

Nyström Regularization for Time Series Forecasting*

Zirui Sun

SUNZIRUI@STU.XJTU.EDU.CN

*Center for Intelligent Decision-Making and Machine Learning
School of Management
Xi'an Jiaotong University
Xi'an, China*

Mingwei Dai

DAIMW@SWUFE.EDU.CN

*Center of Statistical Research and School of Statistics
Southwestern University of Finance and Economics
Chengdu, China*

Yao Wang

YAO.S.WANG@GMAIL.COM

*Center for Intelligent Decision-Making and Machine Learning
School of Management
Xi'an Jiaotong University
Xi'an, China*

Shao-Bo Lin[†]

SBLIN1983@GMAIL.COM

*Center for Intelligent Decision-Making and Machine Learning
School of Management
Xi'an Jiaotong University
Xi'an, China*

Editor: Lorenzo Rosasco

Abstract

This paper focuses on learning rate analysis of Nyström regularization with sequential sub-sampling for τ -mixing time series. Using a recently developed Banach-valued Bernstein inequality for τ -mixing sequences and an integral operator approach based on second-order decomposition, we succeed in deriving almost optimal learning rates of Nyström regularization with sequential sub-sampling for τ -mixing time series. A series of numerical experiments are carried out to verify our theoretical results, showing the excellent learning performance of Nyström regularization with sequential sub-sampling in learning massive time series data. All these results extend the applicable range of Nyström regularization from i.i.d. samples to non-i.i.d. sequences.

Keywords: Time series forecasting, Sub-sampling, Nyström regularization, τ -mixing process.

1. Introduction

Time series is one of the most common data types abounding almost every aspect of human life (Fu, 2011), including clinical medicine, finance data, speech recognition, motion

*. Z. Sun and M. Dai contribute equally and are co-first authors of the paper

†. Corresponding author

capture data, traffic data, music recognition and video data. Besides the time-dependent nature, time series in recent years exhibit additional massiveness and hard-to-model properties, making the existing models such as autoregressive models, linear dynamical systems and hidden Markov models no more efficient (Rakthanmanon et al., 2012). In the time series forecasting community, the massiveness means a sequence over a long period of time while the hard-to-model property refers to the situation that the data structure is too complex to be captured by using the traditional models. It is thus highly desired to develop scalable learning algorithms of high quality to tackle massive time series that are not generated by simple parametric models.

One of the most productive solutions to tackle hard-to-model time series data is the nonparametric approach (Meir, 2000), which aims at finding the intrinsic nature of time series without imposing any structural restrictions on the models. Neural networks (Hagan et al., 1997) and kernel methods (Shawe-Taylor and Cristianini, 2004) are two popular nonparametric schemes in time series forecasting. In particular, it was shown in (Modha and Masry, 1996; Mondha, 1998; Xu and Chen, 2008; Steinwart and Christmann, 2009) that neural networks and kernel methods are memory-universal and possess good generalization performances, provided the time series satisfy the well known α -mixing condition (Doukhan, 1994). The problem is, however, that these existing methods require huge computations, which dampens users' enthusiasms heavily and pushes them to utilize other simple but scalable approaches, especially when the data size is huge.

This paper aims to derive a scalable kernel-based learning algorithm to tackle massive and hard-to-model time series data. Our study is motivated by three important observations. At first, the dependent nature of time series data enhances the dependence among columns of the kernel matrix, making its effective rank¹ much smaller than that of independent and identically distributed (i.i.d.) data (see Figure 2 below for detailed descriptions). Then, numerous studies showed that the dependence of time series data can be maintained via kernelization, i.e., columns of the kernel matrix and input data possess the same mixing property (Bradley, 2005). For example, it can be found in (Bradley, 2005; Sun and Lin, 2021) that the so-called α -mixing condition (Rosenblatt., 1956) is unchanged via kernelization. At last, Nyström regularization (Williams and Seeger, 2000), a special type of learning with sub-sampling that randomly sketches a few columns from the kernel matrix to build up the estimator, has been widely used for i.i.d. data. Nyström regularization (Williams and Seeger, 2000) successfully reduces the computational burden of kernel methods without losing their generalization performance (Rudi et al., 2015; Kriukova et al., 2017), provided the number of sketched columns is larger than the effective rank of the kernel matrix. Taking these interesting observations into accounts, we devote to utilizing Nyström regularization for kernel ridge regression (KRR) (Gittens and Mahoney, 2016) to yield a novel scalable learning strategy for time series.

Different from the classical Nyström regularization for i.i.d. data, Nyström regularization for time series requires strict orders of selected columns to reflect the time-dependent nature, making the widely used sub-sampling schemes including the plain Nyström (Williams and Seeger, 2000), leverage score approach (Gittens and Mahoney, 2016) and random sketching (Yang et al., 2017) no more available. Therefore, suitable sub-sampling strategies are

1. The effective rank in this paper denotes the number of eigenvalues that are larger than a specific threshold.

needed to equip Nyström regularization to tackle time series. This raises two challenges: (i) designing an exclusive sub-sampling mechanism available to time series and (ii) providing theoretical guarantees for the corresponding Nyström regularization approach.

As kernelization maintains the mixing property of time series data, we employ a simple but effective sub-sampling strategy for the first challenge, via selecting continuous columns in the kernel matrix to guarantee the order. As a result, the selected columns possess similar mixing property as the time series. We call such a sub-sampling strategy as Nyström regularization with sequential sub-sampling. There are mainly two advantages of sequential sub-sampling. One is its user-friendly nature, making the sub-sampling easy to be implemented. The other is that the mixing property of selected columns plays a crucial role in providing theoretical guarantees for the corresponding Nyström regularization.

For the second challenge, it should be noted that prior approaches for theoretical guarantees for learning with time series data require quantifying the dependence of time series via some mixing conditions, including α -mixing (Modha and Masry, 1996), β -mixing (Yu, 1994) and ϕ -mixing (Billingsley, 1968). However, most of learning rates established in (Yu, 1994; Xu and Chen, 2008; Steinwart and Christmann, 2009; Sun and Wu, 2010; Alquier and Wintenberger, 2012; Alquier et al., 2013; Hang and Steinwart, 2017) concerning the corresponding mixing data are sub-optimal, since the dependence among data reduces the effective samples and then makes the Bernstein-type inequality established in (Yu, 1994; Modha and Masry, 1996) for dependent data a little bit worse than the classical Bernstein inequality for i.i.d. data. Therefore, to provide optimal theoretical guarantees for Nyström regularization with sequential sub-sampling, it is necessary to develop a novel analysis approach such as the integral operator approach in (Sun and Wu, 2010; Sun and Lin, 2021). Fortunately, for the well known τ -mixing sequences, (Blanchard and Zadorozhnyi, 2019) have derived almost optimal learning rates for KRR via establishing a novel integral operator approach. Noting further that τ -mixing is somewhat weaker than α -mixing (Dedecker and Prieur, 2004), we borrow the idea from (Blanchard and Zadorozhnyi, 2019) to derive almost optimal learning rates of Nyström regularization for τ -mixing time series.

Our main contributions can be summarized in the following three aspects:

- *Methodology*: To tackle massive and hard-to-model time series, we propose a novel Nyström regularization with sequential sub-sampling based on kernel methods. The sequential sub-sampling mechanism succeeds in maintaining the mixing property of time series and the Nyström regularization significantly reduces the computational burden of kernel methods. Meanwhile, compared with i.i.d. data, the dependence nature of time series results in smaller effective rank of kernel matrix and thus requires smaller sub-sampling ratio² in Nyström regularization.

- *Theory*: Utilizing a recently developed Banach-valued Bernstein inequality for τ -mixing sequences (Blanchard and Zadorozhnyi, 2019), a projection-based error decomposition approach in (Rudi et al., 2015) and the second-order decomposition approach for integral operator (Guo et al., 2017), we derive almost optimal learning rates of Nyström regularization with sequential sub-sampling for τ -mixing time series. In particular, we show that, with a small sub-sampling ratio, Nyström regularization with sequential sub-sampling performs the same as KRR, provided the τ -mixing coefficients of time series decay expo-

2. The sub-sampling ratio in this paper means the ratio between the number of selected columns in kernel matrix and size of data.

nentially. This is the first result, to the best of our knowledge, to show almost optimal learning rates for learning non-i.i.d. data with the sub-sampling strategy.

- *Experiments:* Our theoretical assertions were verified by numerous toy simulations and three real-world data experiments including the BITCOIN (BTC) data, WTI data and Western Australia Weather data. Our experimental results show that Nyström regularization with sequential sub-sampling is effective and efficient to reduce the computational burden of KRR without sacrificing its excellent learning performance, provided the sampling ratio is larger than a specific value. Furthermore, we find that Nyström regularization with sequential sub-sampling is a feasible noise-extractor in the sense that it can quantify the random noise in time series. All these results show that Nyström regularization with sequential sub-sampling is a scalable and feasible strategy to tackle massive and hard-to-model time series.

The remainder of this paper is organized as follows. In Section 2, we introduce some basic properties of time series and τ -mixing sequences, and then propose the Nyström regularization with sequential sub-sampling for time series. In Section 3, we study theoretical behaviors of Nyström regularization via presenting its almost optimal learning rates. In Section 4, we compare our results with some related literature and present some discussions. In Section 5, extensive experimental studies are carried out to verify our theoretical assertions. In the last section, we present proofs for the main results.

2. Time Series Forecasting via Nyström Regularization

In this section, we introduce time series forecasting problems and then propose the Nyström regularization with sequential sub-sampling.

2.1 Time series forecasting

We are interested in a standard time series forecasting setting where the learner receives data of the form $D := D_n := \{z_t\}_{t=1}^n = \{(x_t, y_t)\}_{t=1}^n$ with $x_t \in \mathcal{X}$, $y_t \in \mathcal{Y}$ and $z_t \in \mathcal{Z} := \mathcal{X} \times \mathcal{Y}$. The aim is to learn a function $f_n : \mathcal{X} \rightarrow \mathcal{Y}$ such that $f_n(x_{n+1})$ can predict y_{n+1} well. The following are three widely used time series forecasting models.

- *Non-parametric auto-regression:* Let $d \in \mathbb{N}$ be the memory size. Assume $\mathcal{X} = \mathcal{Y}^d$ and there is an $f_0 : \mathcal{X} \rightarrow \mathcal{Y}$ such that

$$x_t = f_0(x_{t-1}, \dots, x_{t-d}) + \varepsilon_t, \quad (1)$$

where $\{\varepsilon_t\}_{t=1}^n$ are independent of x_0 .

- *Non-parametric ARX(Autoregressive Exogenous):* Let $d \in \mathbb{N}$ be the memory size. Assume $\mathcal{X} = \mathcal{Y}^{d+d'}$ for some $d' \in \mathbb{N}$. Let $\{\xi_t\}_{t=1}^n$ with $\xi_t \in \mathcal{Y}$ be a set of auxiliary variants. Assume that there is an $f_0 : \mathcal{X} \rightarrow \mathcal{Y}$ such that x_t , $t = 1, \dots, n$ are generated via

$$x_t = f_0(x_{t-1}, \dots, x_{t-d}, \xi_t, \dots, \xi_{t-d'+1}) + \varepsilon_t, \quad (2)$$

where $\{\varepsilon_t\}_{t=1}^n$ are independent of x_0 .

- *Nonlinear processes:* Let $d \in \mathbb{N}$ be the memory size. Assume $\mathcal{X} = \mathcal{Y}^d$ and there is an $f_0 : \mathcal{X} \rightarrow \mathcal{Y}$ such that

$$x_t = f_0(x_{t-1}, \dots, x_{t-d}; \varepsilon_t),$$




Figure 1: The decomposition of time series.

where $\{\varepsilon_t\}_{t=1}^n$ are independent of x_0 .

More interesting examples of time series forecasting can be found in the monograph (Doukhan, 1994; Fan and Yao, 2008), in which the following boundedness assumption is standard.

Assumption 1 Assume $\mathcal{Y} = [-M, M]$ and $\mathcal{X} = [-M, M]^d$, where $M < \infty$ is a positive constant.

Although the above boundedness assumption is widely used for i.i.d. data (Gyorfi et al., 2002; Cucker and Zhou, 2007; Steinwart and Christmann, 2008), it seems a little bit strict for time series since this assumption excludes numerous models like the non-parametric auto-regression models with Gaussian noise. We highlight that the reasons why adopted such an assumption are two folds. From the implementation side, the gathered time series data are always finitely many and the corresponding outputs are bounded in practice. From the theoretical side, our main tool for analysis is a Banach-valued Bernstein inequality for τ -mixing sequences established in (Blanchard and Zadorozhnyi, 2019), where the boundedness assumption is required. It would be interesting to generalize Assumption 1 to the sub-Gaussian assumption (unbounded samples) and derive corresponding generalization error bounds for different learning algorithms, just as (Caponnetto and DeVito, 2007; Rudi et al., 2015) did for i.i.d. data. We leave it in a future study since some novel and non-trivial concentration inequalities for time series data are required for this purpose.

Due to the dependence nature of the time series, it is not a good choice to formulate time series forecasting problem into the standard regression setting for i.i.d. data (Gyorfi et al., 2002). Instead, as shown in Figure 1, a general time series forecasting problem can be divided into searching a deterministic relation between successive samples and describing a random part that is mainly caused by the random noise. Taking non-parametric auto-regression for example, the deterministic part refers to find a good estimate, f_D , of f_0 in (1) and the random part depends on the distribution of the noise $\{\varepsilon_t\}_{t=1}^n$. Once f_D is good enough, then it can be regarded as a noise-extractor in the sense that $y_t - f_0(x_t)$ is near to ε_t and then the distribution of noise of time series data can be estimated.

2.2 τ -mixing sequences

It is well known that some restrictions on the dependence are necessary to establish satisfactory generalization error bounds for time series data. The α -mixing (or strong mixing)

condition (Modha and Masry, 1996) is one of the most popular restriction to quantify the dependence among time series and is much weaker than the so-called β -mixing condition (Yu, 1994) and ϕ -mixing condition (Billingsley, 1968). For two σ -fields \mathcal{J} and \mathcal{K} , define the α -mixing (or strong mixing) coefficient as

$$\alpha(\mathcal{J}, \mathcal{K}) := \sup_{A \in \mathcal{J}, B \in \mathcal{K}} |P(A \cap B) - P(A)P(B)|. \quad (3)$$

Denote by $\mathcal{M}_{i,j}$ the σ -field generated by random variables $z_{i:j} := (z_i, z_{i+1}, \dots, z_j)$. The α -mixing condition is defined as follows.

Definition 1 *A set of random sequence $\{z_i\}_{i=1}^\infty$ is said to satisfy the α -mixing condition (or strong mixing condition) if*

$$\alpha_j := \sup_{k \geq 1} \alpha(\mathcal{M}_{1,k}, \mathcal{M}_{k+j,\infty}) \rightarrow 0, \quad \text{as } j \rightarrow \infty. \quad (4)$$

We refer the readers to (Doukhan, 1994) for more details and examples for α -mixing sequences. Unfortunately, the α -mixing condition presented in Definition 1 is still a little bit strong, excluding some simple Markov chains and causal linear processes as follows:

- Markov chains with Bernoulli distribution: Let

$$x_t = \frac{1}{2}(x_{t-1} + \varepsilon_t) \quad (5)$$

where $\{\varepsilon_t\}_{t=1}^n$ are i.i.d. drawn according to the Bernoulli distribution $\mathcal{B}(1/2)$ and are independent of x_0 . It was proved in (Andrews, 1984) that $\alpha_j = 1/2$ for any j .

- Causal linear process: Let $(\xi_j)_{j \in \mathbb{Z}}$ be a sequence of i.i.d. random variables with values in \mathbb{R} . Define the time series $\{x_t\}_{t \in \mathbb{N}}$ with

$$x_t = \sum_{j=0}^{\infty} a_j \xi_{t-j}, \quad (6)$$

where $a_j = 2^{-j-1}$. If ξ_0 is drawn from $\mathcal{B}(1/2)$, then it can be found in (Dedecker and Prieur, 2004, P.871) that $\alpha_j = 1/4$ for any j .

Noting this dilemma, (Dedecker and Prieur, 2004, 2005) proposed a slightly weaker τ -mixing condition. Let \mathcal{C}_{Lip} be the set of bounded Lipschitz functions over \mathcal{X} . Consider

$$C_{Lip}(f) := \|f\|_{Lip(\mathcal{X})} := \sup \left\{ \frac{|f(x) - f(x')|}{\|x - x'\|_2} \mid x, x' \in \mathcal{X}, x \neq x' \right\},$$

where $\|\vec{a}\|_2$ denotes the Euclidean norm of the vector \vec{a} . It is easy to see that \mathcal{C}_{Lip} is a semi-norm. Let $(\Omega, \mathcal{M}_{1,i}, P)$ be a probability space, where Ω is a sample space, $\mathcal{M}_{1,i}$ is a σ -algebra of measurable subsets of Ω , and P is a probability measure on $\mathcal{M}_{1,i}$. Denote by $L^p(P) := L^p(\Omega, \mathcal{M}_{1,i}, P)$ the space of p -integrable functions endowed with norm $\|f\|_{L^p(P)} := \|f\|_{L^p(\Omega, \mathcal{M}_{1,i}, P)} := (\int_{\Omega} |f|^p d(P))^{1/p}$. We also consider a norm of \mathcal{C}_{Lip} of the form

$$\|g\|_{\mathcal{C}_{Lip}} := \|g\|_{L^\infty(\mathcal{X})} + C_{Lip}(g).$$

Let \mathcal{C}_1 be the ‘‘semi-ball’’ of functions $g \in \mathcal{C}_{Lip}$ such that $C_{Lip}(g) \leq 1$. Then the τ -mixing sequences can be defined as follows.

Definition 2 The τ -mixing coefficients are defined by

$$\begin{aligned} \tau_j &:= \sup\{E(\eta g(z_{i+j})) - E(\eta)E(g(z_{i+j}))\mid i \in \mathbb{N}, \\ &\quad \eta \text{ is } \mathcal{M}_{1,i}\text{-measurable and } \|\eta\|_{L^1(P)} \leq 1, g \in \mathcal{C}_1\}. \end{aligned} \quad (7)$$

We say that the sequence $\{z_i\}_{i=1}^\infty$ is τ -mixing if $\lim_{j \rightarrow \infty} \tau_j = 0$. In particular, if there are some constants $b_0 > 0, c_0 \geq 0, \gamma_0 > 0$ such that

$$\tau_j \leq c_0 \exp(-(b_0 j)^{\gamma_0}), \quad \forall j \geq 1, \quad (8)$$

then $\{z_t\}_{t=1}^\infty$ is said to be geometrically τ -mixing. If there are some constants $c_1 > 0, \gamma_1 > 0$ such that

$$\tau_j \leq c_1 j^{-\gamma_1}, \quad \forall j \geq 1, \quad (9)$$

then $\{z_t\}_{t=1}^\infty$ is said to be algebraic τ -mixing.

It can be found in (Dedecker and Prieur, 2004) that the Markov chain with Bernoulli distribution (5) and causal linear process (6) are geometrically τ -mixing, showing that τ -mixing is essentially different from α -mixing. Some basic properties of τ -mixing sequences were derived in (Dedecker and Prieur, 2004, 2005), among which the following four properties are important for our analysis.

Property 1 If $\{z_t\}_{t=1}^\infty$ is τ -mixing with coefficient τ_j , then for arbitrary $i, k \in \mathbb{N}$, $\{z_t\}_{t=i}^{i+k}$ is τ -mixing with coefficient τ_j .

Property 2 If $\{z_t\}_{t=1}^\infty$ is τ -mixing with coefficient τ_j and $h \in \mathcal{C}_{Lip}$ with Lipschitz constant C_h , then $\{h(z_t)\}_{t=1}^\infty$ is τ -mixing with coefficient $C_h \tau_j$.

Property 3 If $\{z_t\}_{t=1}^\infty$ is τ -mixing with coefficient τ_j , then for arbitrary $k, \ell \in \mathbb{N}$, $\{z_{k+t(\ell+1)}\}_{t=1}^\infty$ is τ -mixing with coefficient $c' \tau_j$, where c' is a constant satisfying $0 < c' \leq 1$.

Property 4 If $\{z_t\}_{t=1}^\infty$ is α -mixing, then $\{z_t\}_{t=1}^\infty$ is τ -mixing.

Property 1 and Property 2 can be deduced from the definition directly and play crucial roles in developing Nyström Regularization for τ -mixing series. Property 3 can also be deduced from the definition. More precisely, if $\{z_t\}_{t=1}^\infty$ is τ -mixing with coefficient τ_j , then the τ -mixing coefficient of $\{z_{k+t(\ell+1)}\}_{t=1}^\infty$ for any $k, \ell \in \mathbb{N}$ is smaller than τ_j and approaching 0 when $j \rightarrow \infty$. Property 4 established in (Dedecker and Prieur, 2004, Lemma 7) shows that the τ -mixing condition is weaker than the α -mixing condition. We highlight that Property 4 only implies $\tau_j \leq g(\alpha_j)$ for some monotonously increasing function g rather than $\tau_j \leq \alpha_j$.

2.3 Nyström regularization

Let $D := D_n := \{z_t\}_{t=1}^n = \{(x_t, y_t)\}_{t=1}^n$ with $x_t \in \mathcal{X} = [-M, M]^d$ and $y_t \in \mathcal{Y} = [-M, M]$ and $K(\cdot, \cdot)$ be a Mercer kernel and $(\mathcal{H}_K, \|\cdot\|_K)$ be the corresponding reproducing kernel Hilbert space (RKHS). Kernel ridge regression (KRR) (Evgeniou et al., 2000), given a regularization parameter $\lambda > 0$, is defined by

$$f_{D,\lambda} = \arg \min_{f \in \mathcal{H}_K} \left\{ \frac{1}{n} \sum_{t=1}^n (f(x_t) - y_t)^2 + \lambda \|f\|_K^2 \right\}. \quad (10)$$

It is easy to check that the complexities in storage and training of KRR are $\mathcal{O}(n^2)$ and $\mathcal{O}(n^3)$, respectively. As a result, KRR is difficult to tackle massive time series, although it is one of the most popular learning algorithms in the past two decades. Sub-sampling (Gittens and Mahoney, 2016) is a preferable way to reduce the computational burden of KRR.

For any subset $D_m := \{\tilde{z}_i\}_{i=1}^m := \{(\tilde{x}_i, \tilde{y}_i)\}_{i=1}^m$ of D , define

$$\mathcal{H}_{D_m} := \left\{ \sum_{i=1}^m a_i K_{\tilde{x}_i} : a_i \in \mathbb{R} \right\}. \quad (11)$$

where $K_x := K(x, \cdot)$. Nyström regularization with sub-samples D_m is then defined by

$$f_{D, D_m, \lambda} := \arg \min_{f \in \mathcal{H}_{D_m}} \frac{1}{n} \sum_{t=1}^n (f(x_t) - y_t)^2 + \lambda \|f\|_K^2. \quad (12)$$

Direct computation (Rudi et al., 2015) yields

$$f_{D, D_m, \lambda}(\cdot) = \sum_{i=1}^m \alpha_i K_{\tilde{x}_i}(\cdot), \quad (13)$$

where

$$\alpha = (\alpha_1, \dots, \alpha_m)^T = (\mathbb{K}_{nm}^T \mathbb{K}_{nm} + \lambda n \mathbb{K}_{mm})^\dagger \mathbb{K}_{nm}^T y_D, \quad (14)$$

\mathbb{A}^\dagger and \mathbb{A}^T denote the Moore-Penrose pseudo-inverse and transpose of a matrix \mathbb{A} respectively, $(\mathbb{K}_{n,m})_{t,i} = K(x_t, \tilde{x}_i)$, $(\mathbb{K}_{mm})_{k,i} = K(\tilde{x}_k, \tilde{x}_i)$ and $y_D = (y_1, \dots, y_n)^T$. In this way, it requires $\mathcal{O}(nm)$ and $\mathcal{O}(nm^2)$ complexities in storage and training respectively to derive a Nyström regularization estimator. If the sub-sampling ratio, i.e. m/n , is small, then Nyström regularization significantly reduces the computational burden of KRR.

Though theoretical behaviors of Nyström regularization have been rigorously verified for i.i.d. data in (Rudi et al., 2015), it remains open whether Nyström regularization is applicable for time series forecasting and is theoretically sound. Noting further that the basic idea of Nyström regularization (Williams and Seeger, 2000) is that the effective rank of kernel matrix is much smaller than the size of data and sub-sampling aims to find principal component of the kernel matrix, we show the applicability of sub-sampling for time series by comparing eigenvalues of the kernel matrices generated from time series and i.i.d. data respectively. In order to avoid the influence of the kernel functions on the calculation of eigenvalues, we choose two different types of kernel functions. As shown in Figure 2, we see that s_2 is always less than s_1 (s_1 and s_2 represent the number of eigenvalues larger than a fixed value of i.i.d. data and time series respectively), implying that it is easier to retain the main information after sub-sampling time series. That is, time series data allow lower sub-sampling ratio to maintain the spectrum information of kernel matrix than i.i.d. data. Therefore, we claim that time series may be more suitable than i.i.d. data for sub-sampling methods.

2.4 Nyström regularization with sequential sub-sampling

Different from i.i.d. data, time series exhibit additional difficulty in designing the sub-sampling strategy due to their non-i.i.d. nature. In particular, it is unknown whether the




Figure 2: Eigenvalues of kernel matrix. Data are generated with $x_t = 0.5 \sin(x_{t-1}) + \varepsilon_t$ (the blue line) and $f(x) = 0.5 \sin(x) + \varepsilon$ (the red line) where $\varepsilon_t, \varepsilon \sim \mathcal{B}(1/2)$. The number of samples is 3000 and we draw the largest 100 eigenvalues. The left figure and right figure are the eigenvalues under different kernel functions, which are Wendland kernel (27) below and Gaussian kernel with $\sigma = 0.5$ respectively.

selected columns satisfy certain mixing conditions. For this purpose, we propose a concept of sequential sub-sampling to guarantee the τ mixing property of sub-sampling.

Definition 3 Let $m \in \mathbb{N}$ and $m \leq n$. Randomly select $j \in [1, n - m + 1]$ according to the uniform distribution. The set $D_j^* := D_{j,m} := \{z_j, \dots, z_{j+m-1}\} := \{\tilde{z}_i\}_{i=1}^m$ is then defined to be a sequential sub-sampling of size m for D .

The idea of sequential sub-sampling is not difficult to be implemented and is close to the blockbootstrap for time series proposed in (Lahiri, 1999), which focused on sampling the sequence in time order to ensure the dependency among successive samples. It should be mentioned that there is a slight difference between our proposed sequential sub-sampling and the blockbootstrap approach in the sense that the block bootstrap is to resample the observation information in blocks to make a statistical estimate of the distribution characteristics of the population while our sub-sampling scheme devotes to sampling one by one in each sub-sampling stage and is for the purpose of maintaining the mixing property.

Due to Property 1, the τ -mixing property of D implies the τ mixing property of D_j^* . To facilitate the analysis, we also need the following assumption on the kernel.

Assumption 2 Assume $\kappa := \sup_{x \in \mathcal{X}} \sqrt{K(x, x)} \leq 1$ and there exists a $\mathcal{K} > 0$ such that $\max \left\{ \left| \frac{\partial K(x, x')}{\partial x} \right|, \left| \frac{\partial K(x, x')}{\partial x'} \right| \right\} \leq \mathcal{K}$.

Assumption 2 is mild. In particular, the boundedness assumption $\kappa \leq 1$ is satisfied for any continuous kernels with scaling. The restriction on the partial derivatives can also be satisfied for any smooth kernels. It is easy to check that almost all widely used kernels such as the Gaussian kernel, Wendland kernel, Sobolev kernel and multi-quadratic kernel (Steinwart and Christmann, 2008) satisfy Assumption 2. Based on Assumption 2 and Property 2, we obtain the following property directly (Blanchard and Zadorozhnyi, 2019).

Property 5 Under Assumption 2, if $\{z_t\}_{t=1}^n$ is τ -mixing and $\{\tilde{z}_i\}_{i=1}^m$ is a sequential sub-sampling of $\{z_t\}_{t=1}^n$, then $\{K_{\tilde{z}_i}\}_{i=1}^m$ is Banach-valued τ -mixing.

Property 5 plays an important role in guaranteeing the feasibility of sequential sub-sampling. With the proposed sequential sub-sampling scheme, we can develop an exclusive Nyström regularization for time series by using D_j^* in Definition 3 to take place of D_m in (12). The framework of Nyström regularization with sequential sub-sampling is given in Algorithm 1.

Algorithm 1: Nyström regularization with sequential sub-sampling

Input: Given $D := \{z_t\}_{t=1}^n = \{(x_t, y_t)\}_{t=1}^n$, Kernel $K(\cdot, \cdot)$, regularization parameter λ , sub-sampling size m .

Output: $\hat{f}_{D, D_j^*, \lambda}(x)$.

- 1 // Sequential sub-sampling
 - 2 $D_j^* := \{z_j, \dots, z_{j+m-1}\} := \{\tilde{z}_i\}_{i=1}^m \leftarrow \mathbf{Sub-sampling}(m, \{z_t\}_{t=1}^n)$;
 - 3 // Calculate kernel $\mathbb{K}_{nm}, \mathbb{K}_{mm}$
 - 4 $(\mathbb{K}_{nm})_{ti} \leftarrow K(x_t, \tilde{x}_i)$ for all $t = \{1, \dots, n\}, i = \{1, \dots, m\}$;
 - 5 $(\mathbb{K}_{mm})_{ki} \leftarrow K(\tilde{x}_k, \tilde{x}_i)$ for all $k = \{1, \dots, m\}, i = \{1, \dots, m\}$;
 - 6 // Calculate α
 - 7 $\alpha \leftarrow (\mathbb{K}_{nm}^T \mathbb{K}_{nm} + \lambda n \mathbb{K}_{mm})^\dagger \mathbb{K}_{nm}^T y_D$;
 - 8 **return** $\hat{f}_{D, D_j^*, \lambda}(x) \leftarrow \sum_{i=1}^m \alpha_i \cdot K(\tilde{x}_i, x)$.
-

As shown in Algorithm 1, there are two parameters, m and λ , to be tuned. It should be mentioned that λ is a model parameter that balances the bias and variance of the Nyström regularization, while m is an algorithmic parameter that reflects the trade-off between computational burden and prediction accuracy. Our analysis below will show that for appropriately selected λ , the learning performance of Nyström regularization is not sensitive to m , provided m is not extremely small. This shows the feasibility and efficiency of utilizing sub-sampling to tackle time series data. Practically, to reduce the computational burden, it is preferable to set $m = \sqrt{n}$ and then choose λ according to the well-known hold-out scheme (Gyorfi et al., 2002, Chap.7). Besides these two explicit parameters, there are also two hidden parameters in Algorithm 1, the sub-sampling location (j in Definition 3) and sub-sampling interval (ℓ in Property 3) in Algorithm 1. Our theoretical analysis and numerical results below will show that different sub-sampling strategies do not affect the learning performance of Nyström regularization with sequential sub-sampling very much, provided the τ -mixing property of the selected columns is guaranteed.

3. Theoretical Behaviors

In this section, we present our main results on analyzing learning performances of Nyström regularization with sequential sub-sampling for τ -mixing time series.

3.1 Setup and Assumptions

Our analysis is carried out in a standard random setting (Alquier and Wintenberger, 2012; Alquier et al., 2013), where $D = \{z_t\}_{t=1}^n$ is assumed to be drawn according to a joint distribution

$$\rho_{1:n} := \rho_1(\cdot) \times \rho_2(\cdot|z_1) \times \cdots \times \rho_n(\cdot|z_{1:n-1})$$

and $z_{1:k} = \{z_t\}_{t=1}^k$. The only difference is that our analysis is built upon the following identical distribution assumption of $\rho_{1:n}$ rather than the widely used stationary condition.

Assumption 3 *Assume*

$$\rho_1(\cdot) = \int_{\mathcal{Z}} \rho_{1:2}(z_1, \cdot) d\rho(z_1) = \cdots = \int_{\mathcal{Z}} \cdots \int_{\mathcal{Z}} \rho_{1:n}(z_1, \dots, z_{n-1}, \cdot) d\rho(z_1, \dots, z_{n-1}),$$

that is, the marginal distribution ρ_t is independent of t .

It should be mentioned that Assumption 3 can be regarded as the same marginal distribution restriction of $\rho_{1:n}$ and is essentially weaker than the strict (or weak) stationarity (Sun and Lin, 2021) in the sense that strict (or weak) stationarity requires the same joint distribution (or covariance) and Assumption 3 only requires the same marginal distribution.

With the help of Assumption 3, we can describe our theoretical framework as follows. Denote

$$\rho(\cdot) = \rho_1(\cdot) = \int_{\mathcal{Z}} \rho_{1:2}(z_1, \cdot) d\rho(z_1) = \cdots = \int_{\mathcal{Z}} \cdots \int_{\mathcal{Z}} \rho_{1:n}(z_1, \dots, z_{n-1}, \cdot) d\rho(z_1, \dots, z_{n-1}).$$

Noting $z_t = (x_t, y_t)$, we can write $\rho = \rho_X \times \rho(y|x)$, where ρ_X denotes the marginal distribution of ρ and $\rho(y|x)$ denotes the conditional distribution of ρ induced at $x \in \mathcal{X}$. Our aim is then to find a function f to minimize the generalization error $\int_{\mathcal{Z}} (f(x) - y)^2 d\rho$, which is minimized by the well-known regression function (Cucker and Zhou, 2007)

$$f_\rho(x) = \int_{\mathcal{Y}} y d\rho(y|x), \quad x \in \mathcal{X}.$$

Let $L_{\rho_X}^2$ be the Hilbert space of ρ_X square integrable functions on \mathcal{X} , endowed with the norm $\|\cdot\|_\rho$. Our purpose is then to bound

$$\int_{\mathcal{Z}} (f(x) - y)^2 d\rho - \int_{\mathcal{Z}} (f_\rho(x) - y)^2 d\rho = \|f - f_\rho\|_\rho^2,$$

which makes the time series forecasting problem similar as the standard least-squares regression setting in (Gyorfi et al., 2002). The only difference is that the independent assumption of samples is replaced by the following τ -mixing assumption.

Assumption 4 $D = \{z_t\}_{t=1}^n$ is a τ -mixing sequence with mixing coefficient τ_j .

The mixing condition is a standard assumption to describe the dependence of time series (Doukhan, 1994; Bradley, 2005; Alquier and Wintenberger, 2012; Alquier et al., 2013), among which β -mixing (Yu, 1994) and α -mixing (Modha and Masry, 1996) are widely used.

It can be obtained from Property 4 that there are numerous time series satisfying the τ -mixing condition, including the causal linear processes, functional autoregressive processes, and Markov kernel associated to expanding maps (Dedecker and Prieur, 2004, 2005).

Our next assumption is to quantify the regularity of f_ρ . For this purpose, we introduce the well known integral operator associated with the Mercer kernel K . Define L_K on \mathcal{H}_K (or $L^2_{\rho_X}$) by

$$L_K(f) = \int_{\mathcal{X}} K_x f(x) d\rho_X, \quad f \in \mathcal{H}_K \quad (\text{or } f \in L^2_{\rho_X}).$$

It is easy to check that L_K is a compact and positive operator. Denote further L_K^r by the r -th power of $L_K : L^2_{\rho_X} \rightarrow L^2_{\rho_X}$. We are now in position to write the following assumption.

Assumption 5 *There exists an $r > 0$ such that*

$$f_\rho = L_K^r(h_\rho), \quad \text{for some } h_\rho \in L^2_{\rho_X}. \quad (15)$$

Let $\{(\sigma_\ell, \phi_\ell)\}_{\ell=1}^\infty$ be the normalized eigen-pairs of L_K with $\sigma_1 \geq \sigma_2 \geq \dots \geq 0$. The Mercer expansion (Aronszajn, 1950) shows

$$K(x, x') = \sum_{\ell=1}^\infty \phi_\ell(x) \phi_\ell(x') = \sum_{\ell=1}^\infty \sigma_\ell \frac{\phi_\ell(x)}{\sqrt{\sigma_\ell}} \frac{\phi_\ell(x')}{\sqrt{\sigma_\ell}}$$

Then (15) is equivalent to

$$f_\rho = L_K^r h_\rho = \sum_{\ell=1}^\infty \sigma_\ell^{r-1/2} \langle h_\rho, \phi_\ell / \sqrt{\sigma_\ell} \rangle_\rho \phi_\ell. \quad (16)$$

That is, the index r in Assumption 5 determines the smoothness of f_ρ . In particular, (15) with $r = 1/2$ implies $f_\rho \in \mathcal{H}_K$ and (15) with $r = 0$ yields $f_\rho \in L^2_{\rho_X}$. Generally speaking, the larger value of r is, the smoother the function f_ρ is.

Our final assumption is to quantify the capacity of the RKHS \mathcal{H}_K via the effective dimension $\mathcal{N}(\lambda)$ (Zhang, 2005), which is defined to be the trace of the operator $(L_K + \lambda I)^{-1} L_K$, that is,

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

To obtain explicit learning rates for algorithm, we give an assumption on the decaying rate of the effective dimension as follow.

Assumption 6 *There exists some $s \in (0, 1]$ such that*

$$\mathcal{N}(\lambda) \leq C_0 \lambda^{-s}, \quad (17)$$

where $C_0 \geq 1$ is a constant independent of λ .

Condition (17) with $s = 1$ is always satisfied by taking $C_0 = \text{Tr}(L_K) \leq \kappa^2$. For $0 < s < 1$, let

$$L_K = \sum_{\ell=1}^\infty \lambda_\ell \langle \cdot, \phi_\ell \rangle_K \phi_\ell$$

be the spectral decomposition. Assumption 6 is proved in (Fischer, 2020) to be equivalent to the widely used eigenvalue decaying assumption in the literature (Caponnetto and DeVito, 2007; Zhang et al., 2015).

In summary, there are totally six assumptions in our analysis, among which Assumptions 2 and 6 concern properties of the kernel K , Assumptions 1, 3 and 5 refer to properties of the distribution ρ and Assumption 4 describes the mixing property of time series. There are numerous time series satisfying the above assumptions. Taking the one-dimensional non-parametric auto-regression for example, assume that z_0 is drawn from the invariant distribution and

$$z_t = f_0(z_{t-1}) + \varepsilon_t, \quad (18)$$

where $\{\varepsilon_t\}_{t=1}^\infty$ are independent of z_0 and drawn according to the uniform distribution on $[-a, a]$ for some $a > 0$ so that $E[\varepsilon_t] = 0$, and $f_0 \in \mathcal{H}_K$ with K a Sobolev kernel. In this way, it is easy to check that f_0 is uniformly bounded. Then it follows from (Chen and Shen, 1998, Proposition 1) that the time series generated by (18) satisfies Assumptions 1-6.

3.2 Learning rate analysis

By the aid of above assumptions, we are in a position to present our main results. Our first result is the Nyström regularization with sequential sub-sampling given in Algorithm 1 for geometrically τ -mixing time series.

Theorem 4 *Let $0 < \delta \leq 1/2$ and D_j^* be a sequential sub-sampling of size m for D with $j \in [1, n - m + 1]$. Under Assumptions 1-6 with $\frac{1}{2} \leq r \leq 1$ and $0 < s \leq 1$, if (8) holds, $\lambda \sim \left(\frac{n}{(\log n)^{1/\gamma_0}}\right)^{1/(2r+s)}$ and*

$$m \geq n^{\frac{s+1}{2r+s}} (\log m)^{1/\gamma_0} (\log n)^{-\frac{s+1}{(2r+s)\gamma_0}}, \quad (19)$$

then for any $j \in [1, n - m + 1]$, with confidence $1 - \delta$, there holds

$$\|f_{D, D_j^*, \lambda} - f_\rho\|_\rho \leq C^* n^{-\frac{r}{2r+s}} (1 + \log n)^{\frac{r}{(2r+s)\gamma_0}} \log^4 \frac{2}{\delta}, \quad (20)$$

where C^* is a constant independent of m, n, j or δ .

If the samples in D are i.i.d. drawn, the optimal learning rates for KRR $f_{D, \lambda}$, defined by (10), have been established in (Caponnetto and DeVito, 2007) in the sense that there exists a distribution ρ^* satisfying Assumptions 1, 2, 5 and 6 such that with high probability,

$$\|f_{\rho^*} - f_{D, \lambda}\|_{\rho^*} \geq C_1^* n^{-\frac{r}{2r+s}}$$

for some constant $C_1^* > 0$ independent of n . Noting that i.i.d. samples always satisfy Assumptions 3 and 4, the derived error estimate in (20) is optimal up to a logarithmic factor. Therefore, up to a logarithmic factor, the derived learning rate cannot be essentially improved. Furthermore, it should be mentioned that Theorem 4 is an extension of (Rudi et al., 2015, Theorem 1), where optimal learning rates of plain Nyström regularization for i.i.d. samples are deduced, since Theorem 4 with $\gamma_0 \rightarrow \infty$ coincides with (Rudi et al., 2015, Theorem 1). In particular, setting $\mathcal{M}_{1:n}$ as the set of distributions $\rho_{1:n} = \{\rho_1, \dots, \rho_n\}$ satisfying Assumptions 1-6, we can get the following corollary directly.

Corollary 5 *Let D_j^* be a sequential sub-sampling of size m for D with any $j \in [1, n - m + 1]$. If (8), (19) hold, and $\lambda \sim \left(\frac{n}{(\log n)^{1/\gamma_0}}\right)^{1/(2r+s)}$, then for any $j \in [1, n - m + 1]$, there holds*

$$C_1^* n^{-\frac{2r}{2r+s}} \leq \sup_{\rho_{1:n} \in \mathcal{M}_{1:n}} E[\|f_{D, D_j^*, \lambda} - f_\rho\|_{\rho_n}^2] \leq C_2^* n^{-\frac{2r}{2r+s}} (1 + \log(n))^{\frac{2r}{(2r+s)\gamma_0}},$$

where C_2^* is a constant independent of m, n, j .

Corollary 5 shows that up to a logarithmic factor, the proposed Nyström regularization with sequential sub-sampling achieve the optimal learning rates for numerous distributions. In our next corollary, we show that such excellent learning performance of Algorithm 1 is stable to the sub-sampling interval.

Corollary 6 *Let $0 < \delta \leq 1/2$ and $D_{j,k}^*$ be a sequential sub-sampling of interval length $k \in \mathbb{N}$ and size m for D with $j \in [1, n - km + 1]$, $D_{j,k}^* = \{x_j, x_{j+(k+1)}, \dots, x_{j+(m-1)(k+1)}\}$. Under Assumptions 1-6 with $\frac{1}{2} \leq r \leq 1$ and $0 < s \leq 1$, if (8) holds, $\lambda \sim \left(\frac{n}{(\log n)^{1/\gamma_0}}\right)^{1/(2r+s)}$ and*

$$m \geq n^{\frac{s+1}{2r+s}} (\log m)^{1/\gamma_0} (\log n)^{-\frac{s+1}{(2r+s)\gamma_0}}, \quad (21)$$

then for any $j \in [1, n - km + 1]$, with confidence $1 - \delta$, there holds

$$\|f_{D, D_{j,k}^*, \lambda} - f_\rho\|_\rho \leq C^{*'} n^{-\frac{r}{2r+s}} (1 + \log(n))^{\frac{r}{(2r+s)\gamma_0}} \log^4 \frac{2}{\delta}, \quad (22)$$

where $C^{*'}$ is a constant independent of m, n, j, k or δ .

Corollary 6 can be deduced from Theorem 4 and Property 3 directly. It follows from the above corollary that the learning rates of the learning performance of the proposed Nyström regularization with sequential sub-sampling is theoretically not sensitive to the interval length k . In this way, we can set $k = 1$ in practice.

Besides the starting position j and interval length k , there are another two more important tunable parameters, m and λ , in Nyström regularization. Theoretically speaking (Rudi et al., 2015), the regularization parameter λ is introduced to balance the bias and variance, while the sub-sampling size m is involved to balance computational cost and generalization error. The problem is, however, that such an ideal assertion neglects the interaction between parameters. A relation among the generalization error, m and λ is exhibited in Figure 3, where the interaction of λ, m is exhibited in affecting the generalization error of the algorithm. As shown in Figure 3, for an arbitrary fixed $\lambda \geq 0$, the sub-sampling size m can also be utilized to balance the bias and variance, just as (Lin et al., 2021) did. In our theorem, we do not consider the interaction between λ and m , and only use λ for the bias-variance trade-off purpose. Under this circumstance, our derived generalization error in (20) is non-increasing with respect to m .

Finally, we should highlight that the restriction on m in (19) is a bit strict, which makes Theorem 4 trivial for $r = 1/2$ in the sense that m should be not smaller than n . The main reason for this is due to that we adopt the integral algorithm operator based on second order decomposition for operator difference (Lin et al., 2017) in our proof. We believe




Figure 3: Relation among λ, m and generalization error (the data is generated by (25) below with $\varepsilon_t \sim \mathcal{U}(-0.7, 0.7)$). The regularization parameter λ is selected from $[2.5 \times 10^{-3} : 2.5 \times 10^{-4} : 7.5 \times 10^{-3}]$ and sub-sampling ratios are selected from $[0.01 : 0.01 : 0.1]$)

it can be relaxed by using the approach developed by Lin et al. (2020) that estimated the operator difference by using some operator concentration inequalities. Furthermore, compared with our numerical results in Section 5 in which an extremely small m (such as $m = 6$ for $n = 4000$) is sufficient for yielding a comparable generalization error with KRR, our theoretical result in (19) is too pessimistic. The main reason for such an inconsistency is that our theory holds for all distributions satisfying Assumptions 1-6, which is indeed a worst case analysis, while our numerical results only focus on some specific distributions.

Theorem 4 established learning rate analysis for geometrically τ -mixing time series, which is an extension of analysis for the classical i.i.d. samples (Rudi et al., 2015). In our next theorem, we prove that similar results also hold for Nyström regularization with sequential sub-sampling for algebraic τ -mixing time series.

Theorem 7 *Let $0 < \delta \leq 1/2$ and D_j^* be a sequential sub-sampling of size m for D with $j \in [1, n - m + 1]$. Under Assumptions 1-6 with $\frac{1}{2} \leq r \leq 1$ and $0 < s \leq 1$, if (9) holds, $\lambda = n^{-\frac{2\gamma_1}{2\gamma_1(2r+s)+2r+1}}$ and*

$$m \geq n^{\frac{2\gamma_1(s+1)+2}{2\gamma_1(2r+s)+2r+1}}, \quad (23)$$

then for any $j \in [1, n - m + 1]$ with confidence $1 - \delta$, there holds

$$\|f_{D, D_j^*, \lambda} - f_\rho\|_\rho \leq \hat{C} n^{-\frac{2\gamma_1 r}{2\gamma_1(2r+s)+2r+1}} \log^4 \frac{2}{\delta}, \quad (24)$$

where \hat{C} is a constant independent of m, n, j or δ .

As $\gamma_1 \rightarrow \infty$, meaning that D consists of i.i.d. samples, (24) becomes

$$\|f_{D, D_j^*, \lambda} - f_\rho\|_\rho \leq \hat{C} n^{-\frac{r}{2r+s}} \log^4 \frac{2}{\delta},$$

which is the optimal learning rates for KRR. This shows that (24) is a reasonable extension of the classical results for KRR (Caponnetto and DeVito, 2007) and Nyström regularization (Rudi et al., 2015) for i.i.d. samples. It should be mentioned that for $\gamma_1 < \infty$, the derived learning rate is always worse than that for i.i.d. samples. The main reason is that the dependence nature of samples sometimes reduces the effective samples, just as (Yu, 1994) did for β -mixing sequences and (Modha and Masry, 1996) did for α -mixing sequences. It should be mentioned that similar corollary as Corollary 6 also holds for algebraic τ -mixing time series by combining Theorem 7 and Property 3 to show the stability of Nyström regularization with respect to the interval length. We remove it for the sake of brevity.

4. Related Work

In this section, we present some related work to highlight the novelty of our results. For this purpose, we divide our presentation into three parts: related work on scalable kernel methods for i.i.d. samples, related work on learning rates analysis of kernel methods for non-i.i.d. samples and related work on scalable kernel methods for non-i.i.d. samples.

4.1 Scalable kernel methods for massive data

Along with the development of data mining, data of massive size are collected for the learning purpose. Scalable learning algorithms that can tackle these massive data are highly desired in the community of machine learning and numerous scalable schemes including the distributed learning (Zhang et al., 2015; Shi, 2019), localized SVM (Meister and Steinwart, 2016; Thomann et al., 2017) and learning with sub-sampling (Williams and Seeger, 2000; Gittens and Mahoney, 2016) were developed to equip kernel methods to reduce their computational burden. For example, learning with sub-sampling firstly selects centers of kernel with small size either in a random manner or by computing some leverage scores in a data dependent way, and then deduces the final estimator based on the selected centers. The feasibility of these scalable variants have been rigorously verified in (Zhang et al., 2015; Rudi et al., 2015; Meister and Steinwart, 2016; Yang et al., 2017; Lin et al., 2017; Wang et al., 2021), provided the samples are i.i.d. drawn.

The most related work on Nyström regularization for i.i.d. samples is (Rudi et al., 2015), in which learning rates of Nyström regularization with both plain sub-sampling and leverage scores were derived. Our results can be regarded as an extension of the interesting work (Rudi et al., 2015) from i.i.d. samples to τ -mixing time series. It should be mentioned that there are mainly three differences between our work and (Rudi et al., 2015), although some important tools (Lemmas 12 and 13 below) for proofs are borrowed from (Rudi et al., 2015). At first, we are interested in developing Nyström regularization for time series, which requires totally different sub-sampling mechanism. In particular, we propose a method using sequential sub-sampling approach to equip Nyström regularization to guarantee the mixing property of the selected columns of kernel matrix. Due to the non-i.i.d. nature of time series, the well developed integral operator approach in (Caponnetto and DeVito, 2007; Lin et al.,

2017; Guo et al., 2017) for i.i.d. samples is unavailable. We then turn to utilizing a Banach-valued Bernstein inequality established in a recent work (Blanchard and Zadorozhnyi, 2019) for τ -mixing sequences to derive tight bounds for differences between integral operators and their empirical counterparts. Finally, besides establishing almost optimal learning rates for the developed Nyström regularization, we numerically find that learning with sub-sampling is more suitable for time series than i.i.d. samples in the sense that the former requires smaller sub-sampling ratio than the latter due to the dependent nature of samples. In fact, our toy simulations and real data show that a sub-sampling ratio not larger than 0.05 is good enough to maintain the learning performance of KRR for time series, illustrating that Nyström regularization is practically feasible and efficient for massive time series.

4.2 Learning performance of kernel methods for weak dependent samples

It is impossible to quantify the learning rates of kernel methods for time series without presenting any restrictions on the dependence among samples (Chen and Shen, 1998), an extreme case of which is that all samples in the data set are identical. Therefore, some mixing properties (Doukhan, 1994) concerning the weak dependence among samples should be imposed to time series. α -mixing (Modha and Masry, 1996), β -mixing (Yu, 1994) and τ -mixing (Dedecker and Prieur, 2004) are three most widely used conditions in learning theory.

Based on the well developed Bernstein-type inequality for β -mixing sequences (Yu, 1994), α -mixing sequences (Modha and Masry, 1996), and τ -mixing sequences (Hang and Steinwart, 2017), learning rates of kernel methods for α -mixing, β -mixing and τ -mixing sequences have been established in (Alquier and Wintenberger, 2012; McDonald et al., 2017), (Xu and Chen, 2008; Steinwart and Christmann, 2009) and (Dedecker and Prieur, 2005; Hang and Steinwart, 2017), respectively. Unfortunately, most of the derived learning rates for non-i.i.d. data are worse than those for i.i.d. samples. This is mainly due to the dependence of mixing sequences, which reduces the number of valid samples and makes the established Bernstein-type inequality not so tight as that for i.i.d. samples. It is worth mentioning that (Alquier et al., 2013, Theorem 6) provided the optimal rates for ϕ -mixing time series by using a sharp version of Bernstein inequality for ϕ -mixing series (Samson and Paul-Marie, 2000). However, ϕ -mixing is much more restrictive than τ -mixing.

The interesting work (Sun and Wu, 2010), to the best of our knowledge, is the first result to derive learning rates for non-i.i.d. data via using the integral operator approach rather than Bernstein-type inequalities. As a result, the learning rates of KRR for α -mixing sequences can achieve the optimal learning rates for i.i.d. samples, provided the α -mixing sequences decay sufficiently fast and $s = 1$ in Assumption 6. This provides a springboard to study the learning performance of kernel methods for non-i.i.d. data, although the approach developed in (Sun and Wu, 2010) cannot be extended for the general case $0 < s \leq 1$ directly. Recently, (Blanchard and Zadorozhnyi, 2019) established a Banach-valued Bernstein inequality for τ -mixing sequences and developed a novel integral operator approach to deduce almost optimal learning rates for kernel-based spectral algorithms, provided the τ -mixing coefficients decay sufficiently fast. Compared with (Blanchard and Zadorozhnyi, 2019), there are two novelties in our work. From the theoretical side, we consider scalable variant of KRR to reduce its computational burden, while the analysis in

(Blanchard and Zadorozhnyi, 2019) is for KRR. As a result, our results require novel proof skills including the projection strategy for sub-sampling. In a word, our proof skills can be regarded as a combination of (Rudi et al., 2015) for projection strategy, (Guo et al., 2017) for second-order decomposition of operator differences and (Blanchard and Zadorozhnyi, 2019) for Banach-valued Bernstein inequality. From the numerical side, we conduct both toy simulations and three real time-series forecasting experiments to verify our theoretical findings, but (Blanchard and Zadorozhnyi, 2019) is only in a theoretical flavor. It should highlight that the numerical experiments in this paper are important in the sense that they reveal the low sub-sampling ratio in Nyström regularization for time series and imply that Nyström regularization can be used as a noise-extractor in practice.

4.3 Scalable kernel methods for massive time series

There are numerous learning approaches developed to tackle massive time series, including scalable bootstrap (Laptev et al., 2012), sketching (Indyk et al., 2000), and fast approximation correlation (Mueen et al., 2010). Though these approaches have been verified to be feasible in practice, there lack solid theoretical results to demonstrate the running mechanism and reasonability, making these methods sensitive to data.

In our recent work (Sun and Lin, 2021), we propose a distributed kernel ridge regression (DKRR) to handle massive α -mixing time series. Using some covariance inequalities for α -mixing sequences, Sun and Lin (2021) successfully derived almost optimal learning rates for DKRR under similar assumptions as this paper. There are mainly four differences between (Sun and Lin, 2021) and this paper. Firstly, we focus on Nyström regularization while (Sun and Lin, 2021) devoted to distributed learning. It should be noted that there are totally different scalable variants of KRR. In particular, our numerical results show that the proposed Nyström regularization for time series admits extremely small sub-sampling ratio while the numerical results in (Sun and Lin, 2021) showed that DKRR is infeasible if the number of samples in local machines is too small. Secondly, we are interested in τ -mixing time series, which can be regarded as an extension of α -mixing time series in (Sun and Lin, 2021). The direct consequence is that the covariance inequality for α -mixing sequences is unavailable to τ -mixing sequence and novel integral operator approach is required. Thirdly, our results are described in probability while the results in (Sun and Lin, 2021) are in the framework of expectation. It should be mentioned that learning rate analysis in probability is usually stronger than that in expectation in the sense that it is easy to derive an error estimate in expectation based on error in probability, just as our Corollary 5 shows, but not vice-verse. At last, our analysis holds for Assumptions 5 and 6 with all $\frac{1}{2} \leq r \leq 1$ and $0 < s \leq 1$, while the analysis in (Sun and Lin, 2021) imposed an additional $2r + s \geq 1$ in deriving the learning rates.

In summary, Nyström regularization with sequential sub-sampling is a novel scalable learning approach to tackle massive time series data. Different from the widely used distributed learning schemes, our approach admits small sub-sampling ratio, possesses slightly better theoretical behavior and can be used without advanced computing resources. We conclude this section with three important remarks.

Remark 8 *As the Nyström regularization proposed in Algorithm 1 requires to uniformly sample the starting point, it requires to first memorize all time series, which makes it difficult*

to design an incremental version of the algorithm, just as Rudi et al. (2015) did for i.i.d. data. It would be interesting and valuable to develop an incremental version of Algorithm 1 such that the algorithm can make prediction on the fly without keeping in memory all the previous points. The main difficulty to do this, in our opinions, lies in an ordered rank-one Cholesky updates (Rudi et al., 2015) for time series in practice and some novel error decomposition to quantify the rule of incremental implementation in theory. We will keep in studying this interesting topic in future.

Remark 9 *In this paper, the squares loss is employed in our algorithm for the regression purpose. It is natural to arise the question: is our methodology also available for other loss functions? From the algorithmic side, it is not difficult to design a similar Nyström regularization with sequential sub-sampling for general loss by replacing the squares loss in (12) with some convex and Lipschitz functions. However, from the theoretical side, our proof skills depend heavily on the integral operator approach, which requires that the corresponding optimization problem like (12) can be analytically displayed. In this way, our approach is not available to general loss in the sense that we cannot provide similar theoretical assertions as those for the squares loss.*

Remark 10 *As discussed above, under Assumption 5 with $1/2 \leq r \leq 1$, the main idea of our proof is a combination of approaches in (Rudi et al., 2015) for Nyström regularization, (Guo et al., 2017) for integral operator approach and (Blanchard and Zadorozhnyi, 2019) for bounds of operator differences. We believe that our analysis framework is also available to the out-of RKHS setting, i.e., Assumption 5 with $0 < r < 1/2$ by noticing the approaches in (Lu et al., 2019). In a word, although our proof skill is not totally novel, it is important for analyzing Nyström regularization for time series.*

5. Simulation Studies

In this section, we conduct both toy simulations and three real world time series forecasting experiments to verify our theoretical statements and show the excellent learning performance of Nyström regularization with sequential sub-sampling. Our numerical experiments were carried out in Matlab R2018b with Intel(R) Xeon(R) Gold 6248R CPU @3.00GHz 2.99GHz, Windows 10. The code is available at <https://github.com/zirsun/Nyström.git>.

5.1 Toy simulations

In this part, we carry out four simulations to verify the theoretical statements. The first simulation is to study the relationship between the sub-sampling ratio and test error (measured by RMSE: rooted mean squared error) to demonstrate the power of Nyström regularization for long time series. The second simulation focuses on illustrating the effectiveness of the proposed sequential sub-sampling by showing that the learning performance of Nyström regularization is independent of the sampling position or sampling interval. The third simulation aims to verify Theorem 4 and Theorem 7 via showing the relationship between test error and the number of the training samples. The last one is to exhibit the capability of the Nyström regularization with sequential sub-sampling as a noise-extractor.

In all simulations, we consider two time series: nonlinear model (25) with ε_t the independent noise satisfying $\varepsilon_t \sim \mathcal{U}(-0.7, 0.7)$ (here $\mathcal{U}(a, b)$ represents the uniform distribution on (a, b)), and Markov chains with Bernoulli distribution (26), where $\{\varepsilon_t\}_{t=1}^T$ are i.i.d. drawn from the Bernoulli distribution $\mathcal{B}(1/2)$ and are independent of x_0 . It should be mentioned that the time series generated by (25) is an α -mixing sequence (Alquier et al., 2013) while that generated by (26) is a τ -mixing sequence but not an α -mixing sequence (Dedecker and Prieur, 2005).

$$\text{Mechanism 1}(M_1) : \quad x_t = 0.5 \sin(x_{t-1}) + \varepsilon_t, \quad (25)$$

$$\text{Mechanism 2}(M_2) : \quad x_t = \frac{1}{2}(x_{t-1} + \varepsilon_t), \quad (26)$$

We adopt a widely used Wendland kernel (Chang et al., 2017) as follow:

$$K(x, x') = \begin{cases} (1 - \|x - x'\|_2)^4(4\|x - x'\|_2 + 1) & \text{if } 0 < \|x - x'\|_2 \leq 1 \\ 0 & \text{if } \|x - x'\|_2 > 1. \end{cases} \quad (27)$$

It is easy to check that $K(t)$ is three times differentiable.

Let N and N_{test} be the number of training samples and test points respectively. We generate $(N + N_{test} + 1)$ samples via (25) or (26). The training samples are: $\{x_t, x_{t+1}\}_{t=1}^N$ with x_0 drawn randomly according to $\mathcal{U}(0, 1)$. Meanwhile, we construct the test set: $\{x_t, x_{t+1} - \varepsilon_t\}_{t=N+1}^{N+N_{test}}$ for (25) and $\{x_t, x_{t+1} - 0.5\varepsilon_t\}_{t=N+1}^{N+N_{test}}$ for (26). In order to make a more faithful evaluation, we are concerned with one-step prediction, that is, the prediction of k th test sample is built upon $N + (k - 1)$ samples, where k is an integer and varies in the range $[1, N_{test}]$. Since our purpose in all simulations is to verify the theoretical assertions, we use the test set directly to select parameters.

Simulation 1: In this simulation, we aim at studying the relation between learning performance of Nyström regularization and sub-sampling ratio. As shown in our theoretical results in Section 3, the sub-sampling ratio controls not only the generalization performance of the proposed algorithm but also the computational complexities and memory requirements.

The number of training samples N and test samples N_{test} are 2000 and 50, respectively. We repeat the experiments 5 times to obtain the average RMSE. The regularization parameters λ are selected from $[5 \times 10^{-4} : 5 \times 10^{-4} : 0.01]$ (the first value is the lower bound of range, the second value is the step size, and the third one is the upper bound of the range) and $[5 \times 10^{-4} : 5 \times 10^{-5} : 0.001]$ via grid search for M_1 and M_2 respectively. Our numerical results are reported in Figure 4.

From Figure 4, we find three interesting phenomena: 1) for both M_1 and M_2 , the generalization capability does not decrease with respect to the sub-sampling ratio, which verifies our theoretical results in Theorem 4 in the sense that if m is larger than a specific value, then Nyström regularization with sequential sub-sampling reaches the optimal learning rates of KRR; 2) there exists a lower bound of sub-sampling ratio (e.g., about 0.05 for both M_1 and M_2), smaller than which, Nyström regularization with sequential sub-sampling degrades the learning performance of KRR dramatically. Noting that the complexity of training for Nyström regularization with sequential sub-sampling is $O(nm^2)$, this reflects the dilemma in selecting m . From the computational side, it is desired to select m as small as possible.

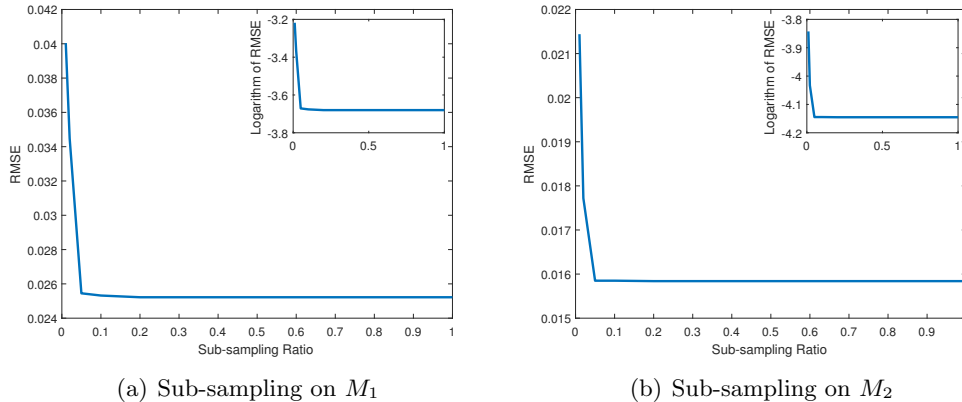


Figure 4: Relation between generalization error and sub-sampling ratio

However, too small m inevitably leads to bad generalization capability; 3) it should be highlighted that the lower bound of sub-sampling ratio to guarantee the learning performance of Nyström regularization is extremely small (about 0.05) with which it is safe to set m to be not so small. In particular, for 2000 samples, Nyström regularization with $m \geq 100$ is good enough to yield an estimator of high quality. It should be mentioned that the effective sub-sampling ratio for time series is much less than that for i.i.d. sampling (Rudi et al., 2015). The main reason, as shown in Figure 2, is that the dependence nature of samples enhances the dependence among columns in the kernel matrix. As a result, the effective rank of kernel matrix of time series is smaller than that of i.i.d. samples.

Simulation 2: In this simulation, we pursue the role of sub-sampling strategy. As shown in Theorem 4, the learning performance of Nyström regularization is independent of the position of sub-sampling. It is thus urgent to verify such an independence. Additionally, we also show the role of sampling intervals in Nyström regularization via comparing the proposed sequentially sub-sampling set D_j^* in Definition 3 with $D_{j,k}^* := \{x_j, x_{j+(k+1)}, x_{j+2(k+1)}, \dots, x_{j+(m-1)(k+1)}\}$.

The number of training samples N is fixed as $N = 2000, 10000$, and the sub-sampling ratio is fixed as 0.01, which means that the sub-sampling size is fixed at 20, 100, respectively. To test on sub-sampling at arbitrary positions, we sub-sample the first 20 (or 100) samples, the middle 20 (or 100) samples and the last 20 (or 100) samples of the training data sequence respectively. Meanwhile, we choose different sub-sampling intervals with k varying in $\{5, 10, 15, 20, 50, 100\}$. The algorithm is employed to predict 5 testing samples and the experiments are repeated 20 times. The regularization parameter λ is set as:

- For M_1 with 2000 training samples, the λ is selected from $[5 \times 10^{-4} : 5 \times 10^{-4} : 0.01]$;
- For M_1 with 10000 training samples, the λ is selected from $[5 \times 10^{-5} : 1 \times 10^{-4} : 0.001]$;
- For M_2 with 2000 training samples, the λ is selected from $[5 \times 10^{-4} : 5 \times 10^{-5} : 0.001]$;
- For M_2 with 10000 training samples, the λ is selected from $[1 \times 10^{-4} : 2 \times 10^{-5} : 4 \times 10^{-4}]$;

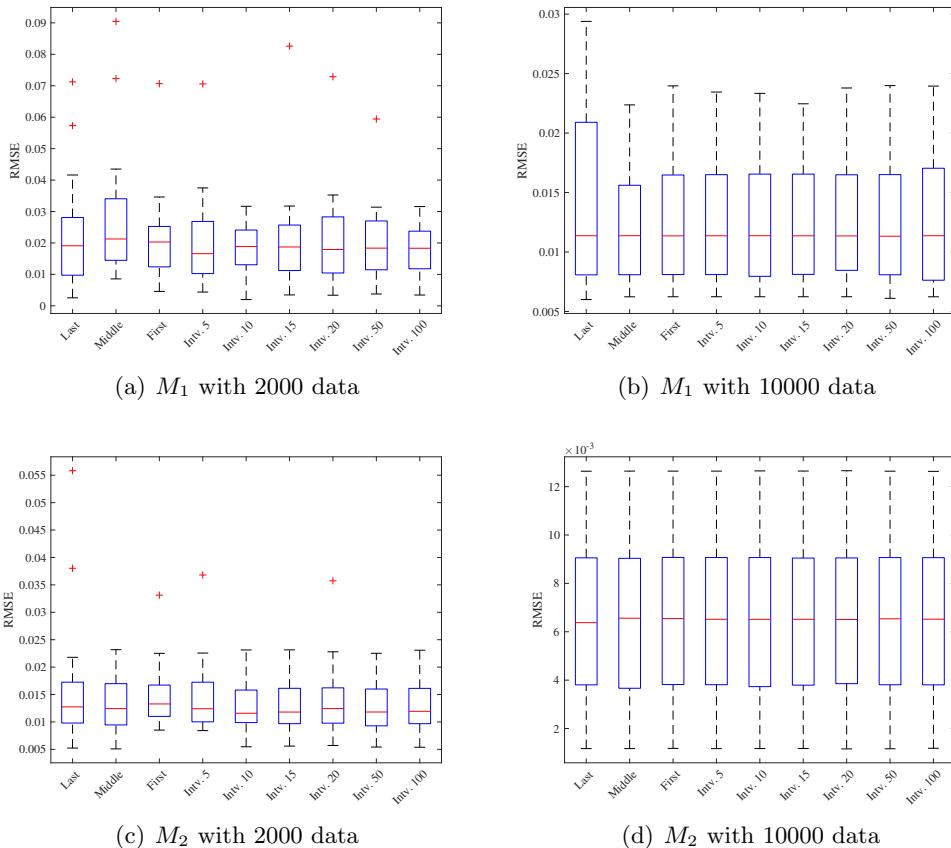


Figure 5: Role of sub-sampling mechanisms in Nyström regularization

Our numerical results are shown in Figure 5, where “Last”, “Middle” and “First” mean $j = N - m + 1$, $j = N/2$ and $j = 1$ respectively, while “Intv.a” means $k = a$ for $a = 5, 10, 15, 20, 50, 100$ in $D_{j,k}^*$. For $N = 2000$, though the mean of RMSE is almost the same for all sub-sampling mechanisms, as shown in Figure 5 (a) and (c), the generalization capability of Nyström regularization changes slightly with respect to different sub-sampling mechanism in each trivial. The reason is that for $N = 2000$, there are only 20 columns selected to build up the estimators. When the number of samples increases, for the same sub-sampling ratio, the number of selected columns also increases. Then, it can be found in Figure 5 (b) and (d) that the generalization capability of Nyström regularization is independent of the sub-sampling mechanism. All the numerical results above show that the proposed sequential sub-sampling strategy is a good choice in practice.

Simulation 3: In this simulation, we study the relation between the generalization capability and number of samples to verify Theorem 4 and Theorem 7. We build the Nyström regularization estimator on data sets of size varying in $\{2000, 5000, 10000, 20000, 50000\}$ and fix the number of testing samples as 10. We consider Nyström regularization with four sub-sampling ratios: 0.002, 0.005, 0.01 and 0.5. The regularization parameter λ is selected

from $[2 \times 10^{-4} : 2 \times 10^{-4} : 0.004]$ and $[1 \times 10^{-4} : 1 \times 10^{-5} : 2 \times 10^{-4}]$ for M_1 and M_2 , respectively. The experimental results can be found in Figure 6.

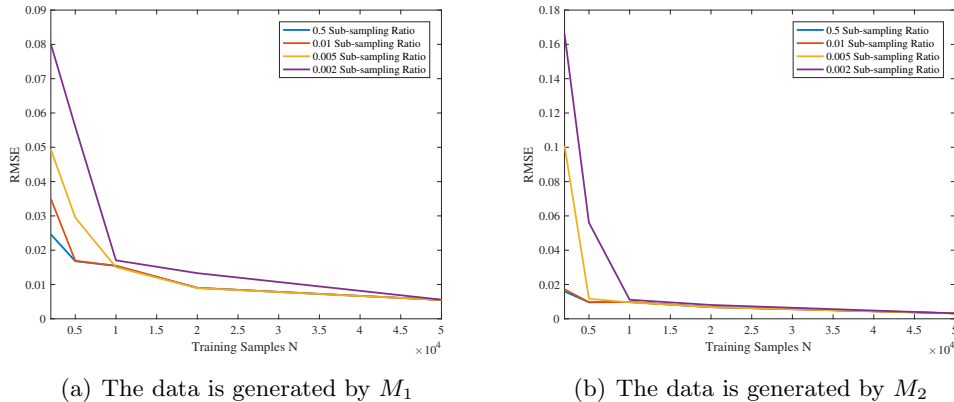


Figure 6: The relation between generalization error and number of samples

Figure 6 exhibits two findings for the proposed Nyström regularization. The one is that RMSE decreases with respect to the number of samples. This partly verifies Theorem 4 and Theorem 7, since our theoretical analysis is for the worst case. It is difficult for KRR to get a similar curve as Figure 6 since KRR requires $\mathcal{O}(N^3)$ computational complexity and thus needs at least $\mathcal{O}(10^{14})$ floating computations for 50000 training samples, which requires advanced computational sources. The other is that Nyström regularization with different sub-sampling ratios perform similarly, provided the number of samples is larger than a specific value. Taking M_2 for example, four sub-sampling ratios perform almost the same when $N \geq 10000$. The main reason is that, for small size training data, too small sampling ratio results in extremely small hypothesis space and then degrades the generalization capability of KRR. Both findings verify our theoretical assertions and show the efficiency of the sequential sub-sampling scheme.

Simulation 4: In previous simulations, the performance of Nyström regularization were tested on clean test data. However, in real-world time series forecasting, it is impossible to neglect the noise, making the machine learning estimator always lag behind the real-world time series (Chen and Shen, 1998). As shown in Figure 1, a real-world time series can be divided into a deterministic part to quantify the relation between samples in successive time and a random part to describe the uncertainty of time series. Our theorems and simulations showed that Nyström regularization is capable of discovering the deterministic part of time series, but similar as all existing machine learning approaches, it is difficult to catch the random part of time series. This makes our approach not so good for time series with large random noise. However, since the deterministic part of time series is well approximated by the proposed Nyström regularization, we can use the derived estimator to be a noise-extractor to pursue the distribution of the random noise of time series. In this simulation, we focus on the performance of Nyström regularization in extracting the random noise.

For this purpose, we generated three time series according to (25), with ε_t i.i.d. drawn from $\mathcal{B}(0.5)$, $\mathcal{U}[-0.2, 0.2]$ and $\mathcal{N}(0, 0.1^2)$, respectively. In this simulation, we set the number of training samples N as 2000, $\lambda = 0.005$ and the sub-sampling ratio as 0.01. The number

Table 1: Parameters of statistical bar

Fig.7(d)	$(-\infty, -1), [-1, -0.8), [-0.8, -0.6), [-0.6, -0.4), [-0.4, -0.2), [-0.2, 0), [0, 0.2), [0.2, 0.4), [0.4, 0.6), [0.6, 0.8), [0.8, 1), [1, \infty)$
Fig.8(d)	$(-\infty, -1), [-1, -0.8), [-0.8, -0.6), [-0.6, -0.4), [-0.4, -0.2), [-0.2, 0), [0, 0.2), [0.2, 0.4), [0.4, 0.6), [0.6, 0.8), [0.8, 1), [1, 1.2), [1.2, \infty)$
Fig.9(d)	$(-\infty, -0.3), [-0.3, -0.2), [-0.2, -0.1), [-0.1, 0), [0, 0.1), [0.1, 0.2), [0.2, 0.3), [0.3, \infty)$

of test samples is 2000 (statistical bar), else 50. Our simulation results can be found in Figure 7, Figure 8 and Figure 9, respectively. For sake of clarity, the division intervals of the statistical bars in panel (d) of Figures 7,8,9 are shown in Table 1.

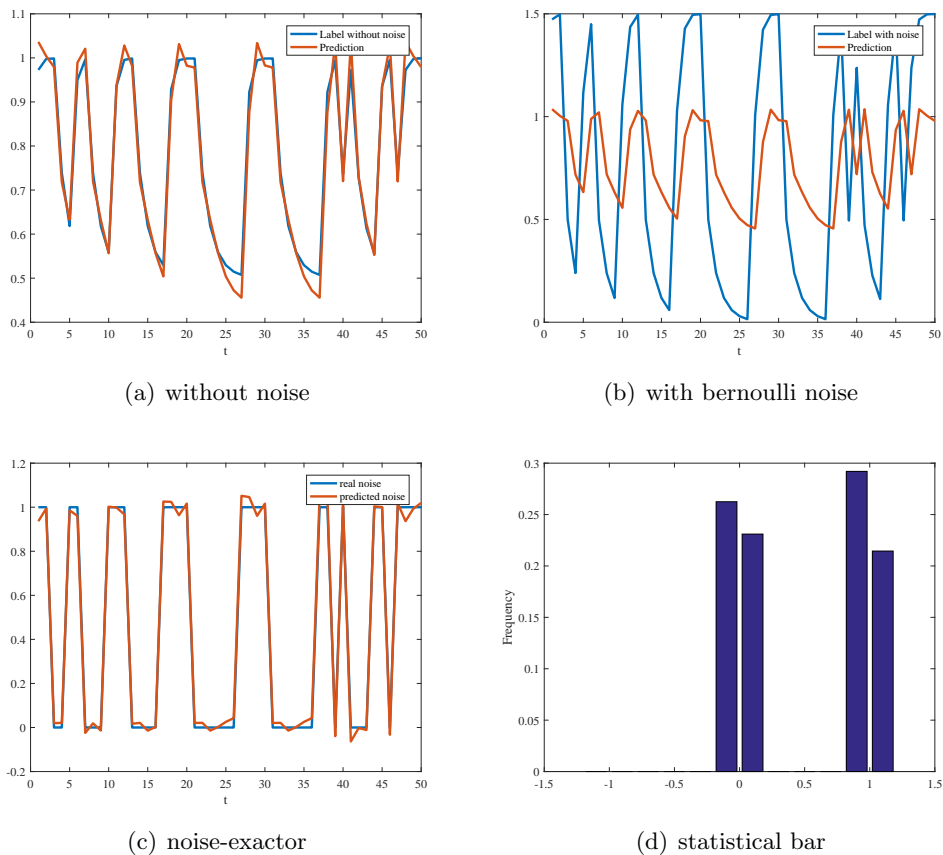


Figure 7: Noise extractor for Nyström regularization ($\varepsilon_t \sim \mathcal{B}(0.5)$).

It can be found in panel (a) of Figures 7,8,9 that the proposed Nyström regularization with sub-sampling ratio 0.01 can precisely catch the deterministic part of the time series, i.e., yielding an estimator that can approximate $x_t = 0.5 \sin(x_{t-1})$ very well, even though the training samples are contaminated by different noises. However, panel (b) of Figures 7,8,9 show that the proposed algorithm is incapable of catching the random noise, making

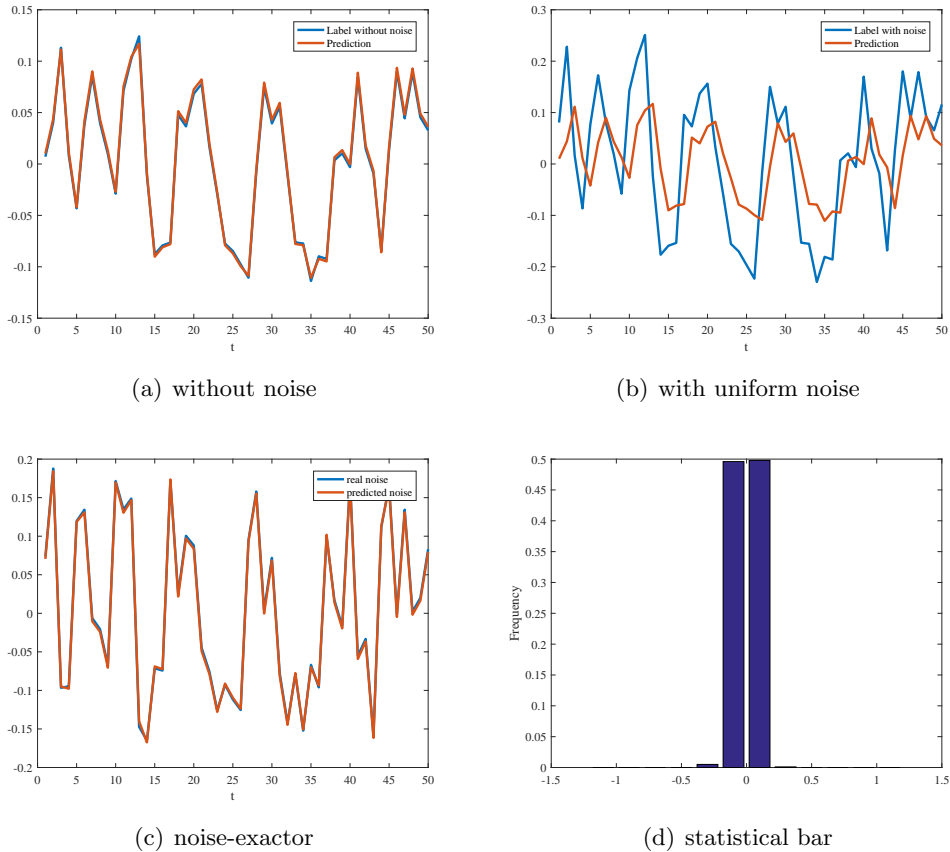
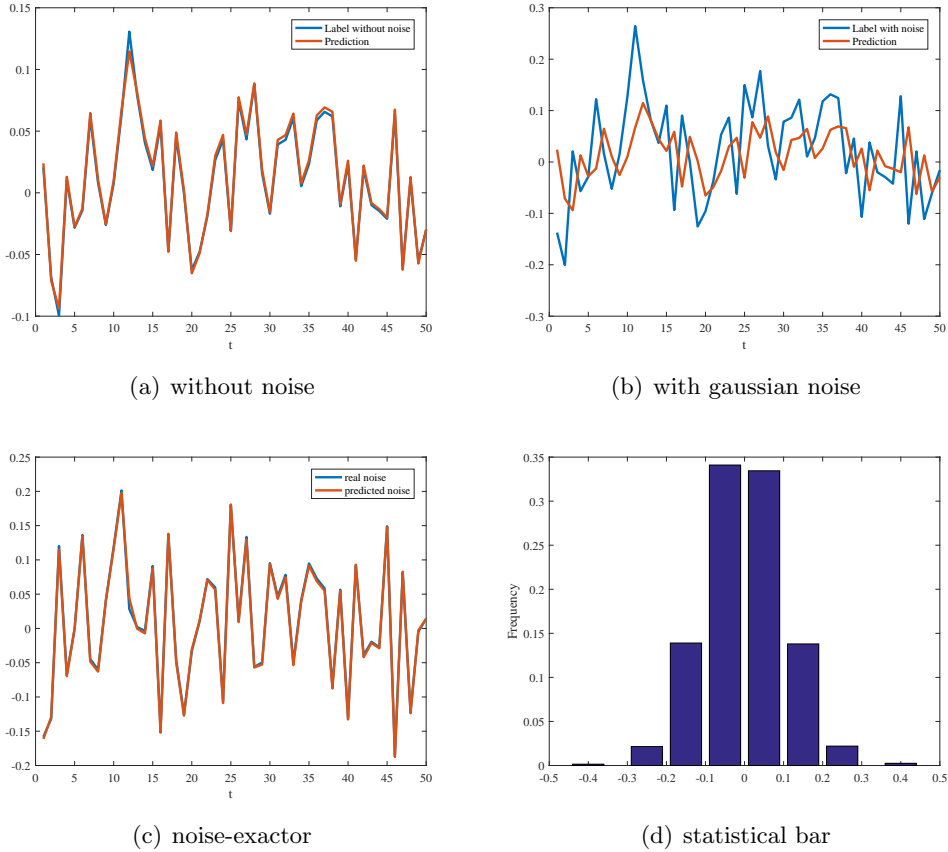


Figure 8: Noise extractor for Nyström regularization ($\varepsilon_t \sim \mathcal{U}[-0.2, 0.2]$).

the derived estimator lag heavily behind the real time series. In panel (c) of Figures 7,8,9, we compare the real noise with $x_{t+1} - f_{D, D_j^*, \lambda}(x_t)$ to regard Nyström regularization as a noise-extractor. It is exhibited that using Nyström regularization is easy to mimic the random noise part of the time series. From panel (d) of Figures 7,8,9, we find that the distribution of the extracted noise via Nyström regularization is similar to the real noise. All these findings yield the following two conclusions: 1) Nyström regularization with sequential sub-sampling can successfully capture the trend information (deterministic part) of time series. 2) Nyström regularization with sequential sub-sampling has strong ability to extract the noise of data (not only the value but also the distribution). Although the noise-extractor property of Nyström regularization seems trivial in toy simulation, it is extremely important in real-world time series forecasting, where the deterministic part and random noise part cannot be divided as those in the toy simulations.

Figure 9: Noise extractor for Nyström regularization ($\varepsilon_t \sim \mathcal{N}(0, 0.1^2)$).

5.2 Real world applications

In this part, we apply the proposed Nyström regularization with sequential sub-sampling to three real-world time series forecasting data, WTI data, BTC data and Western Australia Weather data, to show its learning performance.

WTI data: In the industrial society, oil almost dominates the production of the whole industry. Oil and additional industries account for a high proportion in the national economy. A deep understanding of the international crude oil price fluctuation mechanism and improving the accuracy of international crude oil price forecasts are of great significance to economic development, enterprise production operations and investment. WTI (Western Texas Intermediate) Spot Prices from EIA U.S. (Energy Information Administration) is suitable for refining gasoline, diesel, thermal fuel oil and aircraft fuel, etc., and can increase the output value of refineries. WTI is an important part of the international energy pricing system and has become the benchmark for global crude oil pricing. Meanwhile, WTI has become one of the two most market-oriented crude oils in the world. The data (<https://datahub.io/core/oil-prices>) is daily recorded and from January 2, 1986 to August 31, 2020. Since it is difficult to check the smoothness of the regression function, we use a

less-smoothed kernel function

$$K(x, x') = 1 + \min(x, x')$$

in real-world data experiments.

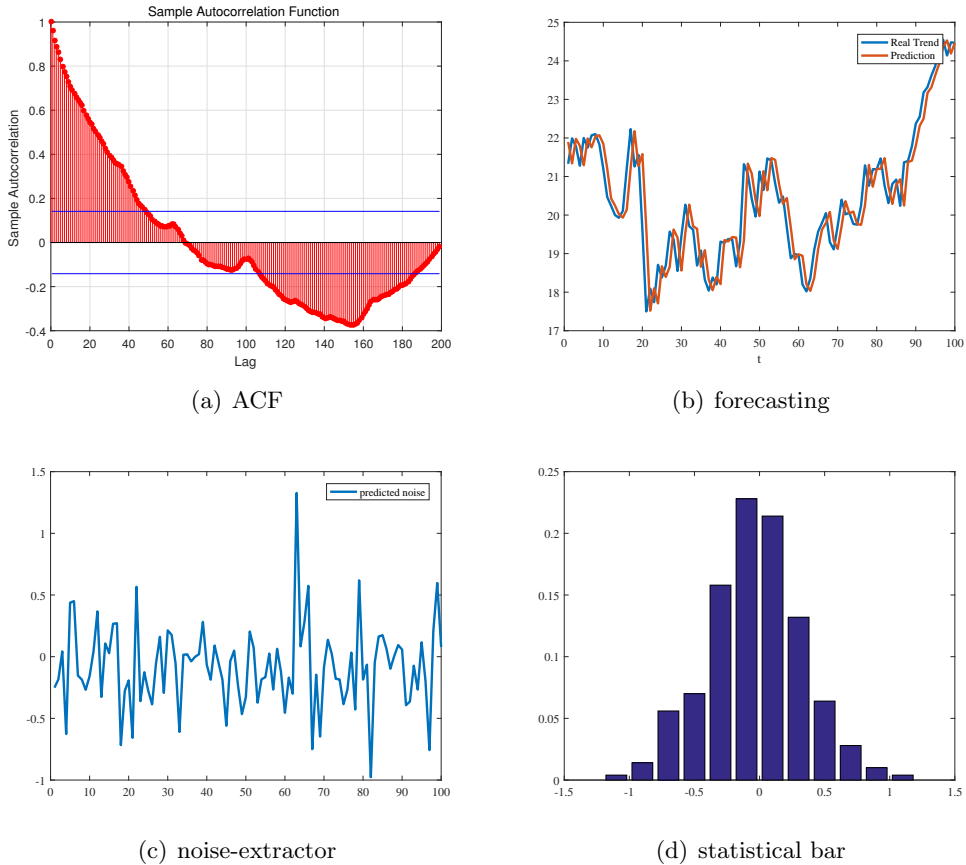


Figure 10: Nyström regularization on WTI data

Before demonstrating the availability of Nyström regularization to WTI data, we present the autocorrelation function (ACF) of the time series to verify their mixing property, just as Figure 10 (a) purports to show. Although it is difficult to numerically compute the τ -mixing coefficient of WTI data, Figure 10 (a) implies that the dependence decreases when the difference of time increases. We employ Nyström regularization on WTI data with the size of samples and regularization parameters as:

- Forecasting: The number of training samples: 4000, test samples: 100, sampling ratio is 0.1, $\lambda = \frac{1}{1000}$.
- Statistical bar: The number of training samples: 3000, test samples: 500, sampling ratio is 0.1, $\lambda = \frac{4}{3000}$.
- Noise-extractor: The number of training samples: 3000, test samples: 100, sampling ratio is 0.1, $\lambda = \frac{4}{3000}$.

The experimental results are reported in Figure 10.

There are three interesting findings illustrated by Figure 10: 1) It can be found in Figure 10 (b) that the proposed Nyström regularization can roughly mimic the trend of WTI data within a relatively small prediction error. However, due to the existence of noise, the predictor lags slightly behind the real time series, making the proposed algorithm a bit wise-after-the-event; 2) As discussed in toy simulations, Nyström regularization can be used as a noise-extractor in practice. According to this property, we extract the noise of WTI data in Figure 10 (b). Furthermore, Figure 10 (d) shows that the random noise part of WTI data obeys the normal distribution; 3) Comparing Figure 10 (b) and Figure 10 (c), it is obvious that the random noise part is far smaller than the deterministic part, exhibiting a noise/signal ratio to be less than 1/15, showing that exploiting the deterministic part of WTI data is important in practice. All these demonstrated the excellent performance of Nyström regularization in WTI data forecasting.

BITCOIN(BTC) data: Bitcoin (BTC) is a decentralized digital currency that can be transformed from user to user on a peer-to-peer BTC network by using the blockchain technique. After its release in 2009, its price increases from about 3\$ to 60000\$ in 2021. Such a surge in price attracts more and more investors' attention. In particular, there are more than 10 million active accounts in October, 2021. The problem is that, however, the price of BTC changes dramatically even for every minute, although the trends of price are roughly increasing. Therefore, it is highly desired to forecast the real trend of price of BTC.

The BTC data collected via “<https://www.kaggle.com/prasoonkottarathil/btcinUSD>” record the price of BTC from September 17, 2014 to April 9, 2020 within each minute. Therefore, there are at least 500000 samples in the data set and the classical kernel methods and neural networks fail to tackle this massive time series. In this part, we focus on applying the proposed Nyström regularization on BTC data to mimic the trend of price and extract the noise. The data and parameters of Nyström regularization are described as follows:

- Forecasting: the number of training samples: 500000, test samples: 100, sampling ratio is 0.001, $\lambda = \frac{1}{500000}$.
- Noise-extractor: the number of training samples: 20000, test samples: 100, sampling ratio is 0.001, $\lambda = \frac{1}{20000}$.
- Statistical bar: the number of training samples: 20000, test samples: 3000, sampling ratio is 0.001, $\lambda = \frac{1}{20000}$.

The numerical results are shown in Figure 11. As shown in Figure 11 (a), the time series concerning BTC price exhibits certain mixing property in the sense that the ACF curve decreases with respect to the difference of time. Like WTI data, it can be found in Figure 11 (b) that the trends predicted by Nyström regularization can approximate the real trend within a relative small error, although there also exists a bit delay in prediction. As discussed above, the reason of delay is the existence of random noise that can be learned by BTC data. We then conduct our experiment on pursuing the distribution of such random noise. Figure 11 (c) and (d) show that the random noise part of BTC data behaves roughly as a normal distribution with standard deviation smaller than 30\$. In summary, even with extremely small sub-sampling ratio (about 0.001), Nyström regularization can still predict

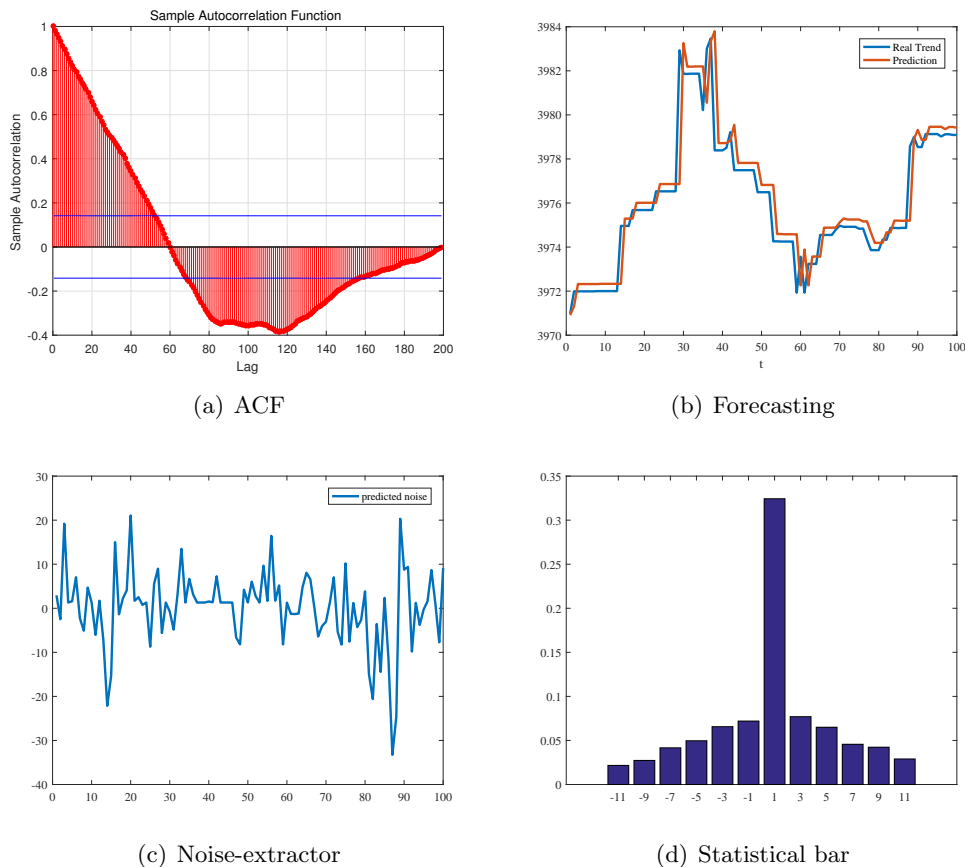


Figure 11: Nyström regularization for BTC min data

the trend well and is capable of describing the distribution of noise for Bitcoin price. All these show the power of the suggested Nyström regularization with sequential sub-sampling in forecasting time series and extracting its noise information.

Western Australia Weather data: Weather forecasting can predict the meteorological changes in a certain period of time, such as when low temperature or drought will occur, so that effective measures can be taken to prevent and reduce disasters. And the weather forecast also helps in other areas of life, such as municipal transportation and travel planning.

Western Australia Weather data (<https://www.kaggle.com/datasets/sveneschlbeck/west-australia-weather-1944-2016>) record Western Australia’s daily average temperatures from June 3, 1944 to December 31, 2016. In this part, We focus on applying the proposed Nyström regularization to weather data to predict daily average temperature. The data and parameters are described as follows:

- Forecasting: the number of training samples: 26400, test samples: 100, sampling ratio is 0.01, $\lambda = \frac{1}{2640}$.

- Noise-extractor: the number of training samples: 20000, test samples: 100, sampling ratio is 0.01, $\lambda = \frac{1}{2000}$.
- Statistical bar: the number of training samples: 20000, test samples: 500, sampling ratio is 0.01, $\lambda = \frac{1}{2000}$.

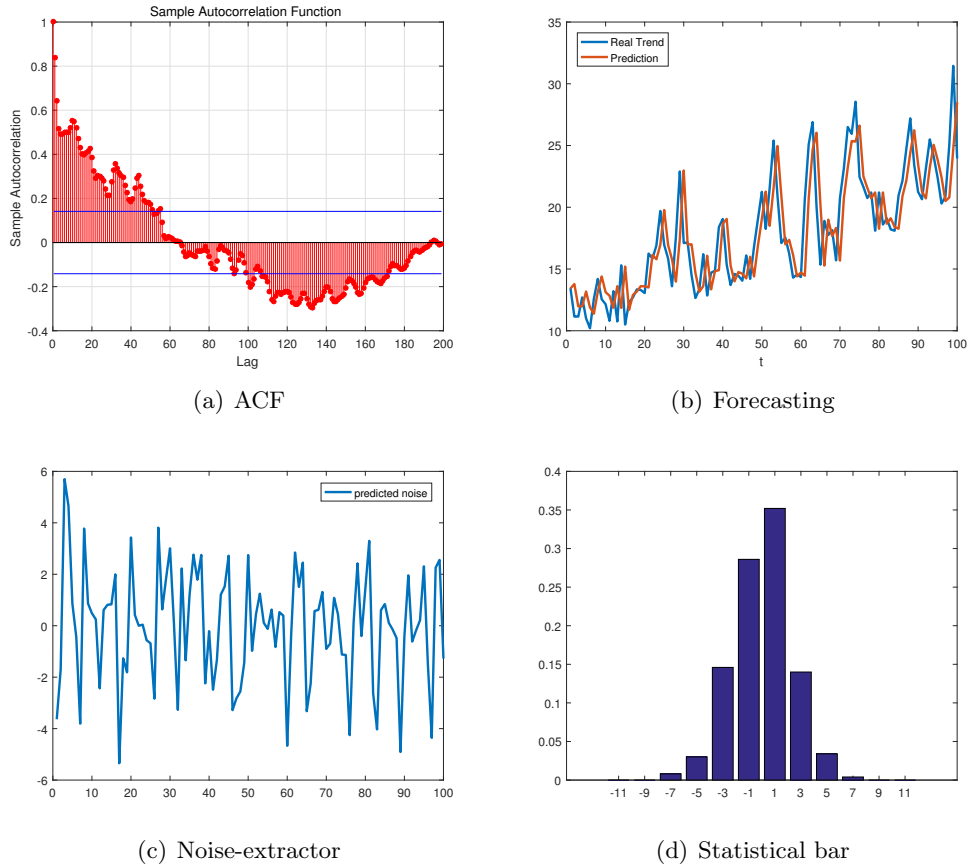


Figure 12: Nyström regularization for Western Australia Weather data

The numerical results are shown in Figure 12. Similar to previous experimental results. As shown in Figure 12 (a), the daily temperature series exhibits some mixing property. Based on this, we show the predictive ability of the Nyström regularization for temperature data, which can be found in Figure 12 (b) that the predicted trend can approximate the true trend within a small error. Figure 12 (c) and (d) show that the random noise part of the temperature data approximately follows a normal distribution. In summary, Nyström regularization with sequential sub-sampling can predict temperature trends well and can capture the information of noise.

6. Proofs

We divide our proof into four steps: error decomposition, error estimates based on operator differences, estimate for operator differences and final error analysis.

6.1 Error decomposition

We conduct our analysis in the popular integral operator approach developed in (Smale and Zhou, 2007; Rudi et al., 2015; Lin et al., 2017). Let $S_D : \mathcal{H}_K \rightarrow \mathbb{R}^n$ be the sampling operator (Smale and Zhou, 2007) defined by

$$S_D f := (\langle f, K_x \rangle_K)_{(x,y) \in D} = (f(x))_{(x,y) \in D}. \quad (28)$$

Its scaled adjoint $S_D^T : \mathbb{R}^n \rightarrow \mathcal{H}_K$ is given by

$$S_D^T \mathbf{c} := \frac{1}{n} \sum_{i=1}^n c_i K_{x_i}, \quad \mathbf{c} := (c_1, c_2, \dots, c_n)^T \in \mathbb{R}^n.$$

Define

$$L_{K,D} f := S_D^T S_D f = \frac{1}{n} \sum_{(x,y) \in D} f(x) K_x.$$

Denote by D_m an arbitrary sequential sub-sampling of D with cardinality m . Let P_{D_m} be the projection from \mathcal{H}_K to \mathcal{H}_{D_m} , where \mathcal{H}_{D_m} is defined by (11). Then for arbitrary $v > 0$, there holds

$$(I - P_{D_m})^v = I - P_{D_m}. \quad (29)$$

Let $\mathcal{S}_D : \mathcal{H}_{D_m} \rightarrow \mathbb{R}^n$ be the sampling operator defined by (28) such that the range of its adjoint operator \mathcal{S}_D^T is exactly \mathcal{H}_{D_m} . Let $U\Sigma V^T$ be the SVD of \mathcal{S}_D . Then we have

$$V_{D_m}^T V_{D_m} = I, \quad V_{D_m} V_{D_m}^T = P_{D_m}. \quad (30)$$

Write

$$g_{D_m, \lambda}(L_{K,D}) := V_{D_m} (V_{D_m}^T L_{K,D} V_{D_m} + \lambda I)^{-1} V_{D_m}^T. \quad (31)$$

Then it can be found in (Rudi et al., 2015) that

$$f_{D, D_m, \lambda} = g_{D_m, \lambda}(L_{K,D}) \mathcal{S}_D^T y_D. \quad (32)$$

Therefore, (12) is similar as the spectral-type algorithms (Lo Gerfo et al., 2008). Under this circumstance, the property of $g_{D_m, \lambda}(L_{K,D})$ plays a crucial role in bounding the generalization error of $f_{D, D_m, \lambda}$. For an arbitrary bounded linear operator B , it follows from (30) and (31) that

$$g_{D_m, \lambda}(L_{K,D})(L_{K,D} + \lambda I) V_{D_m} B V_{D_m}^T = V_{D_m} B V_{D_m}^T. \quad (33)$$

Inserting $B = (V_{D_m}^T L_{K,D} V_{D_m} + \lambda I)^{-1}$ into (33), we obtain

$$\begin{aligned} & \| (L_{K,D} + \lambda I)^{1/2} g_{D_m, \lambda}(L_{K,D})(L_{K,D} + \lambda I)^{1/2} \|^2 \\ &= \| (L_{K,D} + \lambda I)^{1/2} g_{D_m, \lambda}(L_{K,D})(L_{K,D} + \lambda I) g_{D_m, \lambda}(L_{K,D})(L_{K,D} + \lambda I)^{1/2} \|^2 \\ &= \| (L_{K,D} + \lambda I)^{1/2} g_{D_m, \lambda}(L_{K,D})(L_{K,D} + \lambda I)^{1/2} \|^2, \end{aligned}$$

which yields

$$\|(L_{K,D} + \lambda I)^{1/2} g_{D_m, \lambda} (L_{K,D}) (L_{K,D} + \lambda I)^{1/2}\| = 1, \quad (34)$$

where $\|\cdot\|$ denotes the operator norm.

Define

$$f_{D, D_m, \lambda}^* := g_{D_m, \lambda} (L_{K,D}) L_{K,D} f_\rho$$

be the noise-free version of $f_{D, D_m, \lambda}$. The triangle inequality shows

$$\|f_{D, D_m, \lambda} - f_\rho\|_\rho \leq \|f_{D, D_m, \lambda} - f_{D, D_m, \lambda}^*\|_\rho + \|f_{D, D_m, \lambda}^* - f_\rho\|_\rho. \quad (35)$$

Since $\|f_{D, D_m, \lambda} - f_{D, D_m, \lambda}^*\|_\rho$ describes the noise of samples, it is named as the sample error in (Rudi et al., 2015). From (31), we get

$$\begin{aligned} \|f_{D, D_m, \lambda}^* - f_\rho\|_\rho &\leq \|(g_{D_m, \lambda} (L_{K,D}) L_{K,D} - I)(I - P_{D_m}) f_\rho\|_\rho \\ &\quad + \|(g_{D_m, \lambda} (L_{K,D}) L_{K,D} - I) P_{D_m} f_\rho\|_\rho. \end{aligned}$$

It is easy to see that the second term in the righthand side of the above estimate is similar as the classical approximation error (Guo et al., 2017). The first term involves an additional term $I - P_{D_m}$ to show the limitation of sub-sampling and thus is named as the computational error in (Rudi et al., 2015). Plugging the above inequality into (35), we derive the following error decomposition.

Lemma 11 *Let $f_{D, D_m, \lambda}$ be defined by (12), we have*

$$\|f_{D, D_m, \lambda} - f_\rho\|_\rho \leq \mathcal{A}(D, \lambda, m) + \mathcal{S}(D, \lambda, m) + \mathcal{C}(D, \lambda, m), \quad (36)$$

where

$$\begin{aligned} \mathcal{A}(D, \lambda, m) &= \|(g_{D_m, \lambda} (L_{K,D}) L_{K,D} - I) P_{D_m} f_\rho\|_\rho, \\ \mathcal{S}(D, \lambda, m) &= \|f_{D, D_m, \lambda} - f_{D, D_m, \lambda}^*\|_\rho, \\ \mathcal{C}(D, \lambda, m) &= \|(g_{D_m, \lambda} (L_{K,D}) L_{K,D} - I)(I - P_{D_m}) f_\rho\|_\rho \end{aligned}$$

are called the approximation error, sample error and computational error, respectively.

6.2 Error estimates based on operator differences

In this part, we quantify $\|f_{D, D_m, \lambda} - f_\rho\|_\rho$ via differences between operators $L_{K,D}$ and L_K and functions $S_{D_m}^T y_D$ and $L_{K,D} f_\rho$. In particular, we use the products of operators

$$\begin{aligned} \mathcal{Q}_{D, \lambda} &:= \|(L_K + \lambda I)(L_{K,D} + \lambda I)^{-1}\|, \\ \mathcal{Q}_{D, \lambda}^* &:= \|(L_K + \lambda I)^{-1}(L_{K,D} + \lambda I)\|, \end{aligned}$$

and the difference of operators

$$\mathcal{R}_{D, \lambda} := \left\| (L_K + \lambda I)^{-1/2} (L_{K,D} - L_K) \right\|$$

to quantify the similarity of L_K and $L_{K,D}$, while utilize

$$\mathcal{P}_{D, \lambda} := \left\| (L_K + \lambda I)^{-1/2} (L_{K,D} f_\rho - S_{D_m}^T y_D) \right\|_K$$

to quantify the difference between $L_{K,D}f_\rho$ and $S_D^T y_D$. Our main tools are the following two lemmas that can be found in (Rudi et al., 2015, Proposition 3) and (Rudi et al., 2015, Proposition 6), respectively.

Lemma 12 *Let $\mathcal{H}, \mathcal{K}, \mathcal{F}$ be three separable Hilbert spaces. Let $Z : \mathcal{H} \rightarrow \mathcal{K}$ be a bounded linear operator and P be a projection operator on \mathcal{H} such that $\text{range}P = \overline{\text{range}Z^T}$. Then for any bounded linear operator $F : \mathcal{F} \rightarrow \mathcal{H}$ and any $\lambda > 0$ we have*

$$\|(I - P)F\| \leq \lambda^{1/2} \|(Z^T Z + \lambda I)^{-1/2} F\|.$$

Lemma 13 *Let \mathcal{H}, \mathcal{K} be two separable Hilbert spaces, $A : \mathcal{H} \rightarrow \mathcal{H}$ be a positive linear operator, $V : \mathcal{H} \rightarrow \mathcal{K}$ be a partial isometry and $B : \mathcal{K} \rightarrow \mathcal{K}$ be a bounded operator. Then for all $0 \leq r^*, s^* \leq 1/2$, there holds*

$$\|A^{r^*} V B V^T A^{s^*}\| \leq \|(V^T A V)^{r^*} B (V^T A V)^{s^*}\|.$$

With the help of above lemmas and the important properties of $g_{D_m, \lambda}$ in (33) and (34), we derive the following error estimate for $\|f_{D, D_m, \lambda} - f_\rho\|_\rho$.

Proposition 14 *If (15) holds with $1/2 \leq r \leq 1$, then we have*

$$\begin{aligned} \|f_{D, D_m, \lambda} - f_\rho\|_\rho &\leq \mathcal{Q}_{D, \lambda} \mathcal{P}_{D, \lambda} + \lambda^r \mathcal{Q}_{D, \lambda}^r \|h_\rho\|_\rho \\ &\quad + (\mathcal{Q}_{D, \lambda}^{\frac{1}{2}} (\mathcal{Q}_{D, \lambda}^*)^{\frac{1}{2}} + 1) \lambda^r \mathcal{Q}_{D_m, \lambda}^r \|h_\rho\|_\rho. \end{aligned} \quad (37)$$

Proof According to (36), it suffices to bound $\mathcal{A}(D, \lambda, m), \mathcal{S}(D, \lambda, m), \mathcal{C}(D, \lambda, m)$ respectively. We at first use (33) and (34) to bound $\mathcal{S}(D, \lambda, m)$. Due to the well known Codes inequality (Bathis, 1997)

$$\|A^u B^u\| \leq \|AB\|^u, \quad 0 < u \leq 1 \quad (38)$$

for arbitrary positive operators A, B , we have

$$\|(L_K + \lambda I)^{1/2} (L_{K, D} + \lambda I)^{-1/2}\| \leq \mathcal{Q}_{D, \lambda}^{1/2}. \quad (39)$$

Then, it follows from (34) and $\|f\|_\rho = \|L_K^{1/2} f\|_K$ for any $f \in L_{\rho_X}^2$ that

$$\begin{aligned} \mathcal{S}(D, \lambda, m) &= \|g_{D_m, \lambda}(L_{K, D})(S_D^T y_D - L_{K, D} f_\rho)\|_\rho \\ &= \|L_K^{1/2} g_{D_m, \lambda}(L_{K, D})(S_D^T y_D - L_{K, D} f_\rho)\|_K \\ &\leq \|(L_K + \lambda I)^{1/2} (L_{K, D} + \lambda I)^{-1/2}\|^2 \|(L_{K, D} + \lambda I)^{1/2} g_{D_m, \lambda}(L_{K, D})(L_{K, D} + \lambda I)^{1/2}\| \\ &\quad \times \|(L_K + \lambda I)^{-1/2} (S_D^T y_D - L_{K, D} f_\rho)\|_K \\ &\leq \mathcal{Q}_{D, \lambda} \mathcal{P}_{D, \lambda}. \end{aligned} \quad (40)$$

We then turn to bounding $\mathcal{A}(D, \lambda, m)$. It follows from (33) with $B = I$ that

$$P_{D_m} = g_{D_m, \lambda}(L_{K, D})(L_{K, D} + \lambda I)P_{D_m},$$

which together with the definition of $\mathcal{A}(D, \lambda, m)$

$$\mathcal{A}(D, \lambda, m) = \lambda \|g_{D_m, \lambda}(L_{K, D}) P_{D_m} f_\rho\|_\rho$$

Since $L_{K, D}$ is a positive operator, we have $\|(V_{D_m}^T (L_{K, D} + \lambda I) V_{D_m})^{r-1}\| \leq \lambda^{r-1}$ for $r \leq 1$. Noting further (30) and (38), we get from Lemma 13 with $A = (L_{K, D} + \lambda I)$, $V = V_{D_m}$, $B = (V_{D_m}^T L_{K, D} V_{D_m} + \lambda I)^{-1}$, $r^* = 1/2$ and $s^* = r - 1/2$ that

$$\begin{aligned} \mathcal{A}(D, \lambda, m) &\leq \lambda \|L_K^{1/2} g_{D_m, \lambda}(L_{K, D}) V_{D_m} V_{D_m}^T L_K^{r-1/2}\| \|h_\rho\|_\rho \\ &\leq \lambda \mathcal{Q}_{D, \lambda}^r \|h_\rho\|_\rho \|(L_{K, D} + \lambda I)^{1/2} g_{D_m, \lambda}(L_{K, D}) V_{D_m} V_{D_m}^T (L_{K, D} + \lambda I)^{r-1/2}\| \\ &\leq \lambda \mathcal{Q}_{D, \lambda}^r \|(V_{D_m}^T (L_{K, D} + \lambda I) V_{D_m})^{1/2} (V_{D_m}^T (L_{K, D} + \lambda I) V_{D_m})^{-1} \\ &\quad (V_{D_m}^T (L_{K, D} + \lambda I) V_{D_m})^{r-1/2}\| \|h_\rho\|_\rho \\ &\leq \lambda \mathcal{Q}_{D, \lambda}^r \|(V_{D_m}^T (L_{K, D} + \lambda I) V_{D_m})^{r-1}\| \|h_\rho\|_\rho \\ &\leq \lambda^r \mathcal{Q}_{D, \lambda}^r \|h_\rho\|_\rho. \end{aligned} \tag{41}$$

Finally, we aim at bounding $\mathcal{C}(D, \lambda, m)$. Due to Lemma 12 and (39), we have

$$\begin{aligned} &\|(I - P_{D_m})(L_K + \lambda I)^{1/2}\| \\ &\leq \lambda^{1/2} \|(L_{K, D_m} + \lambda I)^{-1/2} (L_K + \lambda I)^{1/2}\| \leq \lambda^{1/2} \mathcal{Q}_{D_m, \lambda}^{1/2}. \end{aligned}$$

Then, it follows from (38), (34) and (29) with $\tau = 2r$ that

$$\begin{aligned} \mathcal{C}(D, \lambda, m) &\leq \|L_K^{1/2} g_{D_m, \lambda}(L_{K, D}) L_{K, D} (I - P_{D_m}) L_K^{r-1/2}\| \|h_\rho\|_\rho \\ &+ \|L_K^{1/2} (I - P_{D_m}) L_K^{r-1/2}\| \|h_\rho\|_\rho \\ &\leq \mathcal{Q}_{D, \lambda}^{1/2} (\mathcal{Q}_{D, \lambda}^*)^{1/2} \|(L_{K, D} + \lambda I)^{1/2} g_{D_m, \lambda}(L_{K, D}) (L_{K, D} + \lambda I)^{1/2}\| \\ &\times \|(L_K + \lambda I)^{1/2} (I - P_{D_m})^{2r} L_K^{r-1/2}\| \|h_\rho\|_\rho \\ &+ \|(L_K + \lambda I)^{1/2} (I - P_{D_m})^{2r} L_K^{r-1/2}\| \|h_\rho\|_\rho \\ &\leq (\mathcal{Q}_{D, \lambda}^{1/2} (\mathcal{Q}_{D, \lambda}^*)^{1/2} + 1) \|h_\rho\|_\rho \|(L_K + \lambda I)^{1/2} (I - P_{D_m})\| \|(I - P_{D_m})^{2r-1} L_K^{r-1/2}\| \\ &\leq (\mathcal{Q}_{D, \lambda}^{1/2} (\mathcal{Q}_{D, \lambda}^*)^{1/2} + 1) \|h_\rho\|_\rho \lambda^r \mathcal{Q}_{D_m, \lambda}^r. \end{aligned} \tag{42}$$

Inserting (41), (40) and (42) into (36), we get (37). This finishes the proof of Proposition 14. \blacksquare

6.3 Bounds for operator differences

In this part, we focus on deriving tight bounds for $\mathcal{Q}_{D, \lambda}$, $\mathcal{Q}_{D, \lambda}^*$ and $\mathcal{P}_{D, \lambda}$ when D is a τ -mixing sequences. Our main tool is the Bernstein-type inequality for Banach-valued sums in (Blanchard and Zadorozhnyi, 2019) and the second-order decomposition for operator differences in (Lin et al., 2017; Guo et al., 2017). Under Assumptions 2 and 4, define

$$n_\gamma := \begin{cases} \frac{n}{2(1 \vee \log(c_1^* n))^{1/\gamma_0}}, & \text{if } D \text{ satisfies (8),} \\ (c_2^* \lambda \mathcal{N}(\lambda))^{\frac{1}{2\gamma_1+1}} n^{\frac{2\gamma_1}{2\gamma_1+1}}, & \text{if } D \text{ satisfies (9),} \end{cases} \tag{43}$$

where

$$\begin{aligned} c_1^* &:= c_0 b_0 \max\{\mathcal{K}, 3(1 + \mathcal{K}M)/(2M)\}, \\ c_2^* &:= 2^{-2\gamma_1} (\min\{1.5M/(1 + \mathcal{K}M), 1/(2\mathcal{K})\}/c_1)^2. \end{aligned}$$

The following lemma can be found in (Blanchard and Zadorozhnyi, 2019, Lemma 4.1).

Lemma 15 *Let $0 < \delta \leq 1/2$. Under Assumptions 1 and 2, then*

$$\mathcal{P}_{D,\lambda} \leq 42M\mathcal{B}(n_\gamma, \lambda) \log \frac{2}{\delta}, \quad (44)$$

$$\mathcal{R}_{D,\lambda} \leq 42\mathcal{B}(n_\gamma, \lambda) \log \frac{2}{\delta}, \quad (45)$$

hold simultaneously with confidence $1 - \delta$, where

$$\mathcal{B}(n_\gamma, \lambda) := \frac{\sqrt{\mathcal{N}(\lambda)}}{\sqrt{n_\gamma}} + \frac{1}{n_\gamma \sqrt{\lambda}}. \quad (46)$$

Then, we use Lemma 15 and approaches in (Guo et al., 2017) to bound operator products $\mathcal{Q}_{D,\lambda}$ and $\mathcal{Q}_{D,\lambda}^*$.

Lemma 16 *Let $0 < \delta \leq 1/2$. Under Assumptions 1 and 2, then*

$$\mathcal{Q}_{D,\lambda} \leq 3528\lambda^{-1}\mathcal{B}^2(n_\gamma, \lambda) \log^2 \frac{2}{\delta} + 2, \quad (47)$$

$$\mathcal{Q}_{D,\lambda}^* \leq 42\lambda^{-1/2}\mathcal{B}(n_\gamma, \lambda) \log \frac{2}{\delta} + 1, \quad (48)$$

hold with confidence $1 - \delta$.

Proof For invertible positive operators A, B , we have

$$A^{-1}B = (A^{-1} - B^{-1})B + I = A^{-1}(B - A) + I \quad (49)$$

$$\begin{aligned} &= (A^{-1} - B^{-1})(B - A) + B^{-1}(B - A) + I \\ &= A^{-1}(B - A)B^{-1}(B - A) + B^{-1}(B - A) + I. \end{aligned} \quad (50)$$

We first use (49) with $A = (L_K + \lambda I)$ and $B = (L_{K,D} + \lambda I)$ to bound $\mathcal{Q}_{D,\lambda}^*$. It follows from (45) that with confidence $1 - \delta$, that holds

$$\begin{aligned} \mathcal{Q}_{D,\lambda}^* &= \|(L_K + \lambda I)^{-1}(L_{K,D} - L_K)\| + 1 \leq \lambda^{-1/2}\mathcal{R}_{D,\lambda} + 1 \\ &\leq 42\lambda^{-1/2}\mathcal{B}(n_\gamma, \lambda) \log \frac{2}{\delta} + 1. \end{aligned}$$

Therefore (48) holds. Then, due to (50) with $A = (L_{K,D} + \lambda I)$ and $B = (L_K + \lambda I)$, we have from (45) again that with confidence $1 - \delta$, there holds

$$\begin{aligned} \mathcal{Q}_{D,\lambda} &\leq \|(L_{K,D} + \lambda I)^{-1}(L_K - L_{K,D})(L_K + \lambda I)^{-1}(L_K - L_{K,D})\| \\ &+ \|(L_K + \lambda I)^{-1}(L_K - L_{K,D})\| + 1 \\ &\leq \lambda^{-1}\|(L_K + \lambda I)^{-1/2}(L_K - L_{K,D})\|^2 + \lambda^{-1/2}\|(L_K + \lambda I)^{-1/2}(L_K - L_{K,D})\| + 1 \\ &\leq 2(\lambda^{-1}\mathcal{R}_{D,\lambda}^2 + 1) \leq 3528\lambda^{-1}\mathcal{B}^2(n_\gamma, \lambda) \log^2 \frac{2}{\delta} + 2. \end{aligned}$$

This completes the proof of Lemma 16. ■

6.4 Error analysis

Based on Lemma 15, Lemma 16 and Proposition 14, we are in a position to prove main results in Section 3. To this end, we present a more general theorem as follows.

Theorem 17 *Let $0 < \delta \leq 1/2$. Under Assumptions 2-5 with $\frac{1}{2} \leq r \leq 1$, for any $j \in [1, n - m + 1]$, with confidence $1 - \delta$, there holds*

$$\|f_{D, D_m, \lambda} - f_\rho\|_\rho \leq \bar{C}(\lambda^{-1} \mathcal{B}^2(n_\gamma, \lambda) + 1) (\mathcal{B}(n_\gamma, \lambda) + \mathcal{B}^{2r}(m_\gamma, \lambda) + \lambda^r) \log^4 \frac{2}{\delta}, \quad (51)$$

where

$$m_\gamma := \begin{cases} \frac{m}{2(1 \vee \log(c_1^* m))^{1/\gamma_0}}, & \text{if } D \text{ satisfies (8),} \\ (c_2^* \lambda \mathcal{N}(\lambda))^{\frac{1}{2\gamma_1+1}} m^{\frac{2\gamma_1}{2\gamma_1+1}}, & \text{if } D \text{ satisfies (9),} \end{cases} \quad (52)$$

and \bar{C} is a constant independent of m, n, λ, j or δ .

Proof Due to Lemma 15 and Lemma 16, we have from $(a+b)^u \leq 2^u(a^u + b^u)$ for $a, b, u \geq 0$ that

$$\begin{aligned} \mathcal{P}_{D, \lambda} \mathcal{Q}_{D, \lambda} &\leq C_1 \mathcal{B}(n_\gamma, \lambda) (42^2 \lambda^{-1} \mathcal{B}^2(n_\gamma, \lambda) + 1) \log^3 \frac{2}{\delta}, \\ \mathcal{Q}_{D, \lambda}^{\frac{1}{2}} (\mathcal{Q}_{D, \lambda}^*)^{\frac{1}{2}} &\leq 2(42 \lambda^{-1/2} \mathcal{B}(n_\gamma, \lambda) + 1)^{3/2} \log^2 \frac{2}{\delta}, \\ \mathcal{Q}_{D, \lambda}^r &\leq (3528^r \lambda^{-r} \mathcal{B}^{2r}(n_\gamma, \lambda) \log^{2r} \frac{2}{\delta} + 2^r) 2^r, \\ \mathcal{Q}_{D_m, \lambda}^r &\leq (3528^r \lambda^{-r} \mathcal{B}^{2r}(m_\gamma, \lambda) \log^{2r} \frac{2}{\delta} + 2^r) 2^r, \end{aligned}$$

where $C_1 := 84M$. Plugging the above estimates into (37), we have

$$\begin{aligned} \|f_{D, D_m, \lambda} - f_\rho\|_\rho &\leq C_1 \mathcal{B}(n_\gamma, \lambda) (42^2 \lambda^{-1} \mathcal{B}^2(n_\gamma, \lambda) + 1) \log^3 \frac{2}{\delta} \\ &+ \lambda^r \|h_\rho\|_\rho \left(3528^r \lambda^{-r} \mathcal{B}^{2r}(n_\gamma, \lambda) \log^{2r} \frac{2}{\delta} + 2^r \right) 2^r, \\ &+ \lambda^r \|h_\rho\|_\rho \left(3528^r \lambda^{-r} \mathcal{B}^{2r}(m_\gamma, \lambda) \log^{2r} \frac{2}{\delta} + 2^r \right) 2^r \\ &\times \left(2(42 \lambda^{-1/2} \mathcal{B}(n_\gamma, \lambda) + 1)^{3/2} \log^2 \frac{2}{\delta} + 1 \right). \end{aligned}$$

Due to (46), we have $\mathcal{B}(\ell, \lambda)$ is monotonously decreasing with respect to ℓ . Therefore, (43) and (52) yield $\mathcal{B}(n_\gamma, \lambda) \leq \mathcal{B}(m_\gamma, \lambda)$. Inserting it into the above estimate, we obtain from $1/2 \leq r \leq 1$ that

$$\begin{aligned} \|f_{D, D_m, \lambda} - f_\rho\|_\rho &\leq C_1 \mathcal{B}(n_\gamma, \lambda) (42^2 \lambda^{-1} \mathcal{B}^2(n_\gamma, \lambda) + 1) \log^3 \frac{2}{\delta} \\ &+ \lambda^r \|h_\rho\|_\rho \left(3528^r \lambda^{-r} \mathcal{B}^{2r}(m_\gamma, \lambda) \log^{2r} \frac{2}{\delta} + 2^r \right) \left(2(42 \mathcal{B}(n_\gamma, \lambda) \lambda^{-1/2} + 1)^{3/2} \log^2 \frac{2}{\delta} + 2 \right) 2^r \\ &\leq C_2 (\mathcal{B}(n_\gamma, \lambda) \lambda^{-1/2} + 1)^{3/2} + 1 (\mathcal{B}(n_\gamma, \lambda) + \mathcal{B}^{2r}(m_\gamma, \lambda) + \lambda^r) \log^4 \frac{2}{\delta}, \end{aligned}$$

where $C_2 := \max\{1764C_1, 7056^r \times 545\|h_\rho\|_\rho\}$. This completes the proof of Theorem 17 with $\bar{C} = C_2$. \blacksquare

With the help of Theorem 17, we can prove our main results easily.

Proof [Proof of Theorem 4] Since $\lambda = n_\gamma^{-\frac{1}{2r+s}}$, it follows from (46) and Assumption 4 with $0 < s \leq 1$ and $1/2 \leq r \leq 1$ that

$$\mathcal{B}(n_\gamma, \lambda) \leq \frac{\sqrt{\mathcal{N}(\lambda)}}{\sqrt{n_\gamma}} + \frac{1}{n_\gamma \sqrt{\lambda}} \leq (\sqrt{C_0} + 1)n_\gamma^{-\frac{r}{2r+s}}$$

and

$$\mathcal{B}(m_\gamma, \lambda) \leq \frac{\sqrt{\mathcal{N}(\lambda)}}{\sqrt{m_\gamma}} + \frac{1}{m_\gamma \sqrt{\lambda}} \leq \sqrt{C_0} n_\gamma^{\frac{s}{4r+2s}} m_\gamma^{-\frac{1}{2}} + m_\gamma^{-1} n_\gamma^{\frac{1}{4r+2s}}.$$

But Assumption 2 holds with τ_j satisfying (8), which together with (43) and (52) yields

$$\mathcal{B}(n_\gamma, \lambda) \leq C_3 n^{-\frac{r}{2r+s}} (1 + \log(n))^{\frac{r}{(2r+s)\gamma_0}}, \quad (53)$$

where $C_3 := (\sqrt{C_0} + 1)(2 + 2 \log c_1)^{r/(2r+s)}$. Furthermore, (19) implies $m_\gamma \geq C_4 n_\gamma^{\frac{s+1}{2r+s}}$ with $C_4 := (2 + 2 \log c_1)^{\frac{1-2r}{(2r+s)\gamma_0}}$. Then, we have

$$\mathcal{B}^{2r}(m_\gamma, \lambda) \leq C_5 n^{-\frac{r}{2r+s}} (1 + \log(n))^{\frac{r}{(2r+s)\gamma_0}}, \quad (54)$$

where $C_5 := C_4^{-r} C_3$. Since $r > 1/2$, there is a constant $C_7 > 0$ depending only on r, s, γ_0 and C_5 such that

$$n^{(1-2r)/2(2r+s)} (1 + \log(n))^{\frac{2r-1}{2(2r+s)\gamma_0}} \leq C_7.$$

Then, plugging (53) and (54) into (51), we obtain that with confidence $1 - \delta$, there holds

$$\|f_{D, D_m, \lambda} - f_\rho\|_\rho \leq C^* n^{-\frac{r}{2r+s}} (1 + \log(n))^{\frac{r}{(2r+s)\gamma_0}} \log^4 \frac{2}{\delta},$$

where $C^* := \bar{C}(C_3 C_7 + 1)^{\frac{3}{2}}(C_3 + C_5 + 1)$. This completes the proof of Theorem 4. \blacksquare

Proof [Proof of Theorem 7] Due to Assumption 4, (9), (43) and (52) yield

$$n_\gamma \leq C_8 \lambda^{\frac{1-s}{2\gamma_1+1}} n^{\frac{2\gamma_1}{2\gamma_1+1}}, \quad \text{and } m_\gamma \leq C_8 \lambda^{\frac{1-s}{2\gamma_1+1}} m^{\frac{2\gamma_1}{2\gamma_1+1}},$$

where $C_8 := (c_2 C_0)^{\frac{1}{2\gamma_1+1}}$. Then (46), (23) and $\lambda = n^{-\frac{2\gamma_1}{2\gamma_1(2r+s)+2r+1}}$ shows

$$\begin{aligned} \mathcal{B}(n_\gamma, \lambda) &\leq C_0^{1/2} C_8^{-1/2} \lambda^{-\frac{2s\gamma_1+1}{4\gamma_1+2}} n^{-\frac{\gamma_1}{2\gamma_1+1}} + C_8^{-1} \lambda^{-\frac{2\gamma_1-2s+3}{4\gamma_1+2}} n^{-\frac{2\gamma_1}{2\gamma_1+1}} \\ &\leq (C_0^{1/2} C_8^{-1/2} + C_8^{-1}) n^{-\frac{2\gamma_1 r}{2\gamma_1(2r+s)+2r+1}}, \end{aligned}$$

and

$$\begin{aligned} \mathcal{B}(m_\gamma, \lambda) &\leq C_0^{1/2} C_8^{-1/2} \lambda^{-\frac{2s\gamma_1+1}{4\gamma_1+2}} m^{-\frac{\gamma_1}{2\gamma_1+1}} + C_8^{-1} \lambda^{-\frac{2\gamma_1-2s+3}{4\gamma_1+2}} m^{-\frac{2\gamma_1}{2\gamma_1+1}} \\ &\leq (C_0^{1/2} C_8^{-1/2} + C_8^{-1}) n^{-\frac{\gamma_1}{2\gamma_1(2r+s)+2r+1}}. \end{aligned}$$

Inserting the above two estimates into (51), we obtain that

$$\|f_{D, D_m, \lambda} - f_\rho\|_\rho \leq \hat{C} n^{-\frac{2\gamma_1 r}{2\gamma_1(2r+s)+2r+1}} \log^4 \frac{2}{\delta}$$

holds with confidence $1 - \delta$, where $\hat{C} := 2C_2(C_0^{1/2} C_8^{-1/2} + C_8^{-1} + 1)^3$. This completes the proof of Theorem 7. \blacksquare

Acknowledge

The work is supported by the National Key R&D Program of China (No.2020YFA0713900) and the National Natural Science Foundation of China (Nos.62276209,618761332,11971374).

References

- P. Alquier and O. Wintenberger. Model selection for weakly dependent time series forecasting. *Bernoulli*, 18(3): 883-913, 2012.
- P. Alquier, X. Li and O. Wintenberger. Prediction of time series by statistical learning: general losses and fast rates. *Dependence Modeling*, 1: 65-93, 2013.
- D. W. K. Andrews. Nonstrong mixing autoregressive processes. *Journal of Applied Probability*, 21: 930-934, 1984.
- N. Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68 (3): 337-404, 1950.
- R. Bathis. Matrix Analysis, Volume 169 of Graduate Texts in Mathematics. Springer, Berlin, 1997.
- P. Billingsley. Convergence of Probability Measure. Wiley, New York, 1968.
- G. Blanchard and O. Zadorozhnyi. Concentration of weakly dependent Banach-valued sums and applications to statistical learning methods. *Bernoulli*, 25(4B): 3421-3458, 2019.
- R. C. Bradley. Basic properties of strong mixing conditions. A survey and some open questions. *Probability Surveys*, 2: 107-144, 2005.
- A. Caponnetto and E. DeVito. Optimal rates for the regularized least squares algorithm. *Foundations of Computational Mathematics*, 7: 331-368, 2007.
- X. Chang, S. B. Lin and D. X. Zhou. Distributed semi-supervised learning with kernel ridge regression. *Journal of Machine Learning Research*, 18(1): 1493-1514, 2017.

- X. Chen and X. Shen. Sieve extremum estimates for weakly dependent data. *Econometrica*, 66: 289-314, 1998.
- F. Cucker and D. X. Zhou. Learning Theory: an Approximation Theory Viewpoint. Cambridge University Press, Cambridge, 2007.
- J. Dedecker and C. Prieur. Coupling for τ -dependent sequences and applications. *Journal of Theoretical Probability*, 17: 861-885, 2004.
- J. Dedecker and C. Prieur. New dependence coefficients. Examples and applications to statistics. *Probability Theory and Related Fields*, 132(2): 203-236, 2005.
- P. Doukhan. Mixing: Properties and Examples. New York: Springer-Verlag, 1994.
- T. Evgeniou, M. Pontil and T. Poggio. Regularization networks and support vector machines. *Advances in Computational Mathematics*, 13: 1-50, 2000.
- J. Fan and Q. Yao. Nonlinear Time Series: Nonparametric and Parametric Methods. Springer Science & Business Media, 2008.
- H. I. Fawaz, G. Forestier, J. Weber, L. Idoumghar and P. A. Muller. Deep learning for time series classification: a review. *Data Mining and Knowledge Discovery*, 33(4): 917-963, 2019.
- T. C. Fu. A review on time series data mining. *Engineering Applications of Artificial Intelligence*, 24(1): 164-181, 2011.
- S. Fischer, I. Steinwart. Sobolev norm learning rates for regularized Least-Squares algorithms. *Journal of Machine Learning Research*, 21: 1-38, 2020.
- A. Gittens and M. W. Mahoney. Revisiting the Nyström method for improved large scale machine learning. *Journal of Machine Learning Research*, 17: 1-65, 2016.
- Z. C. Guo, S. B. Lin and D. X. Zhou. Learning theory of distributed spectral algorithms. *Inverse Problems*, 33: 074009, 2017.
- L. Györfy, M. Kohler, A. Krzyzak and H. Walk. A Distribution-Free Theory of Nonparametric Regression. Springer, Berlin, 2002.
- M. T. Hagan, H. B. Demuth and M. Beale. Neural Network Design. PWS Publishing Company, 1997.
- H. Hang and I. Steinwart. A Bernstein-type inequality for some mixing processes and dynamical systems with an application to learning. *The Annals of Statistics*, 45: 708-743, 2017.
- P. Indyk, N. Koudas and S. Muthukrishnan. Identifying representative trends in massive time series data sets using sketches. In *26th International Conference on Very Large Data Bases, VLDB 2000*, pages 363-372.

- G. Kriukova, S. Pereverzyev and P. Tkachenko. Nyström type subsampling analyzed as a regularized projection. *Inverse Problems*, 33(7): 074001, 2017.
- N. Laptev, C. Zaniolo, T. C. Lu and C. A. Malibu. BOOT-TS: A scalable bootstrap for massive time-series data. NIPS, 2012.
- S. B. Lin, X. Guo and D. X. Zhou. Distributed learning with regularized least squares. *Journal of Machine Learning Research*, 18(92): 1-31, 2017.
- S. B. Lin, D. Wang and D. X. Zhou. Distributed kernel ridge regression with communications. *Journal of Machine Learning Research*, 21(93): 1-38, 2020.
- S. B. Lin, J. Fan and X. Chang. Learning with selected features. *IEEE Transactions on Cybernetics*, 52(4): 2032-2046, 2022.
- S. N. Lahiri. Theoretical comparisons of block bootstrap methods. *Annals of Statistics*, 27(1): 386-404, 1999.
- L. Lo Gerfo, L. Rosasco, F. Odone, E. De Vito and A. Verri. Spectral algorithms for supervised learning. *Neural Computation*, 20: 1873-1897, 2008.
- S. Lu, Mathé, P. and S. Pereverzyev Jr. Analysis of regularized Nyström subsampling for regression functions of low smoothness. *Analysis and Applications*, 17: 931-946, 2019.
- V. Maume-Deschamps. Exponential inequalities and functional estimations for weak dependent data: applications to dynamical systems. *Stochastics and Dynamics*, 6(04): 535-560, 2006.
- D. J. McDonald, C. R. Shalizi and M. Schervish. Nonparametric risk bounds for time-series forecasting. *Journal of Machine Learning Research*, 18(1): 1044-1083, 2017.
- R. Meir. Nonparametric time series prediction through adaptive model selection. *Machine Learning*, 39(1): 5-34, 2000.
- M. Meister and I. Steinwart. Optimal learning rates for localized SVMs. *Journal of Machine Learning Research*, 17(1): 6722-6765, 2016.
- D. S. Modha and E. Masry. Minimum complexity regression estimation with weakly dependent observations. *IEEE Transactions on Information Theory*, 42: 2133-2145, 1996.
- D. S. Modha. Memory-universal prediction of stationary random process. *IEEE Transactions on Information Theory*, 44: 117-133, 1998.
- A. Mueen, S. Nath and J. Liu. Fast approximate correlation for massive time-series data. *In Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, pages 171-182, 2010.
- A. Rahimi and B. Recht. Random features for large-scale kernel machines. *Neural Information Processing Systems*, pages 1177-1184, 2007.

- T. Rakthanmanon, B. Campana, A. Mueen, et al. Searching and mining trillions of time series subsequences under dynamic time warping. *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 262-270, 2012.
- M. Rosenblatt. A central limit theorem and a strong mixing condition. *Proceedings of the National Academy of Sciences*, 42: 43-47, 1956.
- A. Rudi, R. Camoriano and L. Rosasco. Less is more: Nyström computational regularization. *Neural Information Processing Systems*, pages 1657-1665, 2015.
- Samson and Paul-Marie. Concentration of measure inequalities for Markov chains and ϕ -mixing processes. *The Annals of Probability*, 28(1): 416-461, 2000.
- J. Shawe-Taylor and N. Cristianini. Kernel Methods for Pattern Analysis. Cambridge University Press, Cambridge, 2004.
- L. Shi. Distributed learning with indefinite kernels. *Analysis and Applications*, 17(06): 947-975, 2019.
- S. Smale and D. X. Zhou. Learning theory estimates via integral operators and their approximations. *Constructive Approximation*, 26: 153-172, 2007.
- I. Steinwart and A. Christmann. Support Vector Machines. Springer, New York, 2008.
- I. Steinwart and A. Christmann. Fast learning from non-i.i.d. observations. *Neural Information Processing Systems*, pages 1768-1776, 2009.
- H. Sun and Q. Wu. Regularized least square regression with dependent samples. *Advances In Computational Mathematics*, 32: 175-189, 2010.
- Z. Sun and S. B. Lin. Distributed learning with dependent samples. *IEEE Transactions on Information Theory*, 68(9): 6003-6020, 2022.
- P. Thomann, I. Blaschzyk, M. Meister and I. Steinwart. Spatial decompositions for large scale SVMs. In *Artificial Intelligence and Statistics, Proceedings of Machine Learning Research*, 54: 1329-1337, 2017.
- C. K. I. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. *Proceedings of the 13th International Conference on Neural Information Processing Systems*, pages 661-667, 2000.
- D. Wang, J. Zeng and S. B. Lin. Random sketching for neural networks with ReLU. *IEEE Transactions on Neural Networks and Learning Systems*, 32: 748-762, 2021.
- Y. L. Xu and D. R. Chen. Learning rates of regularized regression for exponentially strongly mixing sequence. *Journal of statistical planning and inference*, 138: 2180-2189, 2008.
- Z. Xu, M. Dai and D. Meng. Fast and efficient strategies for model selection of Gaussian support vector machine. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(5): 1292-1307, 2009.

- Y. Yang, M. Pilanci and M. J. Wainwright. Randomized sketches for kernels: Fast and optimal nonparametric regression. *Annals of Statistics*, 45(3): 991-1023, 2017.
- B. Yu. Rates of convergence for empirical processes of stationary mixing sequences. *Annals of Probability*, 22: 94-116, 1994.
- T. Zhang. Learning bounds for kernel regression using effective data dimensionality. *Neural Computation*, 17: 2077-2098, 2005.
- Y. C. Zhang, J. Duchi and M. Wainwright. Communication-efficient algorithms for statistical optimization. *Journal of Machine Learning Research*, 14: 3321-3363, 2013.
- Y. C. Zhang, J. Duchi and M. Wainwright. Divide and conquer kernel ridge regression: A distributed algorithm with minimax optimal rates. *Journal of Machine Learning Research*, 16: 3299-3340, 2015.
- D. X. Zhou. Capacity of reproducing kernel spaces in learning theory, *IEEE Transactions on Information Theory*, 49: 1743-1752, 2003.
- H. H. Zhou and G. Raskutti. Non-parametric sparse additive auto-regressive network models. *IEEE Transactions on Information Theory*, 65: 1473-1492, 2019.