
Efficient online learning with Kernels for adversarial large scale problems

Rémi Jézéquel Pierre Gaillard Alessandro Rudi

INRIA - Département d'Informatique de l'École Normale Supérieure
PSL Research University, Paris, France
{remi.jezequel,pierre.gaillard,alessandro.rudi}@inria.fr

Abstract

We are interested in a framework of online learning with kernels for low-dimensional, but large-scale and potentially adversarial datasets. We study the computational and theoretical performance of online variations of kernel Ridge regression. Despite its simplicity, the algorithm we study is the first to achieve the optimal regret for a wide range of kernels with a per-round complexity of order n^α with $\alpha < 2$.

The algorithm we consider is based on approximating the kernel with the linear span of basis functions. Our contributions are twofold: 1) For the Gaussian kernel, we propose to build the basis beforehand (independently of the data) through Taylor expansion. For d -dimensional inputs, we provide a (close to) optimal regret of order $O((\log n)^{d+1})$ with per-round time complexity and space complexity $O((\log n)^{2d})$. This makes the algorithm a suitable choice as soon as $n \gg e^d$ which is likely to happen in a scenario with small dimensional and large-scale dataset; 2) For general kernels with low effective dimension, the basis functions are updated sequentially, adapting to the data, by sampling Nyström points. In this case, our algorithm improves the computational trade-off known for online kernel regression.

1 Introduction

Nowadays the volume and the velocity of data flows are deeply increasing. Consequently, many applications need to switch from batch to online procedures that can treat and adapt to data on the fly. Furthermore to take advantage of very large datasets, non-parametric methods are gaining increasing momentum in practice. Yet the latter often suffer from slow rates of convergence and bad computational complexities. At the same time, data is getting more complicated and simple stochastic assumptions, such as i.i.d. data, are often not satisfied. In this paper, we try to combine these different aspects due to large scale and arbitrary data. We build a non-parametric online procedure based on kernels, which is efficient for large data sets and achieves close to optimal theoretical guarantees.

Online learning is a subfield of machine learning where a learner sequentially interacts with an environment and tries to learn and adapt on the fly to the observed data as one goes along. We consider the following sequential setting. At each iteration $t \geq 1$, the learner receives some input $x_t \in \mathcal{X}$; makes a prediction $\hat{y}_t \in \mathbb{R}$ and the environment reveals the output $y_t \in \mathbb{R}$. The inputs x_t and the outputs y_t are sequentially chosen by the environment and can be arbitrary. Learner's goal is to minimize his cumulative regret

$$R_n(f) := \sum_{t=1}^n (y_t - \hat{y}_t)^2 - \sum_{t=1}^n (y_t - f(x_t))^2 \quad (1)$$

uniformly over all functions f in a space of functions \mathcal{H} . We will consider Reproducing Kernel Hilbert Space (RKHS) \mathcal{H} , [see next section or Aro50, for more details]. It is worth noting here that

all the properties of a RKHS are controlled by the associated *kernel function* $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, usually known in closed form, and that many function spaces of interest are (or are contained in) RKHS, e.g. when $\mathcal{X} \subseteq \mathbb{R}^d$: polynomials of arbitrary degree, band-limited functions, analytic functions with given decay at infinity, Sobolev spaces and many others [BTA11].

Previous work Kernel regression in a statistical setting has been widely studied by the statistics community. Our setting of online kernel regression with adversarial data is more recent. Most of the existing work focuses on the linear setting (i.e., linear kernel). First work on online linear regression dates back to [Fos91]. [BKM⁺15] provided the minimax rates (together with an algorithm) and we refer the reader to references therein for a recent overview of the literature in the linear case. We only recall relevant work for this paper. [AW01, Vov01] designed the nonlinear Ridge forecaster (denoted AWV). In linear regression (linear kernel), it achieves the optimal regret of order $O(d \log n)$ uniformly over all ℓ_2 -bounded vectors. The latter can be extended to kernels (see Definition (3)) which we refer to as Kernel-AWV. With regularization parameter $\lambda > 0$, it obtains a regret upper-bounded for all $f \in \mathcal{H}$ as

$$R_n(f) \lesssim \lambda \|f\|^2 + B^2 d_{\text{eff}}(\lambda), \quad \text{where} \quad d_{\text{eff}}(\lambda) := \text{Tr}(K_{nn}(K_{nn} + \lambda I_n)^{-1}) \quad (2)$$

is the effective dimension, where $K_{nn} := (k(x_i, x_j))_{1 \leq i, j \leq n} \in \mathbb{R}^{n \times n}$ denotes the *kernel matrix* at time n . The above upper-bound on the regret is essentially optimal (see remark 2.1). Yet the per round complexity and the space complexity of Kernel-AWV are $O(n^2)$. In this paper, we aim at reducing this complexity while keeping optimal regret guarantees.

Though the literature on online contextual learning is vast, little considers non-parametric function classes. Related work includes [Vov06] that considers the Exponentially Weighted Average forecaster or [HM07] which considers bounded Lipschitz function set and Lipschitz loss functions, while here we focus on the square loss. Minimax rates for general function sets \mathcal{H} are provided by [RST13]. RKHS spaces were first considered in [Vov05] though they only obtain $O(\sqrt{n})$ rates which are suboptimal for our problem. More recently, a regret bound of the form (2) was proved by [ZK10] for a clipped version of kernel Ridge regression and by [CLV17b] for a clipped version of Kernel Online Newton Step (KONS) for general exp-concave loss functions.

The computational complexity ($O(n^2)$ per round) of these algorithms is however prohibitive for large datasets. [CLV17b] and [CLV17a] provide approximations of KONS to get manageable complexities. However, these come with deteriorated regret guarantees. [CLV17b] improves the time and space complexities by a factor $\gamma \in (0, 1)$ enlarging the regret upper-bound by $1/\gamma$. [CLV17a] designs an efficient approximation of KONS based on Nyström approximation [SSL00, WS01] and restarts with per-round complexities $O(m^2)$ where m is the number of Nyström points. Yet their regret bound suffers an additional multiplicative factor m with respect to (2) because of the restarts. Furthermore, contrary to our results, the regret bounds of [CLV17b] and [CLV17a] are not with respect to all functions in \mathcal{H} but only with functions $f \in \mathcal{H}$ such that $f(x_t) \leq C$ for all $t \geq 1$ where C is a parameter of their algorithm. Since C comes with a multiplicative factor of their bounds, their results are sensitive to outliers that may lead to large C . Other relevant approximation schemes of Online Kernel Learning have been done by [LHW⁺16] and [ZL19]. The authors consider online gradient descent algorithms which they approximate using different approximation schemes (as Nyström and random features). However since they use general Lipschitz loss functions and consider ℓ_1 -bounded dual norm of functions f , their regret bounds of order $O(\sqrt{n})$ are hardly comparable to ours and seem suboptimal in n in our restrictive setting with square loss and kernels with small effective dimension (such as Gaussian kernel).

Contributions and outline of the paper The main contribution of the paper is to analyse a variant of Kernel-AWV that we call PKAWV (see Definition (4)). Despite its simplicity, it is to our knowledge the first algorithm for kernel online regression that recovers the optimal regret (see bound (2)) with an improved space and time complexity of order $\ll n^2$ per round. Table 1 summarizes the regret rates and complexities obtained by our algorithm and the ones of [CLV17b, CLV17a].

Our procedure consists simply in applying Kernel-AWV while, at time $t \geq 1$, approximating the RKHS \mathcal{H} with a linear subspace \mathcal{H}_t of smaller dimension. In Theorem 3, PKAWV suffers an additional approximation term with respect to the optimal bound of Kernel-AWV which can be made small enough by properly choosing \mathcal{H}_t . To achieve the optimal regret with low computational

Kernel	Algorithm	Regret	Per-round complexity
Gaussian $d_{\text{eff}}(\lambda) \leq (\log \frac{n}{\lambda})^d$	PKAWV	$(\log n)^{d+1}$	$(\log n)^{2d}$
	Sketched-KONS [CLV17b] ($c > 0$)	$c(\log n)^{d+1}$	$(n/c)^2$
	Pros-N-KONS [CLV17a]	$(\log n)^{2d+1}$	$(\log n)^{2d}$
General $d_{\text{eff}}(\lambda) \leq (\frac{n}{\lambda})^\gamma$ $\gamma < \sqrt{2} - 1$	PKAWV	$n^{\frac{\gamma}{\gamma+1}} \log n$	$n^{\frac{4\gamma}{1-\gamma^2}}$
	Sketched-KONS [CLV17b] ($c > 0$)	$cn^{\frac{\gamma}{\gamma+1}} \log n$	$(n/c)^2$
	Pros-N-KONS [CLV17a]	$n^{\frac{4\gamma}{(1+\gamma)^2}} \log n$	$n^{\frac{4\gamma(1-\gamma)}{(1+\gamma)^2}}$

Table 1: Order in n of the best possible regret rates achievable by the algorithms and corresponding per-round time-complexity. Up to $\log n$, the rates obtained by PKAWV are optimal.

complexity, $\tilde{\mathcal{H}}_t$ needs to approximate \mathcal{H} well and to be low dimensional with an easy-to-compute projection. We provide two relevant constructions for $\tilde{\mathcal{H}}_t$.

In section 3.1, we focus on the Gaussian kernel that we approximate by a finite set of basis functions. The functions are deterministic and chosen beforehand by the learner independently of the data. The number of functions included in the basis is a parameter to be optimized and fixes an approximation-computational trade-off. Theorem 4 shows that PKAWV satisfies (up to \log) the optimal regret bounds (2) while enjoying a per-round space and time complexity of $O(\log^{2d}(\frac{n}{\lambda}))$. For the Gaussian kernel, this corresponds to $O(d_{\text{eff}}(\lambda)^2)$ which is known to be optimal even in the statistical setting with i.i.d. data.

In section 3.2, we consider data adaptive approximation spaces $\tilde{\mathcal{H}}_t$ based on Nyström approximation. At time $t \geq 1$, we approximate any kernel \mathcal{H} by sampling a subset of the input vectors $\{x_1, \dots, x_t\}$. If the kernel satisfies the capacity condition $d_{\text{eff}}(\lambda) \leq (n/\lambda)^\gamma$ for $\gamma \in (0, 1)$, the optimal regret is then of order $d_{\text{eff}}(\lambda) = O(n^{\gamma/(1+\gamma)})$ for well-tuned parameter λ . Our method then recovers the optimal regret with a computational complexity of $O(d_{\text{eff}}(\lambda)^{4/(1-\gamma)})$. The latter is $o(n^2)$ (for well-tuned λ) as soon as $\gamma < \sqrt{2} - 1$. Furthermore, if the sequence of input vectors x_t is given beforehand to the algorithm, the per-round complexity needed to reach the optimal regret is improved to $O(d_{\text{eff}}(\lambda)^4)$ and our algorithm can achieve it for all $\gamma \in (0, 1)$.

Finally, we perform in Section 4 several experiments based on real and simulated data to compare the performance (in regret and in time) of our methods with competitors.

Notations We recall here basic notations that we will use throughout the paper. Given a vector $v \in \mathbb{R}^d$, we write $v = (v^{(1)}, \dots, v^{(d)})$. We denote by $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ the set of non-negative integers and for $p \in \mathbb{N}_0^d$, $|p| = p^{(1)} + \dots + p^{(d)}$. By a slight abuse of notation, we denote by $\|\cdot\|$ both the Euclidean norm and the norm for the Hilbert space \mathcal{H} . Write $v^\top w$, the dot product between $v, w \in \mathbb{R}^D$. The conjugate transpose for linear operator Z on \mathcal{H} will be denoted Z^* . The notation \lesssim will refer to approximate inequalities up to logarithmic multiplicative factors. Finally, we will denote $a \vee b = \max(a, b)$ and $a \wedge b = \min(a, b)$, for $a, b \in \mathbb{R}$.

2 Background

Kernels. Let $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a positive definite kernel [Aro50] that we assume to be bounded (i.e., $\sup_{x \in \mathcal{X}} k(x, x) \leq \kappa^2$ for some $\kappa > 0$). The function k is characterized by the existence of a *feature map* $\phi : \mathcal{X} \rightarrow \mathbb{R}^D$, with $D \in \mathbb{N} \cup \{\infty\}$ ¹ such that $k(x, x') = \phi(x)^\top \phi(x')$. Moreover the *reproducing kernel Hilbert space* (RKHS) associated to k is characterized by $\mathcal{H} = \{f \mid f(x) = w^\top \phi(x), w \in \mathbb{R}^D, x \in \mathcal{X}\}$, with inner product $\langle f, g \rangle_{\mathcal{H}} := v^\top w$, for $f, g \in \mathcal{H}$ defined by $f(x) = v^\top \phi(x)$, $g(x) = w^\top \phi(x)$ and $v, w \in \mathbb{R}^D$. For more details and different characterizations of k, \mathcal{H} , see [Aro50, BTA11]. It's worth noting that the knowledge of ϕ is not necessary when working with functions of the form $f = \sum_{i=1}^p \alpha_i \phi(x_i)$, with $\alpha_i \in \mathbb{R}$, $x_i \in \mathcal{X}$ and finite $p \in \mathbb{N}$, indeed $f(x) = \sum_{i=1}^p \alpha_i \phi(x_i)^\top \phi(x) = \sum_{i=1}^p \alpha_i k(x_i, x)$, and moreover $\|f\|_{\mathcal{H}}^2 = \alpha^\top K_{pp} \alpha$, with K_{pp} the kernel matrix associated to the set of points x_1, \dots, x_p .

¹when $D = \infty$ we consider \mathbb{R}^D as the space of squared summable sequences.

Kernel-AWV. The Azoury-Warmuth-Vovk forecaster (denoted *AWV*) on the space of linear functions on $\mathcal{X} = \mathbb{R}^d$ has been introduced and analyzed in [AW01, Vov01]. We consider here a straightforward generalization to kernels (denoted *Kernel-AWV*) of the nonlinear Ridge forecaster (*AWV*) introduced by [AW01, Vov01] on the space of linear functions on $\mathcal{X} = \mathbb{R}^d$. At iteration $t \geq 1$, Kernel-AWV predicts $\hat{y}_t = \hat{f}_t(x_t)$, where

$$\hat{f}_t \in \operatorname{argmin}_{f \in \mathcal{H}} \left\{ \sum_{s=1}^{t-1} (y_s - f(x_s))^2 + \lambda \|f\|^2 + f(x_t)^2 \right\}. \quad (3)$$

A variant of this algorithm, more used in the context of data independently sampled from distribution, is known as *kernel Ridge regression*. It corresponds to solving the problem above, without the last penalization term $f(x_t)^2$.

Optimal regret for Kernel-AWV. In the next proposition, we state a preliminary result which proves that Kernel-AWV achieves a regret depending on the eigenvalues of the kernel matrix.

Proposition 1. *Let $\lambda, B > 0$. For any RKHS \mathcal{H} , for all $n \geq 1$, for all inputs $x_1, \dots, x_n \in \mathcal{X}$ and all $y_1, \dots, y_n \in [-B, B]$, the regret of Kernel-AWV is upper-bounded for all $f \in \mathcal{H}$ as*

$$R_n(f) \leq \lambda \|f\|^2 + B^2 \sum_{k=1}^n \log \left(1 + \frac{\lambda_k(K_{nn})}{\lambda} \right),$$

where $\lambda_k(K_{nn})$ denotes the k -th largest eigenvalue of K_{nn} .

The proof is a direct consequence of the known regret bound of *AWV* in the finite dimensional linear regression setting—see Theorem 11.8 of [CBL06] or Theorem 2 of [GGHS18]. For completeness, we reproduce the analysis for infinite dimensional space (RKHS) in Appendix C.1. In online linear regression in dimension d , the above result implies the optimal rate of convergence $dB^2 \log(n) + O(1)$ (see [GGHS18] and [Vov01]). As shown by the following proposition, Proposition 1 yields optimal regret (up to log) of the form (2) for online kernel regression.

Proposition 2. *For all $n \geq 1$, $\lambda > 0$ and all input sequences $x_1, \dots, x_n \in \mathcal{X}$,*

$$\sum_{k=1}^n \log \left(1 + \frac{\lambda_k(K_n)}{\lambda} \right) \leq \log \left(e + \frac{en\kappa^2}{\lambda} \right) d_{\text{eff}}(\lambda).$$

Combined with Proposition 1, this entails that Kernel-AWV satisfies (up to the logarithmic factor) the optimal regret bound (2). As discussed in the introduction, such an upper-bound on the regret is not new and was already proved by [ZK10] or by [CLV17b] for other algorithms. An advantage of Kernel-AWV is that it does not require any clipping and thus the beforehand knowledge of $B > 0$ to obtained Proposition 1. Furthermore, we slightly improve the constants in the above proposition.

Remark 2.1 (Optimal regret under the capacity condition). Assuming the capacity condition ($d_{\text{eff}}(\lambda) \leq (n/\lambda)^\gamma$ for $0 \leq \gamma \leq 1$), the rate of the regret bound (2) can be made explicit. As we show now, this matches existing minimax lower rates in the stochastic setting. Under the capacity condition, optimizing $\lambda \simeq n^{\gamma/(1+\gamma)}$ to minimize the r.h.s. of (2), the regret bound is then of order $R_n(f) \leq O(n^{\gamma/(1+\gamma)})$ (up to logs). If the data $(x_1, y_1), \dots, (x_n, y_n)$ is i.i.d. according to some distribution ρ over $\mathcal{X} \times \mathbb{R}$, we can apply a standard online to batch conversion (see [CBCG04]). The estimator $\bar{f}_n = \frac{1}{n} \sum_{t=1}^n f_t$ satisfies for any $f \in \mathcal{H}$ the upper-bound on its excess risk

$$\mathcal{E}(\bar{f}_n) - \mathcal{E}(f) \leq \mathbb{E} \left[\frac{R_n(f)}{n} \right] \leq O(n^{-\frac{1}{1+\gamma}}),$$

where $\mathcal{E}(f) := \mathbb{E}_{(X,Y) \sim \rho} [(f(X) - Y)^2]$. This corresponds to the known minimax lower rate in this stochastic setting as shown by Theorem 2 (applied with $c = 1$ and $b = 1/\gamma$) of [CDV07].

It is worth pointing out that in the worst case $d_{\text{eff}}(\lambda) \leq \kappa^2 n/\lambda$ for any bounded kernel. In particular, optimizing the bound yields $\lambda = O(\sqrt{n \log n})$ and a regret bound of order $O(\sqrt{n \log n})$. In the special case of the Gaussian kernel (which we consider in Section 3.1), the latter can be improved to $d_{\text{eff}}(\lambda) \lesssim (\log(n/\lambda))^d$ (see [ABRW18]) which entails $R_n(f) \leq O((\log n)^{d+1})$ for well tuned value of λ .

3 Online Kernel Regression with projections

In the previous section we have seen that Kernel-AWV achieves optimal regret. Yet, it has computational requirements that are $O(n^3)$ in time and $O(n^2)$ in space, for n steps of the algorithm, making it unfeasible in the context of large scale datasets, i.e. $n \gg 10^5$. In this paper, we consider and analyze a simple variation of Kernel-AWV denoted PKAWV. At time $t \geq 1$, for a regularization parameter $\lambda > 0$ and a linear subspace $\tilde{\mathcal{H}}_t$ of \mathcal{H} the algorithm predicts $\hat{y}_t = \hat{f}_t(x_t)$, where

$$\hat{f}_t = \operatorname{argmin}_{f \in \tilde{\mathcal{H}}_t} \left\{ \sum_{s=1}^{t-1} (y_s - f(x_s))^2 + \lambda \|f\|^2 + f(x_t)^2 \right\}. \quad (4)$$

In the next subsections, we explicit relevant approximations $\tilde{\mathcal{H}}_t$ (typically the span of a small number of basis functions) of \mathcal{H} that trade-off good approximation with a low computational cost. Appendix H details how (4) can be efficiently implemented in these cases.

The result below bounds the regret of the PKAWV for any function $f \in \mathcal{H}$ and holds for any bounded kernel and any explicit subspace $\tilde{\mathcal{H}}$ associated with projection P . The cost of the approximation of \mathcal{H} by $\tilde{\mathcal{H}}$ is measured by the important quantity $\mu := \|(I - P)C_n^{1/2}\|^2$, where C_n is the covariance operator.

Theorem 3. *Let $\tilde{\mathcal{H}}$ be a linear subspace of \mathcal{H} and P the Euclidean projection onto $\tilde{\mathcal{H}}$. When PKAWV is run with $\lambda > 0$ and fixed subspaces $\tilde{\mathcal{H}}_t = \tilde{\mathcal{H}}$, then for all $f \in \mathcal{H}$*

$$R_n(f) \leq \lambda \|f\|^2 + B^2 \sum_{j=1}^n \log \left(1 + \frac{\lambda_j(K_{nn})}{\lambda} \right) + (\mu + \lambda) \frac{n\mu B^2}{\lambda^2}, \quad (5)$$

for any sequence $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times [-B, B]$ where $\mu := \|(I - P)C_n^{1/2}\|^2$ and $C_n := \sum_{t=1}^n \phi(x_t) \otimes \phi(x_t)$.

The proof of Thm. 3 is deferred to Appendix D.1 and is the consequence of a more general Thm. 9.

3.1 Learning with Taylor expansions and Gaussian kernel for very large data set

In this section we focus on non-parametric regression with the widely used *Gaussian kernel* defined by $k(x, x') = \exp(-\|x - x'\|^2/(2\sigma^2))$ for $x, x' \in \mathcal{X}$ and $\sigma > 0$ and the associated RKHS \mathcal{H} .

Using the results of the previous section with a fixed linear subspace $\tilde{\mathcal{H}}$ which is the span of a basis of $O(\text{polylog}(n/\lambda))$ functions, we prove that PKAWV achieves optimal regret. This leads to a computational complexity that is only $O(n \text{ polylog}(n/\lambda))$ for optimal regret. We need a basis that (1) approximates very well the Gaussian kernel and at the same time (2) whose projection is easy to compute. We consider the following basis of functions, for $k \in \mathbb{N}_0^d$,

$$g_k(x) = \prod_{i=1}^d \psi_{k_i}(x^{(i)}), \quad \text{where} \quad \psi_t(x) = \frac{x^t}{\sigma^t \sqrt{t!}} e^{-\frac{x^2}{2\sigma^2}}. \quad (6)$$

For one dimensional data, this corresponds to Taylor expansion of the Gaussian kernel. Our theorem below states that PKAWV (see (4)) using for all iterations $t \geq 1$

$$\tilde{\mathcal{H}}_t = \text{Span}(G_M) \quad \text{with} \quad G_M = \{g_k \mid |k| \leq M, k \in \mathbb{N}_0^d\}$$

where $|k| := k_1 + \dots + k_d$, for $k \in \mathbb{N}_0^d$, gets optimal regret while enjoying low complexity. The size of the basis M controls the trade-off between approximating well the Gaussian kernel (to incur low regret) and large computational cost. Theorem 4 optimizes M so that the approximation term of Theorem 3 (due to kernel approximation) is of the same order than the optimal regret.

Theorem 4. *Let $\lambda > 0, n \in \mathbb{N}$ and let $R, B > 0$. Assume that $\|x_t\| \leq R$ and $|y_t| \leq B$. When $M = \lceil \frac{8R^2}{\sigma^2} \vee 2 \log \frac{n}{\lambda \wedge 1} \rceil$, then running PKAWV using G_M as set of functions achieves a regret bounded by*

$$\forall f \in \mathcal{H}, R_n(f) \leq \lambda \|f\|^2 + \frac{3B^2}{2} \sum_{j=1}^n \log \left(1 + \frac{\lambda_j(K_{nn})}{\lambda} \right).$$

Moreover, its per iteration computational cost is $O((3 + \frac{1}{d} \log \frac{n}{\lambda \wedge 1})^{2d})$ in space and time.

Therefore PKAWV achieves a regret-bound only deteriorated by a multiplicative factor of $3/2$ with respect to the bound obtained by Kernel-AWV (see Prop. 1). From Prop. 2 this also yields (up to log) the optimal bound (2).

In particular, it is known [ABRW18] for the Gaussian kernel that

$$d_{\text{eff}}(\lambda) \leq 3 \left(6 + \frac{41}{d} \frac{R^2}{2\sigma^2} + \frac{3}{d} \log \frac{n}{\lambda} \right)^d = O \left(\left(\log \frac{n}{\lambda} \right)^d \right).$$

The upper-bound is matching even in the i.i.d. setting for nontrivial distributions. In this case, we have $|G_M| \lesssim d_{\text{eff}}(\lambda)$. The per-round space and time complexities are thus $O(d_{\text{eff}}(\lambda)^2)$. Though our method is quite simple (since it uses fixed explicit embedding) it is able to recover results -in terms of computational time and bounds in the adversarial setting- that are similar to results obtained in the more restrictive i.i.d. setting obtained via much more sophisticated methods, like learning with (1) Nyström with importance sampling via leverage scores [RCR15], (2) reweighted random features [Bac17, RR17], (3) volume sampling [DWH18]. By choosing $\lambda = (B/\|f\|)^2$, to minimize the r.h.s. of the regret bound of the theorem, we get

$$R_n(f) \lesssim \left(\log \frac{n\|f\|_{\mathcal{H}}^2}{B^2} \right)^{d+1} B^2 + B^2. \quad (7)$$

Note that the optimal λ does not depend on n and can be optimized in practice through standard online calibration methods. For instance, one can run in parallel subroutines of the algorithm, each using a different value of λ in the finite grid $\Lambda := \{n2^k, k = -n^{\frac{1}{d+1}}, \dots, 0\}$. The subroutines can then be sequentially combined with an expert advice algorithm such as the Exponentially Weighted Average forecaster [CBL06] at an additional negligible cost of order $O(B^2 \log |\Lambda|)$ in the regret (using the fact that the squared loss is exp-concave on $[0, B]$). Similarly, though we use a fixed number of features M in the experiments, the latter could be increased slowly over time thanks to online calibration techniques.

3.2 Nyström projection

The previous two subsections considered a deterministic function basis (independent of the data) to approximate specific RKHS. Here, we analyse Nyström projections [RCR15] that are data dependent and work for any RKHS. It consists in sequentially updating a dictionary $\mathcal{I}_t \subset \{x_1, \dots, x_t\}$ and using

$$\tilde{\mathcal{H}}_t = \text{Span} \left\{ \phi(x), x \in \mathcal{I}_t \right\}. \quad (8)$$

If the points included into \mathcal{I}_t are well-chosen, the latter may approximate well the solution of (3) which belongs to the linear span of $\{\phi(x_1), \dots, \phi(x_t)\}$. The inputs x_t might be included in the dictionary independently and uniformly at random. Here, we build the dictionary by following the KORS algorithm of [CLV17a] which is based on approximate leverage scores. At time $t \geq 1$, it evaluates the importance of including x_t to obtain an accurate projection P_t by computing its leverage score. Then, it decides to add it or not, by drawing a Bernoulli random variable. The points are never dropped from the dictionary so that $\mathcal{I}_1 \subset \mathcal{I}_2 \subset \dots \subset \mathcal{I}_n$. With their notations, choosing $\varepsilon = 1/2$ and remarking that $\|\Phi_t^T(I - P_t)\Phi_t\| = \|(I - P_t)C_t^{1/2}\|^2$, their Proposition 1 can be rewritten as follows.

Proposition 5. [CLV17a, Prop. 1] *Let $\delta > 0$, $n \geq 1$, $\mu > 0$. Then, the sequence of dictionaries $\mathcal{I}_1 \subset \mathcal{I}_2 \subset \dots \subset \mathcal{I}_n$ learned by KORS with parameters μ and $\beta = 12 \log(n/\delta)$ satisfies w.p. $1 - \delta$,*

$$\forall t \geq 1, \quad \|(I - P_t)C_t^{1/2}\|^2 \leq \mu \quad \text{and} \quad |\mathcal{I}_t| \leq 9d_{\text{eff}}(\mu) \log(2n/\delta)^2.$$

Furthermore, the algorithm runs in $O(d_{\text{eff}}(\mu)^2 \log^4(n))$ space and $O(d_{\text{eff}}(\mu)^2)$ time per iteration.

Using this approximation result together with Thm. 9 (which is a more general version of Thm. 3), we can bound the regret of PKAWV with KORS. The proof is postponed to Appendix E.1.

Theorem 6. *Let $n \geq 1$, $\delta > 0$ and $\lambda \geq \mu > 0$. Assume that the dictionaries $(\mathcal{I}_t)_{t \geq 1}$ are built according to Proposition 5. Then, probability at least $1 - \delta$, PKAWV with the subspaces $\tilde{\mathcal{H}}_t$ defined in (8) satisfies the regret upper-bound*

$$R_n \leq \lambda \|f\|^2 + B^2 d_{\text{eff}}(\lambda) \log(e + en\kappa^2/\lambda) + 2B^2(|\mathcal{I}_n| + 1) \frac{n\mu}{\lambda},$$

and the algorithm runs in $O(d_{\text{eff}}(\mu)^2)$ space $O(d_{\text{eff}}(\mu)^2)$ time per iteration.




Figure 1: Comparison of the theoretical regret rate $\log R_n / \log n$ according to the size of the dictionary $\log m / \log n$ considered by PKAWV, Sketched-KONS and Pros-N-KONS for optimized parameters when $d_{\text{eff}}(\lambda) \leq (n/\lambda)^\gamma$ with $\gamma = 0.2, \sqrt{2}-1, 0.6$ (from left to right). The value $\gamma/(1+\gamma)$ corresponds to the optimal rate.

The last term of the regret upper-bound above corresponds to the approximation cost of using the approximation (8) in PKAWV. This cost is controlled by the parameter $\mu > 0$ which trades-off between having a small approximation error (small μ) and a small dictionary of size $|\mathcal{I}_n| \approx d_{\text{eff}}(\mu)$ (large μ) and thus a small computational complexity. For the Gaussian Kernel, using that $d_{\text{eff}}(\lambda) \leq O(\log(n/\lambda)^d)$, the above theorem yields for the choice $\lambda = 1$ and $\mu = n^{-2}$ a regret bound of order $R_n \leq O((\log n)^{d+1})$ with a per-round time and space complexity of order $O(|\mathcal{I}_n|^2) = O((\log n)^{2d+4})$. We recover a similar result to the one obtained in Section 3.1.

Explicit rates under the capacity condition Assuming the capacity condition $d_{\text{eff}}(\lambda') \leq (n/\lambda')^\gamma$ for $0 \leq \gamma \leq 1$ and $\lambda' > 0$, which is a classical assumption made on kernels [RCR15], the following corollary provides explicit rates for the regret according to the size of the dictionary $m \approx |\mathcal{I}_n|$.

Corollary 7. *Let $n \geq 1$ and $m \geq 1$. Assume that $d_{\text{eff}}(\lambda') \leq (n/\lambda')^\gamma$ for all $\lambda' > 0$. Then, under the assumptions of Theorem 6, PKAWV with $\mu = nm^{-1/\gamma}$ has a dictionary of size $|\mathcal{I}_n| \lesssim m$ and a regret upper-bounded with high-probability as*

$$R_n \lesssim \begin{cases} n^{\frac{\gamma}{1+\gamma}} & \text{if } m \geq n^{\frac{2\gamma}{1-\gamma^2}} \quad \text{for } \lambda = n^{\frac{\gamma}{1+\gamma}} \\ nm^{\frac{1}{2}-\frac{1}{2\gamma}} & \text{otherwise} \quad \text{for } \lambda = nm^{\frac{1}{2}-\frac{1}{2\gamma}} \end{cases}.$$

The per-round space and time complexity of the algorithm is $O(m^2)$ per iteration.

The rate of order $n^{\frac{\gamma}{1+\gamma}}$ is optimal in this case (it corresponds to optimizing (2) in λ). If the dictionary is large enough $m \geq n^{2\gamma/(1-\gamma^2)}$, the approximation term is negligible and the algorithm recovers the optimal rate. This is possible for a small dictionary $m = o(n)$ whenever $2\gamma/(1-\gamma^2) < 1$, which corresponds to $\gamma < \sqrt{2}-1$. The rates obtained in Corollary 7 can be compared to the one obtained by Sketched-KONS of [CLV17b] and Pros-N-KONS of [CLV17a] which also provide a similar trade-off between the dictionary size m and a regret bound. The forms of the regret bounds in m, μ, λ of the algorithms can be summarized as follows

$$R_n \lesssim \begin{cases} \lambda + d_{\text{eff}}(\lambda) + \frac{nm\mu}{\lambda} & \text{for PKAWV with KORS} \\ \lambda + \frac{n}{m} d_{\text{eff}}(\lambda) & \text{for Sketched-KONS} \\ m(\lambda + d_{\text{eff}}(\lambda)) + \frac{n\mu}{\lambda} & \text{for Pros-N-KONS} \end{cases}. \quad (9)$$

When $d_{\text{eff}}(\lambda) \leq (n/\lambda)^\gamma$, optimizing these bounds in λ , PKAWV performs better than Sketched-KONS as soon as $\gamma \leq 1/2$ and the latter cannot obtain the optimal rate $\lambda + d_{\text{eff}}(\lambda) = n^{\frac{\gamma}{1+\gamma}}$ if $m = o(n)$. Furthermore, because of the multiplicative factor m , Pros-N-KONS can't either reached the optimal rate even for $m = n$. Figure 1 plots the rate in n of the regret of these algorithms when enlarging the size m of the dictionary. We can see that for $\gamma = 1/4$, PKAWV is the only algorithm that achieves the optimal rate $n^{\gamma/(1+\gamma)}$ with $m = o(n)$ features. The rate of Pros-N-KONS cannot beat $4\gamma/(1+\gamma)^2$ and stops improving even when the size of the dictionary increases. This is because Pros-N-KONS is restarted whenever a point is added in the dictionary which is too costly for large dictionaries. It is worth pointing out that these rates are for a well-tuned value of λ . However, such an optimization can be performed at a small cost using expert advice algorithm on a finite grid of λ .

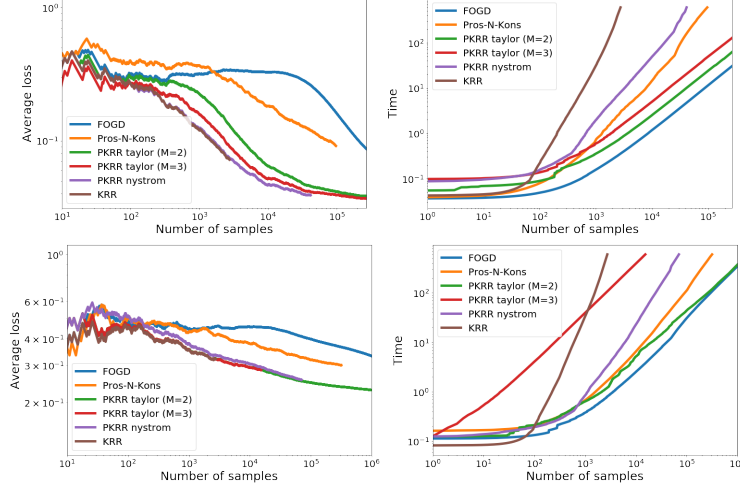


Figure 2: Average classification error and time on: (top) code-rna ($n = 2.7 \times 10^5$, $d = 8$); (bottom) SUSY ($n = 6 \times 10^6$, $d = 22$).

Beforehand known features We may assume that the sequence of feature vectors x_t is given in advance to the learner while only the outputs y_t are sequentially revealed (see [GGHS18] or [BKM⁺15] for details). In this case, the complete dictionary $\mathcal{I}_n \subset \{x_1, \dots, x_n\}$ may be computed beforehand and PKAWV can be used with the fix subspace $\tilde{\mathcal{H}} = \text{Span}(\phi(x), x \in \mathcal{I}_n)$. In this case, the regret upper-bound can be improved to $R_n \lesssim \lambda + d_{\text{eff}}(\lambda) + \frac{n\mu}{\lambda}$ by removing a factor m in the last term (see (9)).

Corollary 8. *Under the notation and assumptions of Corollary 7, PKAWV used with dictionary \mathcal{I}_n and parameter $\mu = nm^{-1/\gamma}$ achieves with high probability*

$$R_n \lesssim \begin{cases} n^{\frac{\gamma}{1+\gamma}} & \text{if } m \geq n^{\frac{2\gamma}{1+\gamma}} \text{ for } \lambda = n^{\frac{\gamma}{1+\gamma}} \\ nm^{-\frac{1}{2\gamma}} & \text{otherwise for } \lambda = nm^{-\frac{1}{2\gamma}} \end{cases}.$$

Furthermore, w.h.p. the dictionary is of size $|\mathcal{I}_n| \lesssim m$ leading to a per-round space and time complexity $O(m^2)$.

The suboptimal rate due to a small dictionary is improved by a factor \sqrt{m} compared to the “sequentially revealed features” setting. Furthermore, since $2\gamma/(1+\gamma) < 1$ for all $\gamma \in (0, 1)$, the algorithm is able to recover the optimal rate $n^{\gamma/(1+\gamma)}$ for all $\gamma \in (0, 1)$ with a dictionary of sub-linear size $m \ll n$. We leave for future work the question whether there is really a gap between these two settings or if this gap from a suboptimality of our analysis.

4 Experiments

We empirically test PKAWV against several state-of-the-art algorithms for online kernel regression. In particular, we test our algorithms in (1) an adversarial setting [see Appendix G], (2) on large scale datasets. The following algorithms have been tested:

- Kernel-AWV for adversarial setting or Kernel Ridge Regression for i.i.d. real data settings;
- Pros-N-Kons [CLV17b];
- Fourier Online Gradient Descent (FOGD, [LHW⁺16]);
- PKAWV(or PKRR for real data settings) with Taylor expansions ($M \in \{2, 3, 4\}$)
- PKAWV(or PKRR for real data settings) with Nyström

The algorithms above have been implemented in python with numpy (the code for our algorithm is in Appendix H.2). The code necessary to reproduce the following experiments is available on GitHub at <https://github.com/Remjez/kernel-online-learning>. For most algorithms, we used hyperparameters from the respective papers. For all algorithms and all experiments, we set $\sigma = 1$ [except for SUSY where $\sigma = 4$, to match accuracy results from RCR17] and $\lambda = 1$. When

using KORS, we set $\mu = 1$, $\beta = 1$ and $\varepsilon = 0.5$ as in [CLV17b]. The number of random-features in FOGD is fixed to 1000 and the learning rate η is $1/\sqrt{n}$. All experiments have been done on a single desktop computer (Intel Core i7-6700) with a timeout of 5-min per algorithm. The results of the algorithms are only recorded up to this time.

Large scale datasets. The algorithms are evaluated on four datasets from UCI machine learning repository. In particular, *casp* (regression) and *ijcnn1*, *cod-rna*, *SUSY* (classification) [see Appendix G for *casp* and *ijcnn1*] ranging from 4×10^4 to 6×10^6 datapoints. For all datasets, we scaled x in $[-1, 1]^d$ and y in $[-1, 1]$. In Figs. 2 and 4 we show the average loss (square loss for regression and classification error for classification) and the computational costs of the considered algorithm.

In all the experiments PKAWV with $M = 2$ approximates reasonably well the performance of kernel forecaster and is usually very fast. We remark that using PKAWV $M = 2$ on the first million examples of *SUSY*, we achieve in 10 minutes on a single desktop, the same average classification error obtained with specific large scale methods for i.i.d. data [RCR17], although Kernel-AWV is using a number of features reduced by a factor 100 with respect to the one used in for FALKON in the same paper. Indeed they used $r = 10^4$ Nyström centers, while with $M = 2$ we used $r = 190$ features, validating empirically the effectiveness of the chosen features for the Gaussian kernel. This shows the effectiveness of the proposed approach for large scale machine learning problems with a moderate dimension d .

References

- [ABRW18] Jason Altschuler, Francis Bach, Alessandro Rudi, and Jonathan Weed. Massively scalable sinkhorn distances via the nyström method. *arXiv preprint arXiv:1812.05189*, 2018.
- [Aro50] Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American mathematical society*, 68(3):337–404, 1950.
- [AW01] Katy S. Azoury and Manfred K. Warmuth. Relative loss bounds for on-line density estimation with the exponential family of distributions. *Machine Learning*, 43(3):211–246, 2001.
- [Bac17] Francis Bach. On the equivalence between kernel quadrature rules and random feature expansions. *Journal of Machine Learning Research*, 18(21):1–38, 2017.
- [BKM⁺15] Peter L. Bartlett, Wouter M. Koolen, Alan Malek, Eiji Takimoto, and Manfred K. Warmuth. Minimax Fixed-Design Linear Regression. *JMLR: Workshop and Conference Proceedings*, 40:1–14, 2015. Proceedings of COLT’2015.
- [BTA11] Alain Berlinet and Christine Thomas-Agnan. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.
- [CBCG04] Nicolo Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.
- [CBL06] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [CDV07] Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.
- [CKS11] Andrew Cotter, Joseph Keshet, and Nathan Srebro. Explicit approximations of the gaussian kernel. *arXiv preprint arXiv:1109.4603*, 2011.
- [CLV17a] Daniele Calandriello, Alessandro Lazaric, and Michal Valko. Efficient second-order online kernel learning with adaptive embedding. In *Neural Information Processing Systems*, 2017.

- [CLV17b] Daniele Calandriello, Alessandro Lazaric, and Michal Valko. Second-order kernel online convex optimization with adaptive sketching. In *International Conference on Machine Learning*, 2017.
- [DWH18] Michał Dereziński, Manfred K Warmuth, and Daniel Hsu. Correcting the bias in least squares regression with volume-rescaled sampling. *arXiv preprint arXiv:1810.02453*, 2018.
- [Fos91] Dean P. Foster. Prediction in the worst case. *The Annals of Statistics*, 19(2):1084–1090, 1991.
- [GGHS18] Pierre Gaillard, Sébastien Gerchinovitz, Malo Huard, and Gilles Stoltz. Uniform regret bounds over \mathbb{R}^d for the sequential linear regression problem with the square loss. *arXiv preprint arXiv:1805.11386*, 2018.
- [HM07] Elad Hazan and Nimrod Megiddo. Online learning with prior knowledge. In *International Conference on Computational Learning Theory*, pages 499–513. Springer, 2007.
- [LHW⁺16] Jing Lu, Steven CH Hoi, Jialei Wang, Peilin Zhao, and Zhi-Yong Liu. Large scale online kernel learning. *The Journal of Machine Learning Research*, 17(1):1613–1655, 2016.
- [OLBC10] Frank WJ Olver, Daniel W Lozier, Ronald F Boisvert, and Charles W Clark. *NIST handbook of mathematical functions*. Cambridge University Press, 2010.
- [RCR15] Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco. Less is more: Nyström computational regularization. In *Advances in Neural Information Processing Systems*, pages 1657–1665, 2015.
- [RCR17] Alessandro Rudi, Luigi Carratino, and Lorenzo Rosasco. Falcon: An optimal large scale kernel method. In *Advances in Neural Information Processing Systems*, pages 3888–3898, 2017.
- [RR17] Alessandro Rudi and Lorenzo Rosasco. Generalization properties of learning with random features. In *Advances in Neural Information Processing Systems*, pages 3215–3225, 2017.
- [RST13] A. Rakhlin, K. Sridharan, and A.B. Tsybakov. Empirical entropy, minimax regret and minimax risk. *Bernoulli*, 2013. To appear.
- [SHS06] Ingo Steinwart, Don Hush, and Clint Scovel. An explicit description of the reproducing kernel hilbert spaces of gaussian rbf kernels. *IEEE Transactions on Information Theory*, 52(10):4635–4643, 2006.
- [SSL00] AJ Smola, B Schölkopf, and P Langley. Sparse greedy matrix approximation for machine learning. In *17th International Conference on Machine Learning, Stanford, 2000*, pages 911–911, 2000.
- [Vov01] Vladimir Vovk. Competitive on-line statistics. *International Statistical Review*, 69(2):213–248, 2001.
- [Vov05] V. Vovk. On-line regression competitive with reproducing kernel hilbert spaces. *arXiv*, 2005.
- [Vov06] V. Vovk. Metric entropy in competitive on-line prediction. *arXiv*, 2006.
- [WS01] Christopher KI Williams and Matthias Seeger. Using the nyström method to speed up kernel machines. In *Advances in neural information processing systems*, pages 682–688, 2001.
- [ZK10] Fedor Zhdanov and Yuri Kalnishkan. An identity for kernel ridge regression. In *International Conference on Algorithmic Learning Theory*, pages 405–419. Springer, 2010.
- [ZL19] Xiao Zhang and Shizhong Liao. Incremental randomized sketching for online kernel learning. In *International Conference on Machine Learning*, pages 7394–7403, 2019.

Supplementary material

The supplementary material is organized as follows:

- Appendix A starts with notations and useful identities that are used in the rest of the proofs
- Appendix B, C, D, E, and F contain the proofs mostly in order of appearance:
 - Appendix B: statement and proof of our main theorem on which are based most of our results.
 - Appendix C: proofs of Section 2 (Propositions 1 and 2)
 - Appendix D: proofs of section 3.1 (Theorem 3 and 4)
 - Appendix E: proofs of section 3.2 (Theorem 7 and Corollaries 7 and 8)
 - Appendix F: proofs of additional lemmas.
- Appendix G provides additional experimental results (adversarial simulated data and large-scale real datasets).
- Appendix H describes efficient implementations of our algorithms together with the Python code used for the experiments.

A Notations and relevant equations

In this section, we give notations and useful identities which will be used in following proofs. We recall that at time $t \geq 1$, the forecaster is given an input $x_t \in \mathcal{X} \subset \mathbb{R}^d$, chooses a prediction function $\hat{f}_t \in \tilde{\mathcal{H}}_t \subset \mathcal{H}$, forecasts $\hat{y}_t = \hat{f}_t(x_t)$ and observes $\hat{y}_t \in [-B, B]$. Moreover, \mathcal{H} is the RKHS associated to the kernel $k : (x, x') \in \mathcal{X} \times \mathcal{X} = \phi(x)^\top \phi(x')$ for some feature map $\phi : \mathcal{X} \rightarrow \mathbb{R}^D$. We also define the following notations for all $t \geq 1$:

- $Y_t = (y_1, \dots, y_t)^\top \in \mathbb{R}^t$ and $\hat{Y}_t = (\hat{y}_1, \dots, \hat{y}_t)^\top \in \mathbb{R}^t$
- $P_t : \mathcal{H} \rightarrow \tilde{\mathcal{H}}_t$ is the Euclidean projection on $\tilde{\mathcal{H}}_t$
- $C_t := \sum_{i=1}^t \phi(x_i) \otimes \phi(x_i)$ is the covariance operator at time $t \geq 1$;
- $A_t := C_t + \lambda I$ is the regularized covariance operator;
- $S_t : \mathcal{H} \rightarrow \mathbb{R}^t$ is the operator such that $[S_t f]_i = f(x_i) = \langle f, \phi(x_i) \rangle$ for any $f \in \mathcal{H}$;
- $L_t := f \in \mathcal{H} \mapsto \|Y_t - S_t f\|^2 + \lambda \|f\|^2$ is the regularized cumulative loss.

The prediction function of PKAWV at time $t \geq 1$ is defined (see Definition 4) as

$$\hat{f}_t = \arg \min_{f \in \tilde{\mathcal{H}}_t} \left\{ \sum_{s=1}^{t-1} (y_s - f(x_s))^2 + \lambda \|f\|^2 + f(x_t)^2 \right\}.$$

Standard calculation shows the equality

$$\hat{f}_t = P_t \tilde{A}_t^{-1} P_t S_{t-1}^* Y_{t-1}. \quad (10)$$

We define also the best functions in the subspace $\tilde{\mathcal{H}}_t$ and $\tilde{\mathcal{H}}_{t+1}$ at time $t \geq 1$,

$$\hat{g}_{t+1} = \arg \min_{f \in \tilde{\mathcal{H}}_t} \{L_t(f)\} = P_t \tilde{A}_t^{-1} P_t S_t^* Y_t, \quad (11)$$

$$\tilde{g}_{t+1} = \arg \min_{f \in \tilde{\mathcal{H}}_{t+1}} \{L_t(f)\} = P_{t+1} (P_{t+1} C_t P_{t+1} + \lambda I)^{-1} P_{t+1} S_t^* Y_t, \quad (12)$$

and the best function in the whole space \mathcal{H}

$$\hat{h}_{t+1} = \arg \min_{f \in \mathcal{H}} \{L_t(f)\} = A_t^{-1} S_t^* Y_t. \quad (13)$$

B Main theorem (statement and proof)

In this appendix, we provide a general upper-bound on the regret of PKAWV that is valid for any sequence of projections P_1, \dots, P_n associated with the sequence $\tilde{\mathcal{H}}_1, \dots, \tilde{\mathcal{H}}_n$. Many of our results will be corollaries of the following theorem for specific sequences of projections.

Theorem 9. *Let $\tilde{\mathcal{H}}_1, \dots, \tilde{\mathcal{H}}_n$ be a sequence of linear subspaces of \mathcal{H} associated with projections $P_1, \dots, P_n \in \mathcal{H} \rightarrow \mathcal{H}$. PKAWV with regularization parameter $\lambda > 0$ satisfies the following upper-bound on the regret: for all $f \in \mathcal{H}$*

$$R_n(f) \leq \sum_{t=1}^n y_t^2 \left\langle \tilde{A}_t^{-1} P_t \phi(x_t), P_t \phi(x_t) \right\rangle + (\mu_t + \lambda) \frac{\mu_t t B^2}{\lambda},$$

for any sequence $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times [-B, B]$ and where $\mu_t := \|(P_{t+1} - P_t)C_t^{1/2}\|^2$ and $P_{n+1} := I$.

Proof. Let $f \in \mathcal{H}$. By definition of \hat{h}_{n+1} (see (13)), we have $L_n(\hat{h}_{n+1}) \leq L_n(f)$ which implies by definition of L_n that

$$\|Y_n - S_n \hat{h}_{n+1}\|^2 - \|Y_n - S_n f\|^2 \leq \lambda \|f\|^2 - \lambda \|\hat{h}_{n+1}\|^2. \quad (14)$$

Now, the regret can be upper-bounded as

$$R_n(f) \stackrel{(1)}{:=} \sum_{t=1}^n (y_t - \hat{y}_t)^2 - \sum_{t=1}^n (y_t - f(x_t))^2 \quad (15)$$

$$\begin{aligned} &= \|Y_n - \hat{Y}_n\|^2 - \|Y_n - S_n f\|^2 \\ &\stackrel{(14)}{\leq} \|Y_n - \hat{Y}_n\|^2 - \|Y_n - S_n \hat{h}_{n+1}\|^2 + \lambda \|f\|^2 - \lambda \|\hat{h}_{n+1}\|^2 \\ &\leq \lambda \|f\|^2 + \underbrace{\|Y_n - \hat{Y}_n\|^2 - \|Y_n - S_n \hat{g}_{n+1}\|^2 - \lambda \|\hat{g}_{n+1}\|^2}_{Z_1} \\ &\quad + \underbrace{\|Y_n - S_n \hat{g}_{n+1}\|^2 + \lambda \|\hat{g}_{n+1}\|^2 - \|Y_n - S_n \hat{h}_{n+1}\|^2 - \lambda \|\hat{h}_{n+1}\|^2}_{\Omega(n+1)} \end{aligned} \quad (16)$$

The first term Z_1 mainly corresponds to the estimation error of the algorithm: the regret incurred with respect to the best function in the approximation space $\tilde{\mathcal{H}}_n$. It also includes an approximation error due to the fact that the algorithm does not use $\tilde{\mathcal{H}}_n$ but the sequence of approximation $\tilde{\mathcal{H}}_1, \dots, \tilde{\mathcal{H}}_n$. The second term $\Omega(n+1)$ corresponds to the approximation error of \mathcal{H} by $\tilde{\mathcal{H}}_n$. Our analysis will focus on upper-bounding both of these terms separately.

Part 1. Upper-bound of the estimation error Z_1 . Using a telescoping argument together with the convention $L_0(\hat{g}_1) = 0$, we have

$$\|Y_n - S_n \hat{g}_{n+1}\|^2 + \lambda \|\hat{g}_{n+1}\|^2 = L_n(\hat{g}_{n+1}) = \sum_{t=1}^n L_t(\hat{g}_{t+1}) - L_{t-1}(\hat{g}_t).$$

Substituted into the definition of Z_1 (see (16)), the latter can be rewritten as

$$\begin{aligned} Z_1 &= \sum_{t=1}^n [(y_t - \hat{y}_t)^2 + L_{t-1}(\hat{g}_t) - L_t(\hat{g}_{t+1})] \\ &= \sum_{t=1}^n \left[\underbrace{(y_t - \hat{y}_t)^2 + L_{t-1}(\hat{g}_t) - L_t(\hat{g}_{t+1})}_{Z(t)} + \underbrace{L_{t-1}(\hat{g}_t) - L_{t-1}(\tilde{g}_t)}_{\Omega(t)} \right]. \end{aligned} \quad (17)$$

where $\tilde{g}_t = P_t(P_t C_{t-1} P_t + \lambda I)^{-1} P_t S_{t-1}^* Y_{t-1}$ is obtained by substituting P_t with P_{t-1} in the definition of \hat{g}_t . Note that with the convention $P_{n+1} = I$ the second term $\Omega(t)$ matches the definition

of $\Omega(n+1)$ of (16) since $\tilde{g}_{n+1} = A_n^{-1} S_n^* Y_n = \hat{h}_{n+1}$. In the rest of the first part we focus on upper-bounding the terms $Z(t)$. The approximation terms $\Omega(t)$ will be bounded in the next part.

Now, we remark that by expanding the square norm

$$\begin{aligned} L_t(f) &= \|Y_t\|^2 - 2Y_t^\top S_t f + \|S_t f\|^2 + \lambda \|f\|^2 = \|Y_t\|^2 - 2Y_t^\top S_t f + \langle f, C_t f \rangle + \lambda \|f\|^2 \\ &= \|Y_t\|^2 - 2Y_t^\top S_t f + \langle f, A_t f \rangle, \end{aligned} \quad (18)$$

where for the second equality, we used

$$\|S_t f\|^2 = \sum_{t=1}^n f(x_t)^2 = \sum_{t=1}^n \langle f, \phi(x_t) \rangle^2 = \sum_{t=1}^n \langle f, \phi(x_t) \otimes \phi(x_t) f \rangle = \langle f, C_t f \rangle.$$

Substituting \hat{g}_{t+1} into (18) we get

$$L_t(\hat{g}_{t+1}) = \|Y_t\|^2 - 2Y_t^\top S_t \hat{g}_{t+1} + \langle \hat{g}_{t+1}, A_t \hat{g}_{t+1} \rangle. \quad (19)$$

But, since $\hat{g}_{t+1} \in \tilde{H}_t$, we have $\hat{g}_{t+1} = P_t \tilde{g}_{t+1}$ which yields

$$Y_t^\top S_t \hat{g}_{t+1} = Y_t^\top S_t P_t \tilde{g}_{t+1} = Y_t^\top S_t \tilde{A}_t^{-1} \tilde{A}_t P_t \tilde{g}_{t+1}.$$

Then, using that $\tilde{A}_t P_t = (P_t C_t P_t + \lambda I) P_t = P_t A_t P_t$, we get

$$Y_t^\top S_t \hat{g}_{t+1} = \underbrace{Y_t^\top S_t P_t \tilde{A}_t^{-1} P_t}_{\hat{g}_{t+1}^\top} A_t \hat{g}_{t+1} = \langle \hat{g}_{t+1}, A_t \hat{g}_{t+1} \rangle.$$

Thus, combining with (19) we get

$$L_t(\hat{g}_{t+1}) = \|Y_t\|^2 - \langle \hat{g}_{t+1}, A_t \hat{g}_{t+1} \rangle.$$

Similarly, substituting \tilde{g}_t into (18) and using $\tilde{g}_t \in \tilde{\mathcal{H}}_t$, we can show

$$L_{t-1}(\tilde{g}_t) = \|Y_{t-1}\|^2 - \langle \tilde{g}_t, A_{t-1} \tilde{g}_t \rangle.$$

Combining the last two equations implies

$$L_{t-1}(\tilde{g}_t) - L_t(\hat{g}_{t+1}) = -y_t^2 + \langle \hat{g}_{t+1}, A_t \hat{g}_{t+1} \rangle - \langle \tilde{g}_t, A_{t-1} \tilde{g}_t \rangle. \quad (20)$$

Furthermore, using the definition of \hat{g}_{t+1} , we have

$$P_t A_t \hat{g}_{t+1} = P_t A_t P_t \tilde{A}_t^{-1} P_t S_t^* Y_t = P_t \tilde{A}_t \tilde{A}_t^{-1} P_t S_t^* Y_t = P_t S_t^* Y_t.$$

The same calculation with \tilde{g}_t yields

$$\begin{aligned} P_t A_{t-1} \tilde{g}_t &= P_t (C_{t-1} + \lambda I) P_t (P_t C_{t-1} P_t + \lambda I)^{-1} S_{t-1}^* Y_{t-1} \\ &= P_t (P_t C_{t-1} P_t + \lambda I) (P_t C_{t-1} P_t + \lambda I)^{-1} S_{t-1}^* Y_{t-1} = P_t S_{t-1}^* Y_{t-1}. \end{aligned} \quad (21)$$

Together with the previous equality, it entails

$$P_t A_t \hat{g}_{t+1} - P_t A_{t-1} \tilde{g}_t = P_t (S_t^* Y_t - S_{t-1}^* Y_{t-1}) = y_t P_t \phi(x_t). \quad (22)$$

Then, because $\hat{f}_t \in \tilde{\mathcal{H}}_t$, we have $\hat{y}_t = \hat{f}_t(x_t) = \langle \hat{f}_t, \phi(x_t) \rangle = \langle \hat{f}_t, P_t \phi(x_t) \rangle$. This yields

$$\begin{aligned} (y_t - \hat{y}_t)^2 &= y_t^2 - 2y_t \hat{y}_t + \hat{y}_t^2 \\ &= y_t^2 - 2 \langle \hat{f}_t, y_t P_t \phi(x_t) \rangle + \langle \hat{f}_t, \phi(x_t) \otimes \phi(x_t) \hat{f}_t \rangle \\ &\stackrel{(22)}{\leq} y_t^2 - 2 \langle \hat{f}_t, P_t A_t \hat{g}_{t+1} - P_t A_{t-1} \tilde{g}_t \rangle + \langle \hat{f}_t, \phi(x_t) \otimes \phi(x_t) \hat{f}_t \rangle \\ &= y_t^2 - 2 \langle \hat{f}_t, A_t \hat{g}_{t+1} - A_{t-1} \tilde{g}_t \rangle + \langle \hat{f}_t, (A_t - A_{t-1}) \hat{f}_t \rangle, \end{aligned} \quad (23)$$

where the last equality uses $\hat{f}_t \in \tilde{\mathcal{H}}_t$ and that $A_t - A_{t-1} = \phi(x_t) \otimes \phi(x_t)$.

Putting equations (20) and (23) together, we get

$$\begin{aligned}
Z(t) &\stackrel{(17)}{=} (y_t - \hat{y}_t)^2 + L_{t-1}(\tilde{g}_t) - L_t(\hat{g}_{t+1}) \\
&\stackrel{(20)+(23)}{\leq} \left(\langle \hat{g}_{t+1}, A_t \hat{g}_{t+1} \rangle - 2 \langle \hat{f}_t, A_t \hat{g}_{t+1} \rangle + \langle \hat{f}_t, A_t \hat{f}_t \rangle \right) \\
&\quad - \left(\langle \tilde{g}_t, A_{t-1} \tilde{g}_t \rangle - 2 \langle P_{t-1} \hat{f}_t, A_{t-1} \tilde{g}_t \rangle + \langle \hat{f}_t, A_{t-1} \hat{f}_t \rangle \right) \\
&= \langle \hat{g}_{t+1} - \hat{f}_t, A_t (\hat{g}_{t+1} - \hat{f}_t) \rangle - \underbrace{\langle \hat{f}_t - \tilde{g}_t, A_{t-1} (\hat{f}_t - \tilde{g}_t) \rangle}_{\geq 0} \\
&\leq \langle \hat{g}_{t+1} - \hat{f}_t, \tilde{A}_t (\hat{g}_{t+1} - \hat{f}_t) \rangle \\
&\stackrel{(10)+(11)}{=} \left\langle P_t \tilde{A}_t^{-1} P_t (S_t^* Y_t - S_{t-1}^* Y_{t-1}), \tilde{A}_t P_t \tilde{A}_t^{-1} P_t (S_t^* Y_t - S_{t-1}^* Y_{t-1}) \right\rangle \\
&= y_t^2 \left\langle P_t \tilde{A}_t^{-1} P_t \phi(x_t), \tilde{A}_t P_t \tilde{A}_t^{-1} P_t \phi(x_t) \right\rangle \\
&= y_t^2 \left\langle \tilde{A}_t^{-1} P_t \phi(x_t), P_t \phi(x_t) \right\rangle
\end{aligned}$$

where the last equality is because $P_t \tilde{A}_t = \tilde{A}_t P_t$ from the definition of $\tilde{A}_t := \tilde{C}_t + \lambda I$ with $\tilde{C}_t := P_t C_t P_t$.

Therefore, plugging back into (17), we have

$$Z_1 \leq \sum_{t=1}^n y_t^2 \left\langle \tilde{A}_t^{-1} P_t \phi(x_t), P_t \phi(x_t) \right\rangle + \Omega(t), \quad (24)$$

where we recall that $\Omega(t) := L_{t-1}(\hat{g}_t) - L_{t-1}(\tilde{g}_t)$.

Part 2. Upper-bound of the approximation terms $\Omega(t)$. We recall that we use the convention $P_{n+1} = I$ which does not change the algorithm. Let $t \geq 1$, expending the square losses we get

$$\begin{aligned}
\Omega(t+1) &= \sum_{s=1}^t [(\hat{g}_{t+1}(x_s) - y_s)^2 - (\tilde{g}_{t+1}(x_s) - y_s)^2 + \lambda \|\hat{g}_{t+1}\|^2 - \lambda \|\tilde{g}_{t+1}\|^2] \\
&= \sum_{s=1}^t [\mathcal{Y}_s^2 - 2 \langle \hat{g}_{t+1}, y_s \phi(x_s) \rangle + \langle \hat{g}_{t+1}, \phi(x_s) \otimes \phi(x_s) \hat{g}_{t+1} \rangle \\
&\quad - \mathcal{Y}_s^2 + 2 \langle \tilde{g}_{t+1}, y_s \phi(x_s) \rangle - \langle \tilde{g}_{t+1}, \phi(x_s) \otimes \phi(x_s) \tilde{g}_{t+1} \rangle + \lambda \|\hat{g}_{t+1}\|^2 - \lambda \|\tilde{g}_{t+1}\|^2] \\
&= 2 \langle \tilde{g}_{t+1} - \hat{g}_{t+1}, S_t^* Y_t \rangle + \langle \hat{g}_{t+1}, A_t \hat{g}_{t+1} \rangle - \langle \tilde{g}_{t+1}, A_t \tilde{g}_{t+1} \rangle
\end{aligned}$$

Since both \tilde{g}_{t+1} and \hat{g}_{t+1} belong to $\tilde{\mathcal{H}}_{t+1}$, we have

$$\Omega(t+1) = 2 \langle \tilde{g}_{t+1} - \hat{g}_{t+1}, P_{t+1} S_t^* Y_t \rangle + \langle \hat{g}_{t+1}, A_t \hat{g}_{t+1} \rangle - \langle \tilde{g}_{t+1}, A_t \tilde{g}_{t+1} \rangle,$$

which using that $P_{t+1} S_t^* Y_t = P_{t+1} A_t \tilde{g}_{t+1}$ by Equality (21) yields

$$\begin{aligned}
\Omega(t+1) &= 2 \langle \tilde{g}_{t+1} - \hat{g}_{t+1}, P_{t+1} A_t \tilde{g}_{t+1} \rangle + \langle \hat{g}_{t+1}, A_t \hat{g}_{t+1} \rangle - \langle \tilde{g}_{t+1}, A_t \tilde{g}_{t+1} \rangle \\
&= 2 \langle \tilde{g}_{t+1} - \hat{g}_{t+1}, A_t \tilde{g}_{t+1} \rangle + \langle \hat{g}_{t+1}, A_t \hat{g}_{t+1} \rangle - \langle \tilde{g}_{t+1}, A_t \tilde{g}_{t+1} \rangle \\
&= -2 \langle \hat{g}_{t+1}, A_t \tilde{g}_{t+1} \rangle + \langle \hat{g}_{t+1}, A_t \hat{g}_{t+1} \rangle + \langle \tilde{g}_{t+1}, A_t \tilde{g}_{t+1} \rangle \\
&= \langle \tilde{g}_{t+1} - \hat{g}_{t+1}, A_t (\tilde{g}_{t+1} - \hat{g}_{t+1}) \rangle.
\end{aligned}$$

Let us denote $B_t = P_{t+1} A_t P_{t+1}$. Then, remarking that $\hat{g}_{t+1} = P_t \tilde{A}_t^{-1} P_t A_t \tilde{g}_{t+1}$ and that $(P_{t+1} - P_t \tilde{A}_t^{-1} P_t A_t) P_t = 0$, we have

$$\begin{aligned}
\Omega(t+1) &= \left\langle (P_{t+1} - P_t \tilde{A}_t^{-1} P_t A_t) \tilde{g}_{t+1}, A_t (P_{t+1} - P_t \tilde{A}_t^{-1} P_t A_t) \tilde{g}_{t+1} \right\rangle \\
&= \left\langle (P_{t+1} - P_t \tilde{A}_t^{-1} P_t B_t) \tilde{g}_{t+1}, B_t (P_{t+1} - P_t \tilde{A}_t^{-1} P_t B_t) \tilde{g}_{t+1} \right\rangle \\
&= \|B_t^{1/2} (P_{t+1} - P_t \tilde{A}_t^{-1} P_t B_t) \tilde{g}_{t+1}\|^2 \\
&= \|B_t^{1/2} (P_{t+1} - P_t \tilde{A}_t^{-1} P_t B_t) (P_{t+1} - P_t) \tilde{g}_{t+1}\|^2 \\
&\leq \|B_t^{1/2} (P_{t+1} - P_t \tilde{A}_t^{-1} P_t B_t)\|^2 \|(P_{t+1} - P_t) \tilde{g}_{t+1}\|^2. \quad (25)
\end{aligned}$$

We now upper-bound the two terms of the right-hand-side. For the first one, we use that

$$\begin{aligned}
& \left\| P_{t+1} - B_t^{1/2} P_t \tilde{A}_t^{-1} P_t B_t^{1/2} \right\|^2 \\
&= \left\| P_{t+1} - 2B_t^{1/2} P_t \tilde{A}_t^{-1} P_t B_t^{1/2} + B_t^{1/2} P_t \tilde{A}_t^{-1} \underbrace{P_t B_t^{1/2} B_t^{1/2} P_t \tilde{A}_t^{-1}}_{P_t} P_t B_t^{1/2} \right\|^2 \\
&= \left\| P_{t+1} - B_t^{1/2} P_t \tilde{A}_t^{-1} P_t B_t^{1/2} \right\|^2 \in \{0, 1\},
\end{aligned} \tag{26}$$

where in the second equality we used that $P_t B_t^{1/2} B_t^{1/2} P_t \tilde{A}_t^{-1} = P_t B_t P_t \tilde{A}_t^{-1} = P_t \tilde{A}_t \tilde{A}_t^{-1} = P_t$. Therefore, using that $B_t^{1/2} P_{t+1} = P_{t+1} B_t^{1/2}$ we get

$$\begin{aligned}
\|B_t^{1/2}(P_{t+1} - P_t \tilde{A}_t^{-1} P_t B_t)\|^2 &= \left\| B_t^{1/2} \left[(P_{t+1} - P_t \tilde{A}_t^{-1} P_t B_t)(P_{t+1} - P_t) \right] \right\|^2 \\
&= \left\| (P_{t+1} - B_t^{1/2} P_t \tilde{A}_t^{-1} P_t B_t^{1/2}) B_t^{1/2} (P_{t+1} - P_t) \right\|^2 \\
&\leq \left\| P_{t+1} - B_t^{1/2} P_t \tilde{A}_t^{-1} P_t B_t^{1/2} \right\|^2 \left\| B_t^{1/2} (P_{t+1} - P_t) \right\|^2 \\
&\stackrel{(26)}{\leq} \left\| B_t^{1/2} (P_{t+1} - P_t) \right\|^2 \\
&\leq \left\| C_t^{1/2} (P_{t+1} - P_t) \right\|^2 + \lambda \\
&\leq \mu_t + \lambda,
\end{aligned}$$

where $\mu_t := \|(P_{t+1} - P_t) C_t^{1/2}\|^2$. Plugging back into (25), this yields

$$\Omega(t+1) \leq (\mu_t + \lambda) \|(P_{t+1} - P_t) \tilde{g}_{t+1}\|^2. \tag{27}$$

Then, substituting \tilde{g}_{t+1} with its definition and using $\|Y_t\|^2 \leq tB^2$, we get

$$\begin{aligned}
\|(P_{t+1} - P_t) \tilde{g}_{t+1}\|^2 &= \|(P_{t+1} - P_t) A_t^{-1} S_t^* Y_t\|^2 \\
&\stackrel{(\text{Cauchy-Schwarz})}{\leq} \|(P_{t+1} - P_t) A_t^{-1} S_t^*\|^2 \|Y_t\|^2 \\
&\leq tB^2 \|(P_{t+1} - P_t) A_t^{-1} S_t^* S_t A_t^{-1} (P_{t+1} - P_t)\| \\
&\stackrel{(C_t = S_t^* S_t)}{=} tB^2 \|(P_{t+1} - P_t) A_t^{-1} C_t A_t^{-1} (P_{t+1} - P_t)\|.
\end{aligned}$$

Because C_t and $A_t = C_t + \lambda I$ are co-diagonalizable, we have

$$C_t^{1/2} A_t^{-1} = A_t^{-1} C_t^{1/2},$$

which, together with $\|A_t^{-2}\| \leq 1/\lambda^2$ leads to

$$\begin{aligned}
\|(P_{t+1} - P_t) \tilde{g}_{t+1}\|^2 &\leq tB^2 \|(P_{t+1} - P_t) C_t^{1/2} A_t^{-2} C_t^{1/2} (P_{t+1} - P_t)\| \\
&\leq \frac{tB^2}{\lambda^2} \|(P_{t+1} - P_t) C_t^{1/2}\|^2 \\
&= \frac{t\mu_t B^2}{\lambda^2}.
\end{aligned}$$

Therefore, Inequality (27) concludes the proof of the second part

$$\Omega(t+1) \leq (\mu_t + \lambda) \frac{t\mu_t B^2}{\lambda^2}. \tag{28}$$

Conclusion of the proof. Combining (16), (24), and (28), we obtain

$$\begin{aligned}
R_n(f) &\leq \sum_{t=1}^n y_t^2 \left\langle \tilde{A}_t^{-1} P_t \phi(x_t), P_t \phi(x_t) \right\rangle + \sum_{t=1}^{n+1} \Omega(t) \\
&\leq \sum_{t=1}^n y_t^2 \left\langle \tilde{A}_t^{-1} P_t \phi(x_t), P_t \phi(x_t) \right\rangle + \sum_{t=1}^{n+1} (\mu_{t-1} + \lambda) \frac{(t-1)\mu_{t-1} B^2}{\lambda^2},
\end{aligned}$$

which concludes the proof of the theorem. \square

C Proofs of Section 2 (Kernel-AWV)

C.1 Proof of Proposition 1

First, remark that Kernel-AWV corresponds to PKAWV with $\tilde{\mathcal{H}}_t = \mathcal{H}$ and thus $P_t = I$ for all $t \geq 1$. Therefore, applying Theorem 9 with $\tilde{P}_t = I$ yields the regret bound,

$$R_n(f) \leq \lambda \|f\|^2 + \sum_{t=1}^n \langle A_t^{-1} \phi(x_t), \phi(x_t) \rangle_{\mathcal{H}},$$

for all $f \in \mathcal{H}$. The rest of the proof consists in upper-bounding the second term in the right hand side. Remarking that $A_t = A_{t-1} + \phi(x_t) \otimes \phi(x_t)$ and applying Lemma 10 stated below we have

$$\langle A_t^{-1} \phi(x_t), \phi(x_t) \rangle_{\mathcal{H}} = 1 - \frac{\det(A_{t-1}/\lambda)}{\det(A_t/\lambda)}.$$

It is worth pointing out that $\det(A_t/\lambda)$ is well defined since $A_t = I + C_t$ with $C_t = \sum_{s=1}^t \phi(x_s) \otimes \phi(x_s)$ at most of rank $t \geq 0$. Then we use $1 - u \leq \log(1/u)$ for $u > 0$ which yields

$$\langle A_t^{-1} \phi(x_t), \phi(x_t) \rangle_{\mathcal{H}} \leq \log \frac{\det(A_t/\lambda)}{\det(A_{t-1}/\lambda)}.$$

Summing over $t = 1, \dots, n$, using $A_0 = \lambda I$ and $A_n = \lambda I + C_n$ we get

$$\begin{aligned} \sum_{t=1}^n \langle A_t^{-1} \phi(x_t), \phi(x_t) \rangle_{\mathcal{H}} &\leq \log \left(\det \left(I + \frac{C_n}{\lambda} \right) \right) \\ &= \sum_{k=1}^{\infty} \log \left(1 + \frac{\lambda_k(C_n)}{\lambda} \right), \end{aligned}$$

which concludes the proof.

The following Lemma is a standard result of online matrix theory (see Lemma 11.11 of [CBL06]).

Lemma 10. *Let $V : \mathcal{H} \rightarrow \mathcal{H}$ be a linear operator. Let $u \in \mathcal{H}$ and let $U = V - u \otimes u$. Then,*

$$\langle V^{-1}u, u \rangle_{\mathcal{H}} = 1 - \frac{\det(U)}{\det(V)}.$$

C.2 Proof of Proposition 2

Using that for $x > 0$

$$\log(1+x) \leq \frac{x}{x+1} (1 + \log(1+x)),$$

and denoting by $a(\lambda)$ the quantity $a(s, \lambda) := 1 + \log(1 + s/\lambda)$, we get for any $n \geq 1$

$$\log \left(1 + \frac{\lambda_k(K_{nn})}{\lambda} \right) \leq \frac{\lambda_k(K_{nn})}{\lambda + \lambda_k(K_{nn})} a(\lambda_k(K_{nn}), \lambda).$$

Therefore, summing over $k \geq 1$ and denoting by λ_1 the largest eigenvalue of K_{nn}

$$\begin{aligned} \sum_{k=1}^n \log \left(1 + \frac{\lambda_k(K_{nn})}{\lambda} \right) &\leq a(\lambda_1, \lambda) \sum_{k=1}^n \frac{\lambda_k(K_{nn})}{\lambda + \lambda_k(K_{nn})} \\ &= a(\lambda_1, \lambda) \text{Tr} (K_{nn} (K_{nn} + \lambda I)^{-1}) \\ &= a(\lambda_1, \lambda) d_{\text{eff}}(\lambda) \end{aligned} \tag{29}$$

where the last equality is from the definition of $d_{\text{eff}}(\lambda)$. Combining with Proposition 1, substituting a and upper-bounding

$$\lambda_1(K_{nn}) \leq \text{Tr}(K_{nn}) = \sum_{t=1}^n \|\phi(x_t)\|^2 \leq n\kappa^2$$

concludes the proof.

D Proofs of Section 3.1 (PKAWV with Taylor's expansion)

D.1 Proof of Theorem 3

Applying Theorem 9 with a fix projection P and following the lines of the proof of Proposition 1 we get

$$R_n(f) \leq \lambda \|f\|^2 + B^2 \sum_{j=1}^n \log \left(1 + \frac{\lambda_j(PC_n P)}{\lambda} \right) + (\mu + \lambda) \frac{n\mu B^2}{\lambda^2},$$

where $\mu = \|(I - P)C_n^{1/2}\|^2$. Moreover we have for all $i = 1, \dots, n$ using that $\tilde{C}_n = PC_n P = PS_n S_n^* P$, we have

$$\lambda_i(\tilde{C}_n) = \lambda_i(PC_n P) = \lambda_i(PS_n^* S_n P) = \lambda_i(S_n P P S_n^*) = \lambda_i(S_n P S_n^*) \leq \lambda_i(K_{nn}).$$

D.2 Proof of Theorem 4

To apply our Thm. 3, we need first (1) to recall that the functions $g_k, k \in \mathbb{N}_0^d$ are in \mathcal{H} , (2) to show that they can approximate perfectly the kernel and (3) to quantify the approximation error of G_M for the kernel function. First we recall some important existing results about the considered set of functions. For completeness, we provide self-contained (and often shorter and simplified) proofs of the following lemmas in Appendix F.

The next lemma states that g_k with $k \in \mathbb{N}$ is an orthonormal basis for \mathcal{H} induced by the Gaussian kernel.

Lemma 11 ([SHS06]). *For any $k, k' \in \mathbb{N}_0^d$,*

$$g_k \in \mathcal{H}, \quad \|g_k\|_{\mathcal{H}} = 1, \quad \langle g_k, g_{k'} \rangle_{\mathcal{H}} = \mathbb{1}_{k=k'}.$$

Note that byproduct of the lemma, we have that $G_M \subset \mathcal{H}$ and moreover that the matrix Q is the identity, indeed $Q_{ij} = \langle g_{k_i}, g_{k_j} \rangle_{\mathcal{H}} = \mathbb{1}_{k_i=k_j}$. This means that the functions in G_M are linearly independent. Moreover the fact that $Q = I_r$ further simplifies the computation of the embedding $\tilde{\phi}$ (see (38)) in the implementation of the algorithm.

The next lemma recalls the expansion of $k(x, x')$ in terms of the given basis.

Lemma 12 ([CKS11]). *For any $x \in \mathcal{X}$,*

$$k(x, x') = \langle \phi(x), \phi(x') \rangle_{\mathcal{H}} = \sum_{k \in \mathbb{N}_0^d} g_k(x) g_k(x'). \quad (30)$$

Finally, next lemma provides approximation error of $k(x, x')$ in terms of the set of functions in G_M , when the data is contained in a ball or radius R .

Lemma 13 ([CKS11]). *Let $R > 0$. For any $x, x' \in \mathbb{R}^d$ such that $\|x\|, \|x'\| \leq R$ we have*

$$\left| k(x, x') - \sum_{g \in G_M} g(x) g(x') \right| \leq \frac{(R/\sigma)^{2M+2}}{(M+1)!}. \quad (31)$$

Now we are ready to prove Thm. 4.

Proof of point 1. First, note that $r := |G_M|$, the cardinality of G_M , corresponds to the number of monomials of the polynomial $(1 + x_1 + \dots + x_d)^d$, i.e. $r := |G_M| = \binom{M+d}{M}$. By recalling that $\binom{n}{k} \leq (en/k)^k$ for any $n, k \in \mathbb{N}$, we have

$$r = \binom{M+d}{M} = \binom{M+d}{d} \leq e^d (1 + M/d)^d.$$

We conclude the proof of the first point of the theorem, by considering that PKAWV used with the set of functions G_M consists in running the online linear regression algorithm of [Vov01, AW01] with $r := |G_M|$ features (see Appendix H for details). It incurs thus a computational cost of $O(nr^2 + nrd)$ in time (no r^3 since we don't need to invert Q which we have proven to be the identity matrix as consequence of Lemma 11) and $O(r^2)$ in memory.

Proof of point 2. By Lemma 11 we have that $G_M \subset \mathcal{H}$ and $Q = I_r$, so the functions in G_M are linearly independent. Then we can apply Thm. 3 obtaining the regret bound in Eq. (5):

$$R_n(f) \leq \lambda \|f\|^2 + B^2 \sum_{j=1}^n \log \left(1 + \frac{\lambda_j(K_{nn})}{\lambda} \right) + B^2 \frac{(\mu + \lambda)n}{\lambda^2} \mu, \quad (32)$$

where $\mu := \|(I - P)C_n^{1/2}\|^2$ and $C_n := \sum_{t=1}^n \phi(x_t) \otimes \phi(x_t)$. The proof consists in upper-bounding the last approximation term $B^2 \frac{(\mu + \lambda)n}{\lambda^2} \mu$. We start by upper-bounding μ as follows

$$\begin{aligned} \mu &:= \|(I - P)C_n^{1/2}\|^2 = \|(I - P)C_n(I - P)\| \\ &= \left\| (I - P) \sum_{t=1}^n \phi(x_t) \otimes \phi(x_t) (I - P) \right\| \\ &\leq \sum_{t=1}^n \|(I - P)\phi(x_t) \otimes \phi(x_t)(I - P)\| \\ &= \sum_{t=1}^n \|(I - P)\phi(x_t)\|^2 \\ &= \sum_{t=1}^n \langle (I - P)\phi(x_t), \phi(x_t) \rangle \\ &= \sum_{t=1}^n \langle \phi(x_t), \phi(x_t) \rangle - \langle P\phi(x_t), P\phi(x_t) \rangle \\ &= \sum_{t=1}^n k(x_t, x_t) - \|P\phi(x_t)\|^2, \end{aligned}$$

where we used that $\langle P\phi(x_t), \phi(x_t) \rangle = \langle P\phi(x_t), P\phi(x_t) \rangle$. Now, since by Lemma 11, the g_k form an orthonormal basis of \mathcal{H} , we have that

$$\|P\phi(x_t)\|^2 = \sum_{g \in G_M} g(x_t)^2.$$

where we recall that P the projection onto G_M . Therefore, by Lemma 13,

$$\mu \leq \frac{(R/\sigma)^{2M+2}n}{(M+1)!} \stackrel{\text{Stirling}}{\leq} \frac{ne^{-(M+1)\log\left(\frac{(M+1)\sigma^2}{eR^2}\right)}}{\sqrt{2\pi(M+1)}} \leq \frac{ne^{-(M+1)}}{\sqrt{2\pi(M+1)}} \stackrel{M \geq 1}{\leq} \frac{n}{9}e^{-M}, \quad (33)$$

where we used the fact that $n!$ is lower bounded by the Stirling approximation as $n! \geq \sqrt{2\pi n}e^{n\log\frac{n}{e}}$, for $n \in \mathbb{N}_0$ and that $M+1 \geq e^2R^2/\sigma^2$, so $\log\frac{M+1}{eR^2/\sigma^2} \geq 1$. Now, since $M \geq 2\log(n/(\lambda \wedge 1))$, we have $M \geq \log(n/\lambda)$ and thus

$$\mu \leq \frac{n}{9}e^{-M} \leq \frac{\lambda}{9} \leq \lambda.$$

Therefore, the approximation term in (32) is upper-bounded as

$$B^2 \frac{(\mu + \lambda)}{\lambda^2} \mu n \leq \frac{2B^2 \mu n}{\lambda} \stackrel{(33)}{\leq} \frac{2B^2 n^2 e^{-M}}{9\lambda}$$

which using again $M \geq 2\log(n/(\lambda \wedge 1))$ entails

$$B^2 \frac{(\mu + \lambda)}{\lambda^2} \mu n \leq \frac{2}{9} B^2 (\lambda \wedge \lambda^{-1}) \leq \frac{4B^2}{9} \log\left(1 + \frac{1}{\lambda}\right), \quad (34)$$

where in the last inequality we used that $(\lambda \wedge \lambda^{-1})/2 \leq \log(1 + 1/\lambda)$ for any $\lambda > 0$. Now, since $\log(1 + x)$ is concave on $[0, \infty)$, by subadditivity

$$\sum_{j=1}^n \log\left(1 + \frac{\lambda_j(K_{nn})}{\lambda}\right) \geq \log\left(1 + \sum_{j=1}^n \frac{\lambda_j(K_{nn})}{\lambda}\right).$$

By definition of trace in terms of eigenvalues and of the diagonal of K_{nn} , we have

$$\sum_{j=1}^n \lambda_j(K_{nn}) = \text{Tr}(K_{nn}) = \sum_{j=1}^n k(x_j, x_j) = n,$$

where the last step is due to the fact that for the Gaussian kernel we have $k(x, x) = 1$, for any $x \in \mathcal{X}$. Then

$$B^2 \log \left(1 + \frac{1}{\lambda} \right) \leq B^2 \log \left(1 + \frac{n}{\lambda} \right) \leq B^2 \sum_{j=1}^n \log \left(1 + \frac{\lambda_j(K_{nn})}{\lambda} \right). \quad (35)$$

Plugging back into Inequality (34) and substituting into (32) concludes the proof of the Theorem.

E Proofs of Section 3.2 (PKAWV with Nyström projections)

E.1 Proof of Theorem 6

The proof consists of a straightforward combination of Proposition 5 and Theorem 9. According to Proposition 5, with probability at least $1 - \delta$, we have for all $t \geq 1$,

$$\mu_t = \|(P_{t+1} - P_t)C_t^{1/2}\|^2 \leq \|(I - P_t)C_t^{1/2}\|^2 \mathbb{1}_{P_{t+1} \neq P_t} \leq \mu \mathbb{1}_{P_{t+1} \neq P_t},$$

with $|\mathcal{I}_n| \leq 9d_{\text{eff}}(\mu) \log(2n/\delta)^2$. Therefore, from Theorem 9, if $\mu \leq \lambda$, the regret is upper-bounded as

$$R_n(f) \leq \lambda \|f\|^2 + B^2 d_{\text{eff}}(\lambda) \log \left(e + \frac{en\kappa^2}{\lambda} \right) + 2 \frac{\mu n(|\mathcal{I}_n| + 1)B^2}{\lambda}.$$

Furthermore, similarly to any online linear regression algorithm in a m -dimensional space, the efficient implementation of the algorithm (see Appendix H) requires $O(m^2)$ space and time per iteration, where $m = |\mathcal{I}_n|$ is the size of the dictionary. This concludes the proof of the theorem.

E.2 Proof of Corollary 7

We recall that the notation \lesssim denotes an approximate inequality which is up to logarithmic multiplicative terms and may depend on unexplained constants. Here, we only consider non-constant quantities n , λ , m and μ and focus on the polynomial dependence on n . Keeping this in mind, the high-probability regret upper-bound provided by Theorem 6 can be rewritten as

$$R_n(f) \lesssim \lambda + \left(\frac{n}{\lambda} \right)^\gamma + \frac{\mu n |\mathcal{I}_n|}{\lambda}, \quad (36)$$

for all $f \in \mathcal{H}$. It only remains to optimize the parameters μ and λ . Choosing $\mu = d_{\text{eff}}^{-1}(m)$ ensures that the size of the dictionary is upper-bounded as $|\mathcal{I}_n| \lesssim d_{\text{eff}}(\mu) = m$.

Moreover, by assumption $m = d_{\text{eff}}(\mu) \leq \left(\frac{n}{\mu} \right)^\gamma$ and thus $\mu \leq nm^{-\frac{1}{\gamma}}$. Therefore, the regret is upper-bounded with high-probability as

$$R_n(f) \lesssim \lambda + \left(\frac{n}{\lambda} \right)^\gamma + \frac{n^2(m^{\frac{\gamma-1}{\gamma}} + 1)}{\lambda}. \quad (37)$$

Now, according to the value of m , two regimes are possible:

- If the dictionary is large enough, i.e., $m \geq n^{\frac{2\gamma}{1-\gamma^2}}$ then, once λ is optimized, the last term of the right-hand side is negligible. The regret upper-bound consists then in optimizing $\lambda + (n/\lambda)^\gamma$ in λ yielding to the choice $\lambda = n^{\frac{\gamma}{1+\gamma}}$. We get the upper-bound

$$R_n(f) \lesssim n^{\frac{\gamma}{1+\gamma}} + n^\gamma n^{-\frac{\gamma^2}{1+\gamma}} + n^2 n^{-\frac{\gamma}{1+\gamma}} n^{\frac{-2}{1+\gamma}} \lesssim n^{\frac{\gamma}{1+\gamma}},$$

which recovers the optimal rate in this case.

- Otherwise, if $m \leq n^{\frac{2\gamma}{1-\gamma^2}}$, then the last term of the r.h.s. of (37) is predominant. The dictionary is too small to recover the optimal regret bound. The parameter λ optimizes the trade-off $\lambda + n^2 m^{(\gamma-1)/\gamma} / \lambda$ which leads to the choice $\lambda = nm^{\frac{1}{2} - \frac{1}{2\gamma}}$. The upper-bound on the regret is then

$$R_n(f) \lesssim nm^{\frac{\gamma-1}{2\gamma}} + m^{\frac{1-\gamma}{2}} + nm^{\frac{1-\gamma}{2\gamma} + \frac{\gamma-1}{\gamma}} \lesssim nm^{\frac{\gamma-1}{2\gamma}}.$$

This concludes the proof.

E.3 Proof of Corollary 8

The proof follows the lines of the one of Theorem 6 and Corollary 7. However, here since the projections are fixed we can apply Theorem 3 instead of Theorem 9. This yields the high-probability regret upper-bound

$$R_n(f) \lesssim \lambda + \left(\frac{n}{\lambda}\right)^\gamma + \frac{\mu n}{\lambda},$$

which improves by a factor $|\mathcal{I}_n|$ the last term of the bound (36). The choice $\mu = d_{\text{eff}}^{-1}(m)$ yields with high probability $|\mathcal{I}_n| \lesssim d_{\text{eff}}(\mu) = m$ and $\mu \leq nm^{-\frac{1}{\gamma}}$ which entails

$$R_n(f) \lesssim \lambda + \left(\frac{n}{\lambda}\right)^\gamma + \frac{n^2 m^{-\frac{1}{\gamma}}}{\lambda}.$$

Similarly to Corollary 7 two regimes are possible. The size of the dictionary decides which term is preponderant in the above upper-bound:

- If $m \geq n^{\frac{2\gamma}{1+\gamma}}$ the dictionary is large enough to recover the optimal rate for the choice $\lambda = n^{\frac{\gamma}{1+\gamma}}$. Indeed it yields

$$R_n(f) \lesssim n^{\frac{\gamma}{1+\gamma}} + n^\gamma n^{-\frac{\gamma^2}{1+\gamma}} + n^{\frac{2\gamma}{1+\gamma}} n^{-\frac{\gamma}{1+\gamma}} \lesssim n^{\frac{\gamma}{1+\gamma}}$$

- Otherwise $m \leq n^{\frac{2\gamma}{1+\gamma}}$ and the choice $\lambda = n^{\frac{\gamma}{1+\gamma}}$ leads to

$$R_n(f) \lesssim n^{\frac{\gamma}{1+\gamma}} + n^\gamma n^{-\frac{\gamma^2}{1+\gamma}} + n^{\frac{2\gamma}{1+\gamma}} n^{-\frac{\gamma}{1+\gamma}} \lesssim n^{\frac{\gamma}{1+\gamma}}.$$

The last inequality is due to $m^{\frac{1}{2}} \leq nm^{-\frac{1}{\gamma} + \frac{1}{2\gamma}}$ because $\gamma \leq 1$ and $m \leq n$.

F Proofs of additional lemmas

F.1 Proof of Lemma 11

Recall the following characterization of scalar product for translation invariant kernels (i.e. $k(x, x') = v(x - x')$ for a $v : \mathbb{R}^d \rightarrow \mathbb{R}$) [see BTA11]

$$\langle f, g \rangle_{\mathcal{H}} = \int \frac{\mathcal{F}[f](\omega) \mathcal{F}[g](\omega)}{\mathcal{F}[v](\omega)},$$

where $\mathcal{F}[f]$ is the unitary Fourier transform of f . Let start from the one dimensional case and denote by \mathcal{H}_0 the Gaussian RKHS on \mathbb{R} . First note that when $d = 1$, we have $g_k = \psi_k$. Now, the Fourier transform of ψ_k is $\mathcal{F}[\psi_k](\omega) = \frac{1}{\sqrt{k!}} H_k(x/\sigma^2) e^{-\omega^2/(2\sigma^2)}$, for any $k \in \mathbb{N}_0^d$, where $H_k(x)$ is the k -th Hermite polynomial [see OLBC10, Eq. 18.17.35 pag. 457], and $\mathcal{F}[v] = e^{-\omega^2/2}$, then, by the fact that Hermite are orthogonal polynomial with respect to $e^{-\omega^2/2}$ forming a complete basis, we have

$$\langle \psi_k, \psi_{k'} \rangle_{\mathcal{H}_0} = \frac{1}{k!} \int H_k(\omega) H_{k'}(\omega) e^{-\omega^2/2} d\omega = \mathbb{1}_{k=k'}.$$

The multidimensional case is straightforward since Gaussian is a product kernel, i.e. $k(x, x') = \prod_{i=1}^d k(x^{(i)}, x'^{(i)})$ and $\mathcal{H} = \otimes_{i=1}^d \mathcal{H}_0$, so $\langle \otimes_{i=1}^d f_i, \otimes_{i=1}^d g_i \rangle_{\mathcal{H}} = \prod_{i=1}^d \langle f_i, g_i \rangle_{\mathcal{H}_0}$ [see Aro50]. Now, since $g_k = \otimes_{i=1}^d \psi_{k_i}$, we have $\langle g_k, g_{k'} \rangle_{\mathcal{H}} = \prod_{i=1}^d \langle \psi_{k_i}, \psi_{k'_i} \rangle_{\mathcal{H}_0} = \mathbb{1}_{k=k'}.$

F.2 Proof of Lemma 12

First, for $j \in \mathbb{N}_0$ define

$$Q_j(x, x') := e^{-\frac{\|x\|^2}{2\sigma^2} - \frac{\|x'\|^2}{2\sigma^2}} \frac{(x^\top x' / \sigma^2)^j}{j!}.$$

First note that, by multinomial expansion of $(x^\top x')^j$,

$$\begin{aligned} Q_j(x, x') &= \frac{e^{-\frac{\|x\|^2 + \|x'\|^2}{2\sigma^2}}}{\sigma^{2j} j!} \sum_{|t|=j} \binom{j}{t_1 \dots t_d} \prod_{i=1}^d (x^{(i)})^{t_i} (x'^{(i)})^{t_i} \\ &= \sum_{|t|=j} g_t(x) g_t(x'). \end{aligned}$$

Now note that, by Taylor expansion of $e^{x^\top x' / \sigma^2}$ we have

$$\begin{aligned} k(x, x') &= \sum_{j=0}^{\infty} Q_j(x, x') = \sum_{j=0}^{\infty} \sum_{|t|=j} g_t(x) g_t(x') \\ &= \sum_{k \in \mathbb{N}_0^d} g_k(x) g_k(x'). \end{aligned}$$

Finally, with ϕ defined as above, and the fact that g_k forms an orthonormal basis for \mathcal{H} , leads to

$$\langle \phi(x), \phi(x') \rangle = \sum_{k \in \mathbb{N}_0^d} g_k(x) g_k(x') = k(x, x').$$

F.3 Proof of Lemma 13

Here we use the same notation of the proof of Lemma 12. Since by Taylor expansion, we have that $k(x, x') = \sum_{j=0}^{\infty} Q_j(x, x')$, by mean value theorem for the function $f(s) = e^{s/\sigma^2}$, we have that there exists $c \in [0, x^\top x']$ such that

$$\begin{aligned} |k(x, x') - \sum_{j=0}^M Q_j(x, x')| &= e^{-\frac{\|x\|^2 + \|x'\|^2}{2\sigma^2}} \frac{c^{M+1}}{(M+1)!} \left. \frac{d^{M+1} e^{\frac{s}{\sigma^2}}}{ds^{M+1}} \right|_{s=c} \\ &\leq \frac{(|x^\top x'| / \sigma^2)^{M+1}}{(M+1)!} \\ &\leq \frac{(R/\sigma)^{2M+2}}{(M+1)!} \end{aligned}$$

where the last step is obtained assuming $\|x\|, \|x'\| \leq R$. Finally note that, by definition of G_M ,

$$\sum_{g \in G_M} g(x) g(x') = \sum_{|k| \leq M} g_k(x) g_k(x') = \sum_{j=0}^M Q_j(x, x').$$

G Additional experiments

Additional large scale datasets (cf. Figure 4). We provides results on two additional datasets from UCI machine learning repository : `casp` (regression) and `ijcnn1`. See section 4 for more details.

Adversarial simulated data (cf. Figure 3) In this experiment we produced the sequence $(x_t, y_t)_{t \in \mathbb{N}}$ adversarially on the regret function. In particular, given the learning algorithm, we use `scipy` as a greedy adversary i.e. at each step an optimization is done on the regret function to find (x_t, y_t) . On the right of Figure 3, we plot the simulations until $n = 80$, with $(x_t, y_t) \in [-1, 1]^d \times [-1, 1]$ where $d = 5$. We see that Kernel-AWV, which does not use any approximation, leads to the best regret.




Figure 3: Regret in adversarial setting.

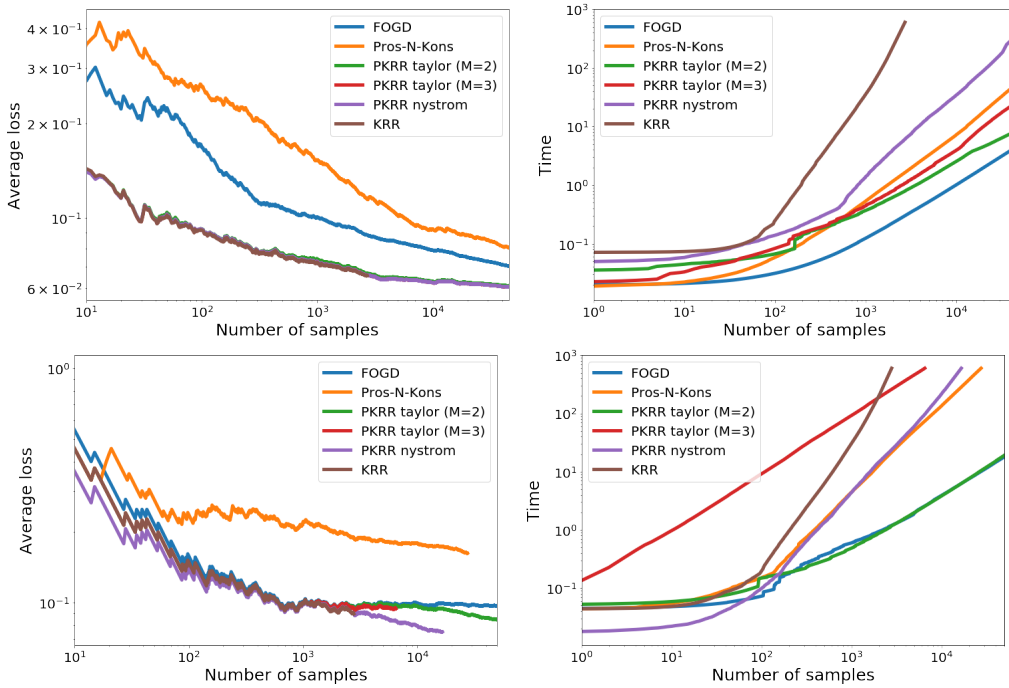


Figure 4: Average loss and time on (top): regression casp ($n = 4.5 \times 10^4$, $d = 9$); (bottom) classification ijcnn1 ($n = 1.5 \times 10^5$, $d = 22$).

Furthermore, PKAWV approximations converges very fast to the regret of Kernel-AWV when M increases. The poor performance of Pros-N-Kons is likely because of its frequent restarts which is harmful when n is small. On the contrary, FOGD has surprisingly good performance. We run the simulations up to $n = 80$ for the high computational cost required by the adversary (especially for algorithms like Kernel-AWV or Pros-N-Kons).

H Efficient implementation of PKAWV

H.1 Pseudo-code

Here, we detail how the formula (4) can be efficiently computed for the projections considered in Section 3.

Fixed embedding We consider fix sub-spaces $\tilde{\mathcal{H}}_t = \tilde{\mathcal{H}}$ induced fixed by the span of a fixed set of functions $G = \{g_1, \dots, g_r\} \subset \mathcal{H}$ as analyzed in Section 3.1. Let denote by $\tilde{\phi} : \mathcal{X} \rightarrow \mathbb{R}^r$ the map

$$\tilde{\phi}(x) = Q^{-1/2}v(x), \quad (38)$$

with $v(x) = (g_1(x), \dots, g_r(x))$, and $Q \in \mathbb{R}^{r \times r}$ defined as $Q_{ij} = \langle g_i, g_j \rangle_{\mathcal{H}}$. Then, computing the prediction $\hat{y}_t = \hat{f}_t(x_t)$ of PKAWV with

$$\hat{f}_t \in \operatorname{argmin}_{f \in \mathcal{H} = \operatorname{Span}(G)} \left\{ \sum_{s=1}^{t-1} (y_s - f(x_s))^2 + \lambda \|f\|^2 + f(x_t)^2 \right\}$$

is equivalent to embedding x_t in \mathbb{R}^r via $\tilde{\phi}$ and then performing linear AWV of [AW01, Vov01] with $\hat{y}_t = \hat{w}_t^\top \tilde{\phi}(x_t)$

$$\hat{w}_t \in \operatorname{argmin}_{w \in \mathbb{R}^r} \left\{ \sum_{s=1}^{t-1} (y_s - w^\top \tilde{\phi}(x_s))^2 + \lambda \|w\|^2 + (w^\top \tilde{\phi}(x_t))^2 \right\}.$$

This reduces the total computational complexity to $O(nr^2 + nrd + r^3)$ in time and $O(r^2)$ in space (see Algorithm 1 for an efficient implementation).

Algorithm 1 PKAWV with fixed embedding

Input: $\lambda > 0$, $\tilde{\phi} : \mathcal{X} \rightarrow \mathbb{R}^r$ for $r \geq 1$

Initialization: $A_0^{-1} = \lambda^{-1}I_r$, $b_0 = 0$

For $t = 1, \dots, n$

- receive $x_t \in \mathcal{X}$
 - compute $v_t = \tilde{\phi}(x_t) \in \mathbb{R}^r$
 - update $A_t^{-1} = A_{t-1}^{-1} - \frac{(A_{t-1}^{-1}v_t)(A_{t-1}^{-1}v_t)^\top}{1 + v_t^\top A_{t-1}^{-1}v_t}$
 - predict $\hat{y}_t = \tilde{\phi}(x_t)^\top A_t^{-1}b_{t-1}$
 - receive $y_t \in \mathbb{R}$
 - update $b_t = b_{t-1} + v_t y_t$
-

Nyström projections Here, we detail how our algorithm can be efficiently implemented with Nyström projections as considered in section 3.2. If we implement naïvely this algorithm, we would compute $\alpha_t = (K_{t,m_t}^T K_{t,m_t} + \lambda K_{m_t,m_t})^{-1} K_{t,m_t}^T Y_t$ at each iteration. However, it would require $nd_{\text{eff}}(\mu) + d_{\text{eff}}(\mu)^3$ operations per iterations. We could have update this inverse with Sherman–Morrison formula and Woodbury formula. However, in practice it leads to numeric instability because the matrix can have small eigenvalues. Here we use a method described in [RCR15]. The idea is to use the cholesky decomposition and cholup which update the cholesky decomposition when adding a rank one matrix i.e. if $A_t = L_t^T L_t$ and $A_{t+1} = A_t + u_{t+1}u_{t+1}^T$ then $L_{t+1} = \text{cholup}(L_t, u_{t+1}, '+')$. Updating the cholesky decomposition with cholup require only $d_{\text{eff}}(\mu)^2$ operations. So, PKAWV with nyström has a $O(nd_{\text{eff}}(\mu) + d_{\text{eff}}(\mu)^2)$ time complexity per iterations.

Algorithm 2 PKAWV with Nystrom projections

Input: $\lambda, \mu, \beta > 0$,

Initialization: $d_1 =$

For $t = 1, \dots, n$

- receive $x_t \in \mathcal{X}$
 - compute z_t with KORS
 - $K_t = (k(x_i, \tilde{x}_j))_{i \leq t, j \in \mathcal{I}_{t-1}}$
 - $\mathcal{I}_t = \mathcal{I}_{t-1}$
 - $a_t = (k(x_t, x_1), \dots, k(x_t, x_t))$
 - $R_t = \text{cholup}(R_t, a_t, \text{'+'})$
 - If** $z_t = 1$
 - $\mathcal{I}_t = \mathcal{I}_t \cup \{t\}$
 - $K_t = (k(x_i, x_j))_{i \leq t, j \in \mathcal{I}_t}$
 - $b_t = (k(x_t, x_j))_{j \in \mathcal{I}_t}$
 - $c_t = K_{t-1}^T a_t + \lambda b_t$
 - $d_t = a_t^T a_t + \lambda k(x_t, x_t)$
 - $g_t = \sqrt{1 + d_t}$
 - $u_t = (c_t / (1 + g_t), g_t)$
 - $v_t = (c_t / (1 + g_t), -1)$
 - $R_t = \begin{pmatrix} R_{t-1} & 0 \\ 0 & 0 \end{pmatrix}$
 - $R_t = \text{cholup}(R_t, u_t, \text{'+'})$
 - $R_t = \text{cholup}(R_t, v_t, \text{'-'})$
 - $\alpha_t = R_t^{-1} R_t^{-T} K_t^T (Y_t, 0)$
 - $b_t = (k(x_t, x_j))_{j \in \mathcal{I}_t}$
 - predict $\hat{y}_t = b_t^T \alpha_t$
 - receive $y_t \in \mathbb{R}$
 - update $Y_t = (Y_{t-1}, y_t)$
-

H.2 Python code

```
import numpy as np
from math import factorial

class PhiAWV:
    def __init__(self, d, sigma=1.0, lbd=1.0, M=2):
        self.b = None
        self.A_inv = None
        self.M = M
        self.lbd = lbd
        self.sigma = sigma

    def taylor_phi(self, x):
        res = np.array([1.])
        mm = np.array([1.])
        for k in range(1, self.M+1):
            mm = (np.outer(mm, x)).flatten()
            q = self.sigma**k*np.sqrt(factorial(k))
            res = np.concatenate((res, mm/q))
        c = 2*self.sigma**2
        res *= np.exp(-np.dot(x, x)/c)
        return np.array(res)

    def predict(self, x):
        z = self.taylor_phi(x)
        if self.b is None:
            r = len(z)
            self.b = np.zeros(r)
            self.A_inv = (1/self.lbd)*np.eye(r)

        v = np.dot(self.A_inv, z)
        v /= np.sqrt(1 + np.dot(z, v))
        self.A_inv -= np.outer(v, v)
        w_hat = np.dot(self.A_inv, self.b)
        self.z = z
        return np.dot(w_hat, z)

    def update(self, y):
        self.b += y*self.z

    def update_inv(A_inv, x):
        B = x[:-1][:, None]
        C = np.transpose(B)
        D = np.array(x[-1])[None, None]
        if A_inv.size == 0:
            return 1./D
        compl = 1./(D - np.dot(np.dot(C, A_inv), B))
        R0 = A_inv + np.dot(np.dot(np.dot(np.dot(A_inv, B), compl), C), A_inv)
        R1 = - np.dot(np.dot(A_inv, B), compl)
        R2 = - np.dot(np.dot(compl, C), A_inv)
        R3 = np.array(compl)
        return np.block([[R0, R1], [R2, R3]])

    def KORS(x, KMM, S, SKS_inv, lbd=1., eps=0.5, beta=1.):
        kS = KMM[:, -1]*S
        SKS = np.array(S)[None, :] * KMM * np.array(S)[:, None]
        en = np.eye(len(kS))[:, -1]
        SKS_inv_tmp = update_inv(SKS_inv, kS + lbd*en)
        tau = (1+eps)/lbd*(KMM[-1, -1] - np.dot(kS, np.dot(SKS_inv_tmp, kS)))
        p = max(min(beta*tau, 1), 0)
        z = np.random.binomial(1, p)
        S = S[:-1]
        if z:
            S.append(1/p)
```

```

        SKS_inv = update_inv(SKS_inv, 1/p*KMM[:, -1]*S + lbd*en)
    return z, S, SKS_inv

class Nystrom_kernel_AWV:
    def __init__(self, d, k, lbd=1.):
        self.c = np.zeros(0)
        self.X = np.zeros((0,d))
        self.Y = np.zeros(0)
        self.k = k
        self.lbd = lbd
        self.R = np.eye(0)
        self.chosen_idx = []
        self.KnM = np.eye(0)
        self.S = []
        self.SKS_inv = np.eye(0)

    def predict(self, x):
        self.X = np.concatenate((self.X, x[None, :]), axis=0)
        n = self.X.shape[0]
        Kn = np.array([self.k(self.X[i, :], x) for i in self.chosen_idx])
        self.KnM = np.concatenate((self.KnM, Kn[None, :]), axis=0)
        K_kors = np.concatenate((self.KnM[self.chosen_idx + [n-1], :],
                                   np.concatenate((Kn, [self.k(x, x)]))[:, None]),
                                   axis=1)
        z, self.S, self.SKS_inv = KORS(x, K_kors, list(self.S)+[1], \
                                       self.SKS_inv, lbd=self.lbd)
        self.R = cholup(self.R, self.KnM[-1, :], '+')
        if z:
            self.chosen_idx.append(n-1)
            M = len(self.chosen_idx)
            KM = np.array([self.k(self.X[i, :], x) for i in range(n)])
            self.KnM = np.concatenate((self.KnM, KM[:, None]), axis=1)
            a = self.KnM[:, -1].T
            d = np.dot(a, a) + self.lbd*self.KnM[-1, -1]
            if M == 1:
                self.R = np.array([[np.sqrt(d)]])
            else:
                b = self.KnM[self.chosen_idx[:-1], -1]
                c = np.dot(self.KnM[:, :-1].T, a) + self.lbd*b
                g = np.sqrt(1 + d)
                u = np.concatenate((c/(1+g), [g]))
                v = np.concatenate((c/(1+g), [-1]))
                self.R = np.block([[self.R, np.zeros((M-1, 1))],
                                   [np.zeros((1, M-1)), 0]])
                self.R = cholup(self.R, u, '+')
                self.R = cholup(self.R, v, '-')
            Yp = np.concatenate((self.Y, [0]))
            if len(self.R) > 0:
                self.c = solve_triangular(self.R,
                                           solve_triangular(self.R.T, np.dot(self.KnM.T, self.Y),
                                                                lower=True))
            Kn = np.array([self.k(self.X[i, :], x) for i in self.chosen_idx])
            return np.dot(Kn, self.c)

    def update(self, y):
        self.Y = np.concatenate((self.Y, np.array(y)[None]), axis=0)

```