EFFECTIVE DISTRIBUTED LEARNING WITH RANDOM FEATURES: IMPROVED BOUNDS AND ALGORITHMS

Anonymous authors
Paper under double-blind review

ABSTRACT

In this paper, we study the statistical properties of distributed kernel ridge regression together with random features (DKRR-RF), and obtain optimal generalization bounds under the basic setting, which can substantially relax the restriction on the number of local machines in the existing state-of-art bounds. Specifically, we first show that the simple combination of divide-and-conquer technique and random features can achieve the same statistical accuracy as the exact KRR in expectation requiring only $\mathcal{O}(|\mathcal{D}|)$ memory and $\mathcal{O}(|\mathcal{D}|^{1.5})$ time. Then, beyond the generalization bounds in expectation that demonstrate the average information for multiple trails, we derive generalization bounds in probability to capture the learning performance for a single trail. Finally, we propose an effective communication strategy to further improve the performance of DKRR-RF, and validate the theoretical bounds via numerical experiments.

1 INTRODUCTION

Kernel ridge regression (KRR) is one of the most popular nonparametric learning methods (Vapnik, 2000). Despite the excellent theoretical guarantees, KRR does not scale well in large scale settings because of high time and memory complexities. Distributed learning (Zhang et al., 2013; Hsieh et al., 2014; Chang et al., 2017b; Li et al., 2019a; Lin et al., 2020), random features (Rahimi & Recht, 2007; Rudi & Rosasco, 2017; Rudi et al., 2018; Liu et al., 2020; Avron et al., 2017a; Yu et al., 2016; Jacot et al., 2020), and Nyström methods (Drineas & Mahoney, 2005; Ding & Liao, 2012; Yang et al., 2012; Camoriano et al., 2016; Si et al., 2016; Musco & Musco, 2017; Kriukova et al., 2017) are the most widely used large scale techniques to address the scalability issues. Recent statistical learning works on KRR together with large scale approaches demonstrate that these large scale approaches can not only obtain great computational gains but also can guarantee the optimal theoretical properties, such as KRR with divide-and-conquer (Zhang et al., 2013; 2015; Chang et al., 2017b;a; Guo et al., 2017; Lin et al., 2017; Li et al., 2019a; Lin et al., 2020), with random features (Rudi & Rosasco, 2017; Li et al., 2019b; Carratino et al., 2018; Yang et al., 2012), and with Nyström methods (Bach, 2013; Alaoui & Mahoney, 2015; Rudi et al., 2015; 2017).

The combinations of distributed learning and other large scale approaches are very intuitive but effective strategies to further improve the effectiveness, such as distributed learning with gradient descent algorithms (Lin & Zhou, 2018; Richards et al., 2020), with multi-pass SGD (Lin & Rosasco, 2017), with random features (Li et al., 2019a), and with Nyström methods (Yin et al., 2020). The optimal generalization performance of these combining approaches has been studied, however, the main theoretical problem is that there is a strict restriction on the number of local machines. For sample, in (Lin & Zhou, 2018; Li et al., 2019a; Yin et al., 2020), to guarantee the optimal generalization performance in the basic setting, the upper bounds of the local machines are restricted to be a constant, which is difficult to be satisfied in real applications.

In this paper, we aim at enlarging the number of local machines by considering communications among different local machines. This paper makes the following three main contributions. Firstly, we improve the existing state-of-art results of the divide-and-conquer technique together with random features. We prove that the optimal generalization performance can be guaranteed even the partitions reach $\Omega(\sqrt{|\mathcal{D}|})$, which are limited to a constant $\Omega(1)$ for the existing bounds in the basic setting, $|\mathcal{D}|$ is the size of the data sets. Secondly, to essentially reflect the generalization performance, beyond the minimax optimal rates in expectation, we derive optimal learning rates in probability, which can

capture the learning performance for a single trail. Finally, we develop a communication strategy to further improve the performance of our proposed method, and validate the effectiveness of the proposed communications via both theoretical assessments and numerical experiments.

Related Work

The most related work includes the statistical analysis of distributed learning and random features.

Distributed learning. Optimal learning rates for divide-and-conquer KRR in expectation were established in the seminal work (Zhang et al., 2013; 2015). An improved bound was derived in (Lin et al., 2017) based on a novel tool of integral operator. Based on the proof techniques proposed in (Zhang et al., 2013; 2015; Lin et al., 2017), optimal learning rates were established for distributed spectral algorithms (Guo et al., 2017), distributed gradient descent algorithms (Lin & Zhou, 2018; Richards et al., 2020), distributed semi-supervised learning (Chang et al., 2017b), distributed local average regression (Chang et al., 2017a), localized SVM (Meister & Steinwart, 2016), etc. The theoretical analysis mentioned above shows that the divide-and-conquer learning can achieve the same statistical accuracy as the exact KRR, however, there is a strict restriction on the number of local machines. To enlarge the number of local machines, Lin et al. (2020) considered the communications among different local machines. However, the communication strategy proposed in (Lin et al., 2020) based on an operator representation, which requires communicating the input data between each local machine. Thus, it is difficult to protect the data privacy of each local machine. Furthermore, for each iteration, the communication complexity of each local machine is $\mathcal{O}(|\mathcal{D}|d)$, where d denotes the dimension, which is infeasible in practice for large scale data sets.

Random Features. The generalization bound of random features was first proposed in (Rahimi & Recht, 2008), which shows that $\mathcal{O}(|\mathcal{D}|)$ random features are needed for $\mathcal{O}(1/\sqrt{|\mathcal{D}|})$ learning rate. Some works further studied its theoretical performance (Cortes et al., 2010; Yang et al., 2012; Bach, 2017; Avron et al., 2017b). By applying the standard integral operator framework (Smale & Zhou, 2007; Caponnetto & Vito, 2007), the optimal generalization bounds of KRR with random features were established in (Rudi & Rosasco, 2017), which requires only $\mathcal{O}(\sqrt{|\mathcal{D}|\log(|\mathcal{D}|)})$ random features. To decrease the size of random features, an improved approach was proposed based on a novel leverage score sampling strategy (Rudi et al., 2018). Sun et al. (2018) extended the result of (Rudi & Rosasco, 2017) to SVM. In (Li et al., 2019b), they further devised a simple framework for the unified analysis of random Fourier features, which can be applied to KRR, as well as SVM and logistic regression. To further improve the effectiveness, recently, Li et al. (2019a) considered the simple combination of divide-and-conquer and random features. However, to guarantee the optimal generalization performance, the number of local machines should be restricted to a constant, degenerating it into a single random features-based large scale KRR.

2 BACKGROUND

In a standard framework of supervised learning, there is a probability space $\mathcal{X} \times \mathcal{Y}$ with an unknown distribution ρ, where $\mathcal{X}$ is the input and $\mathcal{Y}$ is the output space. The sample set $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ of size n is drawn i.i.d from $\mathcal{X} \times \mathcal{Y}$ with respect to ρ. Let $K : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ be a Mercer kernel, and $\mathcal{H}_K$ be its reproducing kernel Hilbert space (RKHS) (Steinwart & Christmann, 2008; Vapnik, 2000), and assume that $K(\mathbf{x}, \mathbf{x}') \leq \kappa, \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$. Throughout, we will denote the inner product in $\mathcal{H}_K$ by $\langle \cdot, \cdot \rangle_K$, and corresponding norm by $\| \cdot \|_K$.

Kernel Ridge Regression (KRR)
KRR is one of the most popular nonparametric learning methods (Shawe-Taylor & Cristianini, 2000; Vapnik, 2000), which can be stated as

$$f_{\mathcal{D},\lambda} = \underset{f \in \mathcal{H}_K}{\arg\min} \left\{ \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} (f(\mathbf{x}_i) - y_i)^2 + \lambda \|f\|_K^2 \right\}, \tag{1}$$

where $\lambda > 0$ is the regularization parameter, $|\mathcal{D}|$ is the size of $\mathcal{D}$. Using the representation theorem (Shawe-Taylor & Cristianini, 2000; Vapnik, 2000), $f_{\mathcal{D},\lambda}$ can be written as $f_{\mathcal{D},\lambda}(\mathbf{x}) = \sum_{i=1}^{|\mathcal{D}|} \alpha_i K(\mathbf{x}_i, \mathbf{x})$ with $\boldsymbol{\alpha} = (\mathbf{K}_{\mathcal{D}} + \lambda \mathbf{I})^{-1} \mathbf{y}_{\mathcal{D}}$, where $\mathbf{K}_{\mathcal{D}} = \frac{1}{|\mathcal{D}|} \left[K_{\mathcal{D}}(\mathbf{x}_i, \mathbf{x}_j) \right]_{i,j=1}^{|\mathcal{D}|}$ is the $|\mathcal{D}| \times |\mathcal{D}|$ kernel matrix and $\mathbf{y}_{\mathcal{D}} = \frac{1}{|\mathcal{D}|}(y_1, \ldots, y_{|\mathcal{D}|})^{\mathrm{T}}$.

Despite the excellent theoretical guarantees (Blanchard & Krämer, 2010; Caponnetto & Vito, 2007), KRR requires $\mathcal{O}(|\mathcal{D}|^2)$ memory to store $\mathbf{K}_{\mathcal{D}}$, and $\mathcal{O}(|\mathcal{D}|^3)$ time to solve inverse of $\mathbf{K}_{\mathcal{D}} + \lambda\mathbf{I}$, which is infeasible for large scale settings.

KRR with Random Features (KRR-RF)
Assuming the spectral measure has a density function $\varrho(\cdot)$, the corresponding shift-invariant kernel can be written as $K(\mathbf{x}, \mathbf{x}') = \int_\Omega \psi(\mathbf{x}, \boldsymbol{\omega})\psi(\mathbf{x}', \boldsymbol{\omega})\varrho(\boldsymbol{\omega})d\boldsymbol{\omega}$, where $\psi : \mathcal{X} \times \Omega \to \mathbb{R}$ is a continuous and bounded function with respect to $\boldsymbol{\omega}$ and $\mathbf{x}$. The main idea behind random Fourier features is to approximate the kernel function $K(\mathbf{x}, \mathbf{x}')$ by its Monte-Carlo estimation (Rahimi & Recht, 2007):

$$K_M(\mathbf{x}, \mathbf{x}') = \frac{1}{M}\sum_{i=1}^{M} \psi(\mathbf{x}, \boldsymbol{\omega}_i)\psi(\mathbf{x}', \boldsymbol{\omega}_i) = \langle \boldsymbol{\phi}_M(\mathbf{x}), \boldsymbol{\phi}_M(\mathbf{x}')\rangle,$$

where $\boldsymbol{\phi}_M(\mathbf{x}) = \frac{1}{\sqrt{M}}\left(\psi(\mathbf{x}, \boldsymbol{\omega}_1), \ldots, \psi(\mathbf{x}, \boldsymbol{\omega}_M)\right)^{\mathrm{T}}$. The solution of KRR with random features can be written as

$$f_{M,\mathcal{D},\lambda}(\mathbf{x}) = \mathbf{w}_{M,\mathcal{D},\lambda}^{\mathrm{T}}\boldsymbol{\phi}_M(\mathbf{x}) \text{ with } \mathbf{w}_{M,\mathcal{D},\lambda} = (\boldsymbol{\Phi}_{M,\mathcal{D}}\boldsymbol{\Phi}_{M,\mathcal{D}}^{\mathrm{T}} + \lambda\mathbf{I})^{-1}\boldsymbol{\Phi}_{M,\mathcal{D}}\bar{\mathbf{y}}_{\mathcal{D}}, \tag{2}$$

where $\boldsymbol{\Phi}_{M,\mathcal{D}} = \frac{1}{\sqrt{|\mathcal{D}|}}(\boldsymbol{\phi}_M(\mathbf{x}_1), \ldots, \boldsymbol{\phi}_M(\mathbf{x}_{|\mathcal{D}|}))$ and $\bar{\mathbf{y}}_{\mathcal{D}} = \frac{1}{\sqrt{|\mathcal{D}|}}(y_1, \ldots, y_{|\mathcal{D}|})^{\mathrm{T}}$.

KRR-RF requires $\mathcal{O}(M|\mathcal{D}|)$ to store $\boldsymbol{\Phi}_{M,\mathcal{D}}$, $\mathcal{O}(M^3)$ and $\mathcal{O}(M^2|\mathcal{D}|)$ time to solve the inverse of $(\boldsymbol{\Phi}_{M,\mathcal{D}}\boldsymbol{\Phi}_{M,\mathcal{D}}^{\mathrm{T}} + \lambda\mathbf{I})$ and the matrix multiplication $\boldsymbol{\Phi}_{M,\mathcal{D}}\boldsymbol{\Phi}_{M,\mathcal{D}}$, respectively. Thus, the total space and time complexity of KRR-RF are $\mathcal{O}(M|\mathcal{D}|)$ and $\mathcal{O}(M^2|\mathcal{D}|)$, $M \ll |\mathcal{D}|$, respectively.

Distributed KRR with Random Features (DKRR-RF)
Let $\{\mathcal{D}_j\}_{j=1}^m$ be m disjoint subsets with $\mathcal{D} = \cup_{j=1}^m \mathcal{D}_j$. The distributed KRR with random features (DKRR-RF) is defined as

$$\bar{f}_{M,\mathcal{D},\lambda}^0 = \sum_{j=1}^{m} \frac{|\mathcal{D}_j|}{|\mathcal{D}|} f_{M,\mathcal{D}_j,\lambda}, \tag{3}$$

where $f_{M,\mathcal{D}_j,\lambda}(\mathbf{x}) = \mathbf{w}_{M,\mathcal{D}_j,\lambda}^{\mathrm{T}}\boldsymbol{\phi}_M(\mathbf{x})$ with $\mathbf{w}_{M,\mathcal{D}_j,\lambda} = (\boldsymbol{\Phi}_{M,\mathcal{D}_j}\boldsymbol{\Phi}_{M,\mathcal{D}_j}^{\mathrm{T}} + \lambda\mathbf{I})^{-1}\boldsymbol{\Phi}_{M,\mathcal{D}_j}\bar{\mathbf{y}}_{\mathcal{D}_j}$. The space complexity, time complexity and communication complexity of DKRR-RF for each local machine are $\mathcal{O}(M|\mathcal{D}_j|)$, $\mathcal{O}(M^3 + M^2|\mathcal{D}_j|)$ and $\mathcal{O}(M)$, respectively.

3 DKRR-RF WITH COMMUNICATIONS (DRKK-RF-CM)

In this section, we will present an effective communication strategy to enlarge the number of local machines. We first give the motivation of our communication strategy, and then propose a communication-based method, called DRKK-RF-CM.

Motivation. Let $g_{M,\mathcal{D},\lambda} : \mathbb{R}^M \to \mathbb{R}^M$ be $g_{M,\mathcal{D},\lambda}(\mathbf{w}) := \left[\boldsymbol{\Phi}_{M,\mathcal{D}}\boldsymbol{\Phi}_{M,\mathcal{D}}^{\mathrm{T}} + \lambda\mathbf{I}\right]\mathbf{w} - \boldsymbol{\Phi}_{M,\mathcal{D}}\bar{\mathbf{y}}_{\mathcal{D}}$. One can see that $2g_{M,\mathcal{D},\lambda}(\mathbf{w})$ is the gradient of the empirical risk of $\frac{1}{|\mathcal{D}|}\sum_{(\mathbf{x}_i,y_i)\in\mathcal{D}}\left(\mathbf{w}^{\mathrm{T}}\boldsymbol{\phi}_M(\mathbf{x}_i) - y_i\right)^2 + \lambda\|\mathbf{w}\|^2$ on $\mathbf{w}$. From Eq.2, we know that for any $\mathbf{w}$, the following equation holds:

$$\begin{aligned}\mathbf{w}_{M,\mathcal{D},\lambda} &= \mathbf{w} - \left[\boldsymbol{\Phi}_{M,\mathcal{D}}\boldsymbol{\Phi}_{M,\mathcal{D}}^{\mathrm{T}} + \lambda\mathbf{I}\right]^{-1}\left[\left[\boldsymbol{\Phi}_{M,\mathcal{D}}\boldsymbol{\Phi}_{M,\mathcal{D}}^{\mathrm{T}} + \lambda\mathbf{I}\right]\mathbf{w} - \boldsymbol{\Phi}_{M,\mathcal{D}}\bar{\mathbf{y}}_{\mathcal{D}}\right] \\ &= \mathbf{w} - \left[\boldsymbol{\Phi}_{M,\mathcal{D}}\boldsymbol{\Phi}_{M,\mathcal{D}}^{\mathrm{T}} + \lambda\mathbf{I}\right]^{-1} g_{M,\mathcal{D},\lambda}(\mathbf{w}).\end{aligned} \tag{4}$$

Define $\bar{\mathbf{w}}_{M,\mathcal{D},\lambda}^0 = \sum_{j=1}^m \frac{|\mathcal{D}_j|}{|\mathcal{D}|}\mathbf{w}_{M,\mathcal{D}_j,\lambda}$, it is easy to verify that

$$\bar{\mathbf{w}}_{M,\mathcal{D},\lambda}^0 = \mathbf{w} - \sum_{j=1}^{m}\frac{|\mathcal{D}_j|}{|\mathcal{D}|}\left[\boldsymbol{\Phi}_{M,\mathcal{D}_j}\boldsymbol{\Phi}_{M,\mathcal{D}_j}^{\mathrm{T}} + \lambda\mathbf{I}\right]^{-1} g_{M,\mathcal{D}_j,\lambda}(\mathbf{w}). \tag{5}$$

Comparing 4 and 5, and noting that the global gradient $g_{M,\mathcal{D},\lambda}(\mathbf{w})$ can be achieved via the communications of each local gradient $g_{M,\mathcal{D}_j,\lambda}(\mathbf{w})$, i.e., $g_{M,\mathcal{D},\lambda}(\mathbf{w}) = \sum_{j=1}^m \frac{|\mathcal{D}_j|}{|\mathcal{D}|} g_{M,\mathcal{D}_j,\lambda}(\mathbf{w})$, thus, we consider the following Newton Raphson iteration-based communication strategy:

$$\bar{\mathbf{w}}_{M,\mathcal{D},\lambda}^t = \bar{\mathbf{w}}_{M,\mathcal{D},\lambda}^{t-1} - \sum_{j=1}^{m}\frac{|\mathcal{D}_j|}{|\mathcal{D}|}\left[\boldsymbol{\Phi}_{M,\mathcal{D}_j}\boldsymbol{\Phi}_{M,\mathcal{D}_j}^{\mathrm{T}} + \lambda\mathbf{I}\right]^{-1} g_{M,\mathcal{D},\lambda}(\bar{\mathbf{w}}_{M,\mathcal{D},\lambda}^{t-1}). \tag{6}$$

Algorithm 1 Distributed KRR with Random Features and Communications (DKRR-RF-CM)

Initialize: $\bar{\mathbf{w}}^0_{M,\mathcal{D},\lambda} = \mathbf{0}$
for $t = 1$ **to** p **do**
  Local machine: compute the local gradient $g_{M,\mathcal{D}_j,\lambda}(\bar{\mathbf{w}}^{t-1}_{M,\mathcal{D},\lambda})$, and communicate back to GM.
  Global machine: get the global gradient $g_{M,\mathcal{D},\lambda}(\bar{\mathbf{w}}^{t-1}_{M,\mathcal{D},\lambda}) = \sum_{j=1}^m \frac{|\mathcal{D}_j|}{|\mathcal{D}|} g_{M,\mathcal{D}_j,\lambda}(\bar{\mathbf{w}}^{t-1}_{M,\mathcal{D},\lambda})$ and communicate to each local machine.
  Local machine: compute $\boldsymbol{\beta}_j^{t-1} = \left[\boldsymbol{\Phi}_{M,\mathcal{D}_j}\boldsymbol{\Phi}^{\mathrm{T}}_{M,\mathcal{D}_j} + \lambda\mathbf{I}\right]^{-1} g_{M,\mathcal{D},\lambda}(\bar{\mathbf{w}}^{t-1}_{M,\mathcal{D},\lambda})$ and communicate back to the global machine.
  Global machine: compute $\bar{\mathbf{w}}^t_{M,\mathcal{D},\lambda} = \bar{\mathbf{w}}^{t-1}_{M,\mathcal{D},\lambda} - \frac{|\mathcal{D}_j|}{|\mathcal{D}|}\sum_{j=1}^m \boldsymbol{\beta}_j^{t-1}$, and communicate to each local machine.
end for
Output: $\bar{\mathbf{w}}^p_{M,\mathcal{D},\lambda}$ **and** $\bar{f}^p_{M,\mathcal{D},\lambda} = \langle \bar{\mathbf{w}}^p_{M,\mathcal{D},\lambda}, \boldsymbol{\phi}_M(\cdot)\rangle$

We propose an iterative procedure to implement the communication strategy Eq.6, which can be broken down into 4 steps. At first, each local machine computes the local gradient and communicates back to the global machine. Then the global machine computes the global gradient based on the local gradient, and communicates it to each local machine. In the third step, each local machine computes $\boldsymbol{\beta}_j^{t-1}$, and communicates back to the global machine. Finally, the global machine obtains the solution $\bar{\mathbf{w}}^t_{M,\mathcal{D},\lambda}$. More details can be seen in Algorithm 1.

Complexity analysis. Space complexity: each local machine only needs to store $\boldsymbol{\Phi}_{M,\mathcal{D}_j}$ and the local gradient $g_{M,\mathcal{D}_j,\lambda}$, thus the space complexity of each local machine is $\mathcal{O}(M|\mathcal{D}_j| + M) = \mathcal{O}(M|\mathcal{D}_j|)$; Time complexity: for each local machine, we only need to compute the matrix multiplication $\boldsymbol{\Phi}_{M,\mathcal{D}_j}\boldsymbol{\Phi}^{\mathrm{T}}_{M,\mathcal{D}_j}$ and the inverse of $\boldsymbol{\Phi}_{M,\mathcal{D}_j}\boldsymbol{\Phi}^{\mathrm{T}}_{M,\mathcal{D}_j} + \lambda\mathbf{I}$ once. For each iteration, we need to compute the local gradient $g_{M,\mathcal{D}_j,\lambda}$ and $\boldsymbol{\beta}_j$ for each local machine. Therefore, the total time complexity of each local machine is $\mathcal{O}(M^3 + M^2|\mathcal{D}_j| + pM|\mathcal{D}_j|)$, where p is the number of communication; Communication complexity: for each iteration, we only communicate the local gradient $g_{M,\mathcal{D}_j,\lambda}$ and $\boldsymbol{\beta}_j$ to the global machine, and receive the gradient $g_{M,\mathcal{D},\lambda}$ and $\bar{\mathbf{w}}^{t-1}_{M,\mathcal{D},\lambda}$ from the global machine, so the total communication complexity is $\mathcal{O}(pM)$.

Remark 1. *From the complexity analysis above, we can see that if the number of the communication p satisfying $p \leq M$ or $p \leq |\mathcal{D}_j|$, then the time and space complexity of DKRR-RF-CM are the same as DKRR-RF. Only the communication complexity is slightly increased from $\mathcal{O}(M)$ to $\mathcal{O}(pM)$.*

4 Theoretical Analysis

In this section, we analyze the generalization performances of DKRR-RF and DKRR-RF-CM. The performance of the algorithm is usually measured by the expected risk

$$\mathcal{E}(f) = \int_{\mathcal{X}\times\mathcal{Y}} (f(\mathbf{x}) - y)^2 d\rho(\mathbf{x}, y). \tag{7}$$

The optimal hypothesis $f_{\mathcal{H}_K}$ in $\mathcal{H}_K$ is denoted by $f_{\mathcal{H}_K} = \arg\min_{f\in\mathcal{H}_K} \mathcal{E}(f)$, and we assume $f_{\mathcal{H}_K}$ exists in the paper.

4.1 Optimal learning rate for DKRR in Expectation

Theorem 1. *Suppose that $|\psi(\mathbf{x}, \boldsymbol{\omega})| \leq \tau$ almost surely, $\tau \in [1, \infty)$ and $|y| \leq \zeta$. If $\lambda = \Omega(|\mathcal{D}|^{-\frac{1}{2}})$, $|\mathcal{D}_1| = \ldots |\mathcal{D}_m|$, the number of partitions m and the number of random features M respectively correspond to*

$$m \lesssim \sqrt{|\mathcal{D}|} \text{ and } M \gtrsim \sqrt{|\mathcal{D}|}, \tag{8}$$

then, for every $\delta \in (0, 1]$, with probability at least $1 - \delta$, we have

$$\mathbb{E}[\mathcal{E}(\bar{f}^0_{M,\mathcal{D},\lambda})] - \mathcal{E}(f_{\mathcal{H}_K}) = \mathcal{O}\left(|\mathcal{D}|^{-\frac{1}{2}} \log^2(1/\delta)\right).$$

From Theorem 1, one can see that if $m \lesssim \sqrt{|\mathcal{D}|}$ and $M \gtrsim \sqrt{|\mathcal{D}|}$, the learning rate of the generalization bound can reach $\mathcal{O}(1/\sqrt{|\mathcal{D}|})$, which is optimal in a minmax sense (Rudi & Rosasco, 2017; Caponnetto & Vito, 2007). It means that, in this basic setting, as long as the number of partitions and random features are in order $\Omega(\sqrt{|\mathcal{D}|})$, the corresponding ridge regression estimator has optimal generalization properties. The assumption of $|y| \leq \zeta$ can be related to the Bernstein condition or moment assumption (Blanchard & Krämer, 2016), but for simplicity, in this paper, we only consider that $\mathcal{Y}$ is bounded.

Optimal learning rates for divide-and-conquer KRR in expectation have been established in (Zhang et al., 2013; 2015; Lin et al., 2017), etc. However, there is a strict restriction on the number of local machines m. Specifically, in (Lin et al., 2017), to reach the optimal rate, m should to restrict to a constant $m = \Omega(1)$. In (Li et al., 2019a), the authors have studied the generalization performance of the combination of divide-and-conquer technique and random features. Using the same setting as Theorem 1 (that is $r = 1/2$ and $\gamma = 1$ in Theorem 8 of Li et al. (2019a)), they prove that, if $M \gtrsim \sqrt{|\mathcal{D}|}$ and $m \lesssim \Omega(1)$, then $\mathbb{E}[\mathcal{E}(\bar{f}^0_{M,\mathcal{D},\lambda})] - \mathcal{E}(f_{\mathcal{H}_K}) = \mathcal{O}(|\mathcal{D}|^{-\frac{1}{2}} \log^2(1/\delta))$. It means that to guarantee the optimal generalization properties, the number of partitions should be restricted to a constant, but for our result is $\Omega(\sqrt{|\mathcal{D}|})$. In (Li et al., 2019a), they also considered using the unlabeled data to enlarge the number of partitions. They have proved that (see Corollary 12 of Li et al. (2019a) for detail), if $M \gtrsim \sqrt{|\mathcal{D}|}$ and $m \lesssim |\mathcal{D}^*|/|\mathcal{D}|$, then $\mathbb{E}[\mathcal{E}(\bar{f}^0_{M,\mathcal{D},\lambda})] - \mathcal{E}(f_{\mathcal{H}_K}) = \mathcal{O}(|\mathcal{D}|^{-\frac{1}{2}} \log^2(1/\delta))$, where $\mathcal{D}^*$ is the dataset includes both labeled and unlabeled data. Thus, if we want the number of partitions m to reach $\Omega(\sqrt{|\mathcal{D}|})$ as the same as our Theorem 1, the size of $|\mathcal{D}^*|$ should be $\Omega(\sqrt{|\mathcal{D}|})$ times of $|\mathcal{D}|$. In this case, the data size of each local machine is $|\mathcal{D}| = |\mathcal{D}|^{3/2}/\sqrt{|\mathcal{D}|}$, so the time and space complexity are the same as KRR with a single random feature technique.

4.2 Optimal learning rates for DKRR in Probability

Note that $\mathbb{E}[\mathcal{E}(\bar{f}^0_{M,\mathcal{D},\lambda})] - \mathcal{E}(f_{\mathcal{H}_K}) = \mathbb{E}[\|\bar{f}^0_{M,\mathcal{D},\lambda} - f_{\mathcal{H}_K}\|^2_\rho]$ (Caponnetto & Vito, 2007), thus Theorem 1 proposes the optimal learning rate *in expectation*, which demonstrates the average information for multiple trails, but may fail to capture the learning performance for a single trail. To essentially reflect the generalization performance for a single trail, we derive the optimal learning rate *in probability* as follows:

Theorem 2. *Under the same assumptions as Theorem 1. If $\lambda = \Omega(|\mathcal{D}|^{-\frac{1}{2}})$, $|\mathcal{D}_1| = \ldots |\mathcal{D}_m|$, $m \lesssim |\mathcal{D}|^{\frac{1}{4}}$ and $M \gtrsim |\mathcal{D}|^{\frac{1}{2}}$, then, for every $\delta \in (0, 1]$, with probability at least $1 - \delta$, we have*

$$\left\|\bar{f}^0_{M,\mathcal{D},\lambda} - f_{\mathcal{H}_K}\right\|^2_\rho = \mathcal{O}\left(|\mathcal{D}|^{-\frac{1}{2}} \log^2(1/\delta)\right).$$

To guarantee the optimal generalization properties in probability, the number of partitions should be restricted to $\Omega(|\mathcal{D}|^{1/4})$, which is stricter than $\Omega(|\mathcal{D}|^{1/2})$ in Theorem 1. This is because the generalization error in expectation can be decomposed into approximation error, sample error and distributed error (more details can be seen in Proposition 1 in Appendix), but the error decomposition in probability is not easy to separate a distributed error in probability to control the number of local machines. The derive the optimal learning rate, we provide a novel decomposition, please see details in Proposition 9 in Appendix.

The following result demonstrates that the proposed communication strategy can enlarge the number of partitions in probability.

Theorem 3. *Under the same assumptions as Theorem 1, If $\lambda = \Omega(|\mathcal{D}|^{-\frac{1}{2}})$, $|\mathcal{D}_1| = \ldots |\mathcal{D}_m|$, $m \lesssim |\mathcal{D}|^{\frac{p+1}{2(p+2)}}$ and $M \gtrsim |\mathcal{D}|^{\frac{1}{2}}$, then, for every $\delta \in (0, 1]$, with probability at least $1 - \delta$, we have*

$$\left\|\bar{f}^p_{M,\mathcal{D},\lambda} - f_{\mathcal{H}_K}\right\|^2_\rho = \mathcal{O}\left(|\mathcal{D}|^{-\frac{1}{2}} \log^{p+2}(1/\delta)\right),$$

where $\bar{f}^p_{M,\mathcal{D},\lambda}$ is returned by Algorithm 1 under p-th iterations.

Compared Theorem 3 with Theorem 2, it is clear that the proposed communication strategy can relax the restriction on m from $\Omega\left(|\mathcal{D}|^{1/4}\right)$ to $\Omega\left(|\mathcal{D}|^{(p+1)/(2(p+2))}\right)$. Note that m is monotonically increasing with the number of communications p, which can demonstrate the power of the proposed

Table 1: Computational complexity required by different algorithms for the optimal learning rate $\mathcal{O}\big(1/\sqrt{|\mathcal{D}|})\big)$ in the basic setting. Logarithmic terms are not showed.

Methods	Partitions m	Random M	Types	Space	Time	Communication
KRR (Caponnetto & Vito, 2007)	/	/	In probability	$\|\mathcal{D}\|^2$	$\|\mathcal{D}\|^3$	/
KRR-RF (Rudi & Rosasco, 2017)	/	$\|\mathcal{D}\|^{0.5}$	In probability	$\|\mathcal{D}\|^{1.5}$	$\|\mathcal{D}\|^2$	/
KRR-Nystöm Rudi et al. (2015)	/	$\|\mathcal{D}\|^{0.5}$	In probability	$\|\mathcal{D}\|^{1.5}$	$\|\mathcal{D}\|^2$	/
DKRR (Zhang et al., 2015; Chang et al., 2017b)	$\|\mathcal{D}\|^{0.5}$	/	In expectation	$\|\mathcal{D}\|$	$\|\mathcal{D}\|^2$	$\|\mathcal{D}\|^{0.5}$
DKRR (Lin et al., 2020)	$\|\mathcal{D}\|^{0.25}$	/	In probability	$\|\mathcal{D}\|^{1.5}$	$\|\mathcal{D}\|^{2.25}$	$\|\mathcal{D}\|^{0.75}$
DKRR-CM (Lin et al., 2020)	$\|\mathcal{D}\|^{\frac{p+1}{2(p+2)}}$	/	In probability	$\|\mathcal{D}\|^{\frac{p+3}{p+2}}$	$\|\mathcal{D}\|^{\frac{3(p+3)}{2(p+2)}}$	$pd\|\mathcal{D}\|$
DKRR-RF (Li et al., 2019a)	$\Omega(1)$	$\|\mathcal{D}\|^{0.5}$	In expectation	$\|\mathcal{D}\|$	$\|\mathcal{D}\|^2$	$\|\mathcal{D}\|^{0.5}$
DKRR-RF (Theorem 1)	$\|\mathcal{D}\|^{0.5}$	$\|\mathcal{D}\|^{0.5}$	In expectation	$\|\mathcal{D}\|$	$\|\mathcal{D}\|^{1.5}$	$\|\mathcal{D}\|^{0.5}$
DKRR-RF (Theorem 2)	$\|\mathcal{D}\|^{0.25}$	$\|\mathcal{D}\|^{0.5}$	In probability	$\|\mathcal{D}\|^{1.25}$	$\|\mathcal{D}\|^{1.75}$	$\|\mathcal{D}\|^{0.5}$
DKRR-RF-CM (Theorem 3)	$\|\mathcal{D}\|^{\frac{p+1}{2(p+2)}}$	$\|\mathcal{D}\|^{0.5}$	In probability	$\|\mathcal{D}\|^{\frac{2p+5}{2p+4}}$	$\|\mathcal{D}\|^{\frac{3p+7}{2p+4}}$	$p\|\mathcal{D}\|^{0.5}$

communications. When $p \to \infty$, the partitions can reach $\Omega(\sqrt{|\mathcal{D}|})$, which is the same as the generalization bound in expectation.

Remark 2. *In the main text of this paper, we only give the optimal rates of DKRR-RF in the basic setting. The fast learning rates can be achieved under favorable conditions, see in Appendix.*

4.3 Compared with the Related Work

4.3.1 Comparisons of the Time and Space Complexities

Table 1 reports the statistical and computational properties of the related approaches and our theoretical findings under the basic setting. We see that our DKRR-RF can guarantee the optimal generalization performance in expectation only requiring $|\mathcal{D}|$ memory and $|\mathcal{D}|^{1.5}$ time, which is more effective than other methods. For DKRR-RF-CM, we can also see that it can guarantee optimal generalization performance in probability requiring less complexity than the communication-based method of DKRR-CM (Lin et al., 2020).

Remark 3. *In (Rudi et al., 2017), the authors considered combining the Nyström method and preconditioned conjugate gradient (PCG) (Cutajar et al., 2016) to scale up KRR. As far as we know, it is the only existing work that can guarantee optimal statistical accuracy, only requiring $|\mathcal{D}|$ memory and $|\mathcal{D}|^{1.5}$ time for KRR. In this paper, we consider combining distributed learning and random features, a completely different path from (Rudi et al., 2017). Note that in our proposed method, we need to compute the inverse of $\mathbf{\Phi}_{M,\mathcal{D}_j}\mathbf{\Phi}_{M,\mathcal{D}_j}^{\mathrm{T}} + \lambda\mathbf{I}$, which requires $|\mathcal{D}|^{1.5}$ time. Inspired by (Rudi et al., 2017), we can also adopt PCG to avoid the inverse calculation, which can further speed up our proposed DKRR-RF-CM. The combination of DKRR-RF-CM and PCG may open a path to reach the* ***linear time complexity*** *for optimal learning rate.*

4.3.2 Novelty and Proof Techniques

The most related works of our paper are (Li et al., 2019a), (Lin et al., 2020) and (Rudi & Rosasco, 2017). We discuss the novel techniques adopted to derive the improved results compared with them.

Compared with (Li et al., 2019a). (a) To derive the learning bounds, $\mathcal{Q}_{M,\mathcal{D}} := \|(C_M + \lambda I)^{-1/2}(C_M - C_{M,\mathcal{D}})(C_M + \lambda I)^{-1/2}\|$ is required to be estimated, where C_M and $C_{M,\mathcal{D}}$ are self-adjoint and positive operators defined in Definition 1 (see in Appendix). In (Li et al., 2019a), they used a classical approach from (Chang et al., 2017b; Guo et al., 2017) to estimate $\mathcal{Q}_{M,\mathcal{D}}$ (see Lemmas 21 and 22 in (Li et al., 2019a)), and obtain that

$$\mathcal{Q}_{M,\mathcal{D}} \leq \frac{1}{\sqrt{\lambda}}\|(C_M - C_{M,\mathcal{D}})(C_M + \lambda I)^{-1/2}\| = \mathcal{O}\left(\frac{1}{\lambda|\mathcal{D}|} + \sqrt{\frac{\mathcal{N}(\lambda)}{\lambda|\mathcal{D}|}}\right),$$

where $\mathcal{N}(\lambda)$ is the effective dimension defined in Assumption 1 (see in Appendix). However, in our paper, we directly estimate $\mathcal{Q}_{M,\mathcal{D}}$ based on the concentration inequality for self-adjoint operators (Caponnetto & Yao, 2006), and prove that

$$\mathcal{Q}_{M,\mathcal{D}} = \mathcal{O}\left(\frac{1}{\lambda|\mathcal{D}|} + \sqrt{\frac{1}{\lambda|\mathcal{D}|}}\right) \text{ (see in Proposition 6).}$$

Thus, our estimation of $\mathcal{Q}_{M,\mathcal{D}}$ is $\sqrt{\mathcal{N}(\lambda)}$ tighter than that in (Li et al., 2019a). This is the one of the key reasons why we can substantially relax the restriction on the number of local machines compared with (Li et al., 2019a); (b) We not only present the bounds in expectation, but also in probability. To derive the tight bound in probability, we provide new decompositions of $\|\bar{f}^0_{M,\mathcal{D},\lambda} - f_{M,\mathcal{D},\lambda}\|$ and $\|\bar{\mathbf{w}}^0_{M,\mathcal{D},\lambda} - \mathbf{w}_{M,\mathcal{D},\lambda}\|$, please see Proposition 9 in detail. As far as we know, these decompositions are novel; (c) We also consider the Newton Raphson iteration-based communication strategy. To derive the improved high-probability bounds with communication, we introduce a novel decomposition of $\|\bar{f}^t_{M,\mathcal{D},\lambda} - f_{M,\mathcal{D},\lambda}\|_\rho$ (see in Proposition 10).

Compared with (Lin et al., 2020). (a) In (Lin et al., 2020), they also considered a communication strategy to enlarge the number of local machines. However, the communications strategy proposed in (Lin et al., 2020) is based on an operator representation, which is usually infeasible in practice. The authors presented a realization for the proposed strategy by communicating the data among each local machine. Thus, the data privacy of each local machine is difficult to be protected. Furthermore, since it requires communicating data $\mathcal{D}_j$, $j = 1, \ldots, m$, among each local machine, for each iteration, the communication complexity of each local machine is $\mathcal{O}(|\mathcal{D}|d)$, which is too high for large scale data sets. However, the communication strategy proposed in this paper only requires communicating the gradient $g_{M,\mathcal{D}_j,\lambda}(\bar{\mathbf{w}}^{t-1}_{M,\mathcal{D},\lambda})$ and the model parameters $\boldsymbol{\beta}^{t-1}_j$, rather than the data, therefore our proposed strategy do better on privacy protection. Moreover, the communication complexity is only $\mathcal{O}(M)$ for each local machine, $M \ll |\mathcal{D}|$, which is suitable for large scale data sets; (b) At first it seems that the proof techniques of (Lin et al., 2020) can be easily extended to our paper, but it is not true. If we use the same proof technical of (Lin et al., 2020), we can only obtain that $\|f^\diamond_{M,\mathcal{D},\lambda} - f_{M,\lambda}\|_\rho = \|(L_{M,\mathcal{D}} + \lambda I)^{-1}(L_{M,\mathcal{D}} - L_M)(f_\rho - f_{M,\lambda})\| = \mathcal{O}((1/(\sqrt{\lambda}|\mathcal{D}|) + \sqrt{\mathcal{N}(\lambda)/|\mathcal{D}|})\|f_\rho - f_{M,\lambda}\|_\rho/\sqrt{\lambda})$, where $f^\diamond_{M,\mathcal{D},\lambda}$, $f_{M,\lambda}$ and f_ρ are defined in Definition 14. Combing with Proposition 2, one can only obtain that

$$\|f_{M,\mathcal{D},\lambda} - f_{M,\lambda}\|_\rho = \mathcal{O}\left(\left(\frac{1}{\lambda|\mathcal{D}|} + \sqrt{\frac{\mathcal{N}(\lambda)}{\lambda|\mathcal{D}|}}\right)\|f_\rho - f_{M,\lambda}\|_\rho\right).$$

However, in our paper, we introduce new decompositions of $\|f^\diamond_{M,\mathcal{D},\lambda} - f_{M,\lambda}\|_\rho$ and $\|f_{M,\mathcal{D},\lambda} - f^\diamond_{M,\mathcal{D},\lambda}\|_\rho$ (see Propositions 2 and 3 in detail), and further obtain that

$$\|f_{M,\mathcal{D},\lambda} - f_{M,\lambda}\|_\rho = \mathcal{O}\left(\left(\frac{1}{\sqrt{\lambda}|\mathcal{D}|} + \sqrt{\frac{\mathcal{N}(\lambda)}{|\mathcal{D}|}}\right)\|f_\rho - f_{M,\lambda}\|_\rho\right) \quad \text{(see in Proposition 4)},$$

which is $\Omega(1/\sqrt{\lambda})$ tighter than the directly use of the techniques of (Lin et al., 2020). The novel decompositions $\|f^\diamond_{M,\mathcal{D},\lambda} - f_{M,\lambda}\|_\rho$ and $\|f_{M,\mathcal{D},\lambda} - f^\diamond_{M,\mathcal{D},\lambda}\|_\rho$ are the key reasons why we can guarantee the optimal performance even under $m = \Omega(\sqrt{|\mathcal{D}|})$. We only give an example here, but the novel decompositions have also been embedded in Propositions 5, 9, 10, etc.

Compared with (Rudi & Rosasco, 2017). (a) We study the statistical properties of the **combination** of distributed learning and random features, but in (Rudi & Rosasco, 2017), they only consider random features. As mentioned above, to estimate the tight bound of the distributed error, we introduce a novel decomposition of $\|\bar{f}^0_{M,\mathcal{D},\lambda} - f_{M,\mathcal{D},\lambda}\|$ to derive the tight bound in expectation, and novel decompositions of $\|\bar{f}^0_{M,\mathcal{D},\lambda} - f_{M,\mathcal{D},\lambda}\|$ and $\|\bar{f}^p_{M,\mathcal{D},\lambda} - f_{M,\mathcal{D},\lambda}\|$ to derive tight bounds in probability; (b) The combination not only brings distributed error, but also brings some other problems. If you compare the proofs of (Rudi & Rosasco, 2017) with these of our paper in detail, you can find that the decompositions of (Rudi & Rosasco, 2017) and ours are very different.

Overall, in this paper, we improve the existing state-of-art bounds in expectation, and provide novel (communication-based) distributed bounds with random features in probability. Moreover, we introduce some novel techniques and decompositions to substantially relax the restriction on the number of local machines, which are non-trivial extensions of (Li et al., 2019a; Lin et al., 2020; Rudi & Rosasco, 2017).

5 EXPERIMENTS

In this section, we validate our theoretical findings by performing experiments on both simulated and real datasets.

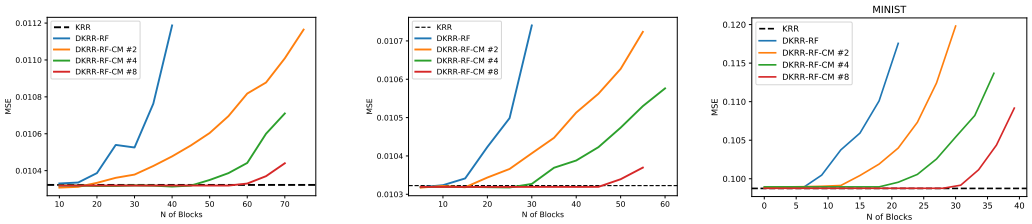

(a) simulated data, $q = 2$, $t = 6$ (b) simulated data, $q = 2$, $t = 5$ (c) real data, MINIST

Figure 1: The mean square error on the test set with different partitions on KRR, DKRR-RF and our DKRR-RF-CM. # represents the number of communications.

Numerical Experiments. Inspired by numerical experiments in (Rudi & Rosasco, 2017; Li et al., 2019b), we consider a spline kernel of order q:

$$K_{2q}(\mathbf{x}, \mathbf{x}') = 1 + \sum_{k=1}^{\infty} \cos(2\pi k(\mathbf{x} - \mathbf{x}'))/(k^{2q}).$$

If the marginal distribution of $\mathcal{X}$ is uniform on [0,1], then $K_{2q}(\mathbf{x}, \mathbf{x}') = \int_0^1 \psi(\mathbf{x}, \boldsymbol{\omega})\psi(\mathbf{x}, \boldsymbol{\omega})\varrho(\boldsymbol{\omega})d\boldsymbol{\omega}$, where $\psi(\mathbf{x}, \boldsymbol{\omega}) = K_q(\mathbf{x}, \boldsymbol{\omega})$ and $\varrho(\boldsymbol{\omega})$ is also uniform on [0,1]. The random features of the spline kernel are $\boldsymbol{\phi}_M(\mathbf{x}) = (\psi(\mathbf{x}, \boldsymbol{\omega}_1), \ldots, \psi(\mathbf{x}, \boldsymbol{\omega}_M))^{\mathrm{T}}/\sqrt{M}$. According to Theorem 1, 2 and 3, we set the size of the random features to be $M = \sqrt{|\mathcal{D}|}$, and fine tune λ around $|\mathcal{D}|^{-1/2}$ using 5-fold cross validation, the tuned set is $\{2^{-5}, 2^{-3}, \ldots 2^5\}|\mathcal{D}|^{-1/2}$. We let the target function f_* be a Gaussian random variable with mean $\mu = K_t(x, 0)$ and variance $\sigma^2 = 0.01$.

We generate 10000 samples for training and 10000 samples for testing. We use the exact KRR as a baseline, which trains all samples in a batch. We compare our proposed DKRR-RF-CM ($p = 2, 4, 8$) with KRR and DKRR-RF. We repeat the training 5 times and estimate the averaged error on testing data. The mean square error on the test set with different partitions is given in Figure 1(a, b), which can be summarized as follows: 1) When m is not too large, the distributed methods (DRKK-RF and DRKK-RF-CM) are always comparable to original KRR. There exists an upper bound of m, when larger than it, the error increases dramatically and is far from the original KRR. This verifies the theoretical statement in Theorem 1, 2 and 3; 2) The upper bound m of DKRR-RF-CM is much larger than DKRR-RF. This result is aligned with Theorem 3, which demonstrates that the proposed communication strategy can enlarge the upper bound m; 3) The upper bound m of DKRR-RF-CM monotonically increases with the number of communications, which verifies Theorem 3.

Real Data. In this experiment, we consider the performance on real data. The empirical evaluations with Gaussian kernel, $\exp(-\|\mathbf{x} - \mathbf{x}'\|^2/\sigma)$, on MINIST (random selected 10000 samples for training and 10000 for testing) is given in Figure 1(c), where the optimal σ is selected by 5-fold cross-validation, $\sigma \in \{2^i, i = -10, -8, \ldots, 10\}$, and the number of random features is 300. From Figure 1(c), one can find that the result is similar as that on the simulated data, which demonstrates that our proposed communication strategy is also effective on the real dataset.

6 Conclusion

In this paper, we study the generalization properties of the combination of distributed learning and random features for ridge regression. We first improve the existing results of divide-and-conquer KRR with random features in expectation. Then, beyond the expectation, we derive generalization error bound in probability. Finally, we propose a novel effective communication strategy to further improve the learning performance of the combination method, and demonstrate the power of communications via both theoretical assessments and numerical experiments. Our results may open several venues for both theoretical and empirical work: (a) combine the approach with gradient algorithms such as preconditioned conjugate gradient Avron et al. (2017a) and multi-pass SGD (Carratino et al., 2018); (b) replace synchronous distributed methods with asynchronous ones (Suresh et al., 2017); (c) consider the loss functions other than quadratic loss (Li et al., 2019b).

REFERENCES

Ahmed Alaoui and Michael Mahoney. Fast randomized kernel ridge regression with statistical guarantees. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 775–783, 2015.

Miguel Arcones. A Bernstein-type inequality for U-statistics and U-processes. *Statistics & probability letters*, 22(3):239–247, 1995.

Haim Avron, Kenneth Clarkson, and David Woodruff. Faster kernel ridge regression using sketching and preconditioning. *SIAM Journal on Matrix Analysis and Applications*, 38(4):1116–1138, 2017a.

Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, and Amir Zandieh. Random fourier features for kernel ridge regression: Approximation bounds and statistical guarantees. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pp. 253–262, 2017b.

Francis Bach. Sharp analysis of low-rank kernel matrix approximations. In *Proceedings of the 26th Annual Conference on Learning Theory Conference on Learning Theory (COLT)*, pp. 185–209, 2013.

Francis Bach. On the equivalence between kernel quadrature rules and random fekernel expansions. *Journal of Machine Learning Research*, 21:1–38, 2017.

Gilles Blanchard and Nicole Krämer. Optimal learning rates for kernel conjugate gradient regression. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 226–234, 2010.

Gilles Blanchard and Nicole Krämer. Convergence rates of kernel conjugate gradient for random design regression. *Analysis and Applications*, 14(6):763–794, 2016.

Raffaello Camoriano, Tomás Angles, Alessandro Rudi, and Lorenzo Rosasco. Nytro: When subsampling meets early stopping. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 1403–1411, 2016.

Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.

Andrea Caponnetto and Yuan Yao. Adaptation for regularization operators in learning theory. Technical report, Technical Report CBCL Paper# 265/AI Technical Report, 2006.

Luigi Carratino, Alessandro Rudi, and Lorenzo Rosasco. Learning with SGD and random features. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 10192–10203, 2018.

Xiangyu Chang, Shao-Bo Lin, and Yao Wang. Divide and conquer local average regression. *Electronic Journal of Statistics*, 11(1):1326–1350, 2017a.

Xiangyu Chang, Shao-Bo Lin, and Ding-Xuan Zhou. Distributed semi-supervised learning with kernel ridge regression. *The Journal of Machine Learning Research*, 18(1):1493–1514, 2017b.

Corinna Cortes, Mehryar Mohri, and Ameet Talwalkar. On the impact of kernel approximation on learning accuracy. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 113–120, 2010.

Kurt Cutajar, Michael Osborne, John Cunningham, and Maurizio Filippone. Preconditioning kernel matrices. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pp. 2529–2538, 2016.

Lizhong Ding and Shizhong Liao. Nyström approximate model selection for LSSVM. In *Advances in Knowledge Discovery and Data Mining - 16th Pacific-Asia Conference (PAKDD)*, pp. 282–293, 2012.

Petros Drineas and Michael W Mahoney. On the Nyström method for approximating a Gram matrix for improved kernel-based learning. *Journal of Machine Learning Research*, 6:2153–2175, 2005.

Zheng-Chu Guo, Shao-Bo Lin, and Ding-Xuan Zhou. Learning theory of distributed spectral algorithms. *Inverse Problems*, 33(7):074009, 2017.

Cho-Jui Hsieh, Si Si, and Inderjit Dhillon. A divide-and-conquer solver for kernel support vector machines. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pp. 566–574, 2014.

Arthur Jacot, Berfin Şimşek, Francesco Spadaro, Clément Hongler, and Franck Gabriel. Implicit regularization of random feature models. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.

Galyna Kriukova, Sergiy Pereverzyev, and Pavlo Tkachenko. Nyström type subsampling analyzed as a regularized projection. *Inverse Problems*, 33(7):074001, 2017.

Jian Li, Yong Liu, and Weiping Wang. Distributed learning with random features. *arXiv preprint arXiv:1906.03155*, 2019a.

Zhu Li, Jean-Franois Ton, Dino Oglic, and Dino Sejdinovic. Towards a unified analysis of random fourier features. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pp. 3905–3914, 2019b.

Junhong Lin and Lorenzo Rosasco. Optimal rates for multi-pass stochastic gradient methods. *The Journal of Machine Learning Research*, 18(1):3375–3421, 2017.

Shao-Bo Lin and Ding-Xuan Zhou. Distributed kernel-based gradient descent algorithms. *Constructive Approximation*, 47(2):249–276, 2018.

Shao-Bo Lin, Xin Guo, and Ding-Xuan Zhou. Distributed learning with regularized least squares. *The Journal of Machine Learning Research*, 18(1):3202–3232, 2017.

Shao-Bo Lin, Di Wang, and Ding-Xuan Zhou. Distributed kernel ridge regression with communications. *arXiv preprint arXiv:2003.12210*, 2020.

Fanghui Liu, Xiaolin Huang, Yudong Chen, and Johan AK Suykens. Random features for kernel approximation: A survey in algorithms, theory, and beyond. *arXiv preprint arXiv:2004.11154*, 2020.

Mona Meister and Ingo Steinwart. Optimal learning rates for localized svms. *The Journal of Machine Learning Research*, 17(1):6722–6765, 2016.

Cameron Musco and Christopher Musco. Recursive sampling for the Nyström method. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3833–3845, 2017.

Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1177–1184, 2007.

Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1313–1320, 2008.

Dominic Richards, Patrick Rebeschini, and Lorenzo Rosasco. Decentralised learning with random features and distributed gradient descent. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.

Lorenzo Rosasco, Mikhail Belkin, and Ernesto De Vito. On learning with integral operators. *Journal of Machine Learning Research*, 11:905–934, 2010.

Alessandro Rudi and Lorenzo Rosasco. Generalization properties of learning with random features. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3215–3225, 2017.

Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco. Less is more: Nyström computational regularization. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1648–1656, 2015.

Alessandro Rudi, Luigi Carratino, and Lorenzo Rosasco. Falkon: An optimal large scale kernel method. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3888–3898, 2017.

Alessandro Rudi, Daniele Calandriello, Luigi Carratino, and Lorenzo Rosasco. On fast leverage score sampling and optimal learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 5672–5682, 2018.

John Shawe-Taylor and Nello Cristianini. *An Introduction to Support Vector Machines and other Kernel-based Learning Methods*. Cambridge University Press, 2000. ISBN 0521780195.

Si Si, Cho-Jui Hsieh, and Inderjit Dhillon. Computationally efficient nyström approximation using fast transforms. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pp. 2655–2663, 2016.

Steve Smale and Ding-Xuan Zhou. Shannon sampling ii: connections to learning theory. *Applied and Computational Harmonic Analysis*, 19(3):285–302, 2005.

Steve Smale and Ding-Xuan Zhou. Learning theory estimates via integral operators and their approximations. *Constructive Approximation*, 26(2):153–172, 2007.

Ingo Steinwart and Andreas Christmann. *Support Vector Machines*. Springer Verlag, New York, 2008.

Yitong Sun, Anna Gilbert, and Ambuj Tewari. But how does it work in theory? linear SVM with random features. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3379–3388, 2018.

Ananda Theertha Suresh, Felix Yu, Sanjiv Kumar, and Brendan McMahan. Distributed mean estimation with limited communication. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pp. 3329–3337, 2017.

Vladimir Vapnik. *The Nature of Statistical Learning Theory*. Springer Verlag, 2000.

Tianbao Yang, Yu-Feng Li, Mehrdad Mahdavi, Rong Jin, and Zhi-Hua Zhou. Nyström method vs random Fourier features: A theoretical and empirical comparison. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 476–484, 2012.

Rong Yin, Yong Liu, Lijing Lu, Weiping Wang, and Dan Meng. Divide-and-conquer learning with Nyström: optimal rate and algorithm. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 6696–6703, 2020.

Felix Xinnan X Yu, Ananda Theertha Suresh, Krzysztof M Choromanski, Daniel N Holtmann-Rice, and Sanjiv Kumar. Orthogonal random features. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1975–1983, 2016.

Yuchen Zhang, John Duchi, and Martin Wainwright. Divide and conquer kernel ridge regression. In *Proceedings of the 26th Annual Conference on Learning Theory (COLT)*, pp. 592–617, 2013.

Yuchen Zhang, John Duchi, and Martin Wainwright. Divide and conquer kernel ridge regression: A distributed algorithm with minimax optimal rates. *The Journal of Machine Learning Research*, 16 (1):3299–3340, 2015.