpha_m^2\mu_m^2\sigma^2,$$
$$\bar{\varepsilon}_{m+1,i} = \alpha_m(1 - \rho_i\mu_m)\bar{\varepsilon}_{m,i} + \bar{\alpha}_m c_i,$$

$m = 0,1,\ldots,$ with starting values $\varepsilon_{0,i} = c_i^2$ and $\bar{\varepsilon}_{0,i} = c_i$. In principle, this system can be solved explicitly. For example, we have

$$\bar{\varepsilon}_{m,i} = \frac{1}{m+1}c_i S_m, \qquad S_m := \sum_{k=0}^m \Pi_k^{m-1},$$

where the notation

$$\Pi_k^{m-1} = (1 - \frac{a}{m^t})\cdot\ldots\cdot(1 - \frac{a}{(k+1)^t}), \quad 0 \leq k \leq m-1, \quad \Pi_m^{m-1} = 1, \quad (45)$$

is used with $a = A\rho_i$. Similarly, we get

$$\varepsilon_{m,i} \leq \frac{1}{(m+1)^2} \left(2c_i^2 \sum_{k=0}^{m} \Pi_k^{m-1} S_k + \rho_i \sigma^2 \sum_{k=0}^{m-1} (k+1)^2 \mu_k^2 \Pi_{k+1}^{m-1} \right).$$

A matching lower bound for $\varepsilon_{m,i}$ can be obtained by using a slightly different value of a in the definition of the products Π_k^{m-1}. The remainder of the argument for Theorem 2 first requires the substitution of tight upper bounds for Π_k^{m-1} and S_k in dependence on t and a into the bounds for $\varepsilon_{m,i}$. Next, after substituting the estimates for $\varepsilon_{m,i}$ into (44), the resulting series has to be estimated separately for the index sets $I_1 := \{i : A\rho_i \leq (m+1)^{t-1}\}$ and $I_2 := \mathbb{N}\backslash I_1$ followed by choosing the indicated optimal value of $t = t_{s,\bar{s}}$. This leads to the bound (42) in Theorem 2. For the proof of Theorem 3, lower bounds for Π_k^{m-1}, S_k and consequently for ε_m are needed, combined with choosing suitable discrete probability distributions ρ. Regarding lower bounds for Π_k^{m-1}, see the considerations in Subsection 5.3.

Acknowledgement

References

1. Caponnetto, A., De Vito, A.: Optimal rates for the regularized least-squares algorithm. Found. Comput. Math. **7:3**, 331–368 (2007)
2. Carmeli, C., De Vito, E., Toigo, A.: Vector valued reproducing kernel Hilbert spaces of integrable functions and Mercer theorem. Anal. Appl. **4:4**, 377–408 (2006)
3. Carmeli, C., De Vito, E., Toigo, A., Umanita, V.: Vector valued reproducing kernel Hilbert spaces and universality. Anal. Appl. **8:1**, 19–61 (2010)
4. Cucker, F., Smale, S.: On the mathematical foundations of learning. Bull. Amer. Math. Soc. **39:1**, 1–49 (2002)
5. Cucker, F., Zhou, D.-X.: Learning Theory: An Approximation Theory Viewpoint. Cambridge Univ. Press (2007)
6. Dieuleveut, A., Flammarion, N., Bach, F.: Harder, better, faster, stronger convergence rates for least-squares regression. J. Machine Learning Research **18**, 1–51 (2017)
7. Dieuleveut, A., Bach, F.: Nonparametric stochastic approximation with large step-sizes, Ann. Statistics **44:4**, 1363–1399 (2016)
8. Griebel, M., Oswald, P., Stochastic subspace correction in Hilbert space, Constr. Approx. **48:3**, 501–521 (2018)
9. Guo, X., Lin, J., Zhou, D.-X.: Rates of convergence of randomized Kaczmarz algorithms in Hilbert spaces. Appl. Comput. Harm. Anal. **61:3**, 288–318 (2022)
10. Guo, Z.-C., Shi, L.: Fast and strong convergence of online learning algorithms. Adv. Comput. Math. **45**, 2745–2770 (2019)
11. Micchelli, C., Pontil, M.: On learning vector-valued functions. Neural. Comput. **17:1**, 177–204 (2005)
12. Schölkopf, B., Smola, A.: Learning with Kernels. MIT Press, Cambridge MA, (2002)
13. Steinwart, I., Christmann, A.: Support Vector Machines. Springer, New York (2008)
14. Tarrés, P., Yao, Y.: Online learning as stochastic approximation of regularization paths: Optimality and almost-sure convergence. IEEE Transactions on Information Theory **60:9**, 5716–5735 (2014)
15. Vapnik, V.: Statistical Learning Theory. Wiley, New York (1998)
16. Ying, Y., Pontil, M.: Online gradient descent learning algorithms. Found. Comput. Math. **8:5**, 561–596 (2008)