

Stability-based Generalization Analysis of Randomized Coordinate Descent for Pairwise Learning

Liang Wu^{1, 2}, Ruixi Hu^{1*}, Yunwen Lei³

¹Center of Statistical Research, School of Statistics, Southwestern University of Finance and Economics, Chengdu, China

²Big Data Laboratory on Financial Security and Behavior, SWUFE (Laboratory of Philosophy and Social Sciences, Ministry of Education), Chengdu, China

³Department of Mathematics, University of Hong Kong, Pok Fu Lam, Hong Kong
wuliang@swufe.edu.cn, hu_ruixi@163.com, leiyw@hku.hk

Abstract

Pairwise learning includes various machine learning tasks, with ranking and metric learning serving as the primary representatives. While randomized coordinate descent (RCD) is popular in various learning problems, there is much less theoretical analysis on the generalization behavior of models trained by RCD, especially under the pairwise learning framework. In this paper, we consider the generalization of RCD for pairwise learning. We measure the on-average argument stability for both convex and strongly convex objective functions, based on which we develop generalization bounds in expectation. The early-stopping strategy is adopted to quantify the balance between estimation and optimization. Our analysis further incorporates the low-noise setting into the excess risk bound to achieve the optimistic bound as $O(1/n)$, where n is the sample size.

Introduction

The paradigm of pairwise learning has found wide applications in machine learning. Several popular examples are shown as the following. In ranking, we aim to find a model that can predict the ordering of instances (Cléménçon, Lugosi, and Vayatis 2008; Rejchel 2012). In metric learning, we wish to build a model to measure the distance between instances (Cao, Guo, and Ying 2016; Ye, Zhan, and Jiang 2019; Dong et al. 2020). Besides, various problems such as AUC maximization (Cortes and Mohri 2003; Gao et al. 2013; Ying, Wen, and Lyu 2016; Liu et al. 2018) and learning tasks with minimum error entropy (Hu et al. 2015) can also be formulated as this paradigm. For all these pairwise learning tasks, the performance of models needs to be measured on pairs of instances. In contrast to pointwise learning, this paradigm is characterized by pairwise loss functions $f : \mathcal{H} \times \mathcal{Z} \times \mathcal{Z} \mapsto \mathbb{R}$, where \mathcal{H} and \mathcal{Z} denote the hypothesis space and the sample space respectively. To understand and apply the paradigm better, there is a growing interest in the study under the uniform framework of pairwise learning.

Randomized coordinate descent (RCD) is one of the most commonly used first-order methods in optimization. In each iteration, RCD updates a randomly chosen coordinate along

the negative direction of the gradient and keeps other coordinates unchanged. This makes RCD especially effective for large-scale problems (Nesterov 2012), where the computational cost is rather hard to handle.

The extensive applications of RCD have motivated some interesting theoretical analysis on its empirical behavior (Nesterov 2012; Richtárik and Takáč 2014; Beck and Teboulle 2021; Chen, Li, and Lu 2023), which focuses on iteration complexities and empirical risks in the optimization process. However, there is much less work considering the generalization performance of RCD, i.e., how models trained by RCD would behave on testing samples. It is notable that the relative analysis only considers the case of pointwise learning (Wang, Wu, and Lei 2021), which is different from pairwise learning in the structure of loss functions. Besides, this work fails to establish the generalization bound based on the ℓ_2 on-average argument stability in a strongly convex case. Therefore, the existing theoretical analysis of RCD is not enough to describe the discrepancy between training and testing for pairwise learning. How to quantify the balance between statistics and optimization under this setting still remains a challenge. In this paper, we develop a more systematical and fine-grained generalization analysis of RCD for pairwise learning to refine the above study. Our analysis can lead to a more appropriate design for the optimization algorithm and the machine learning model.

In this paper, we present the generalization analysis based on the concept of algorithmic stability (Bousquet and Elisseeff 2002). The comparison between the existing work and this paper is presented in Table 1. Our contributions are summarized as follows.

1. Under general assumptions on L -smoothness of loss functions, coordinate-wise smoothness and convexity of objective functions, we study the ℓ_2 on-average argument stability and the corresponding generalization bounds of RCD for pairwise learning. To achieve optimal performance, we consider the balance between the generalization error and the optimization error. The result shows that the early stopping strategy is beneficial to the generalization. The excess risk bounds enjoy the order of $O(1/\sqrt{n})$ and $O(\sqrt{\log(n)}/n)$ for convex and strongly-convex objective functions respectively, where n denotes the sample size and σ is the strong-convexity parameter.

2. We use the low noise condition $F(\mathbf{w}^*) = O(1/n)$ to

*The corresponding author is Ruixi Hu.

Learning Paradigm	Algorithm	Reference	Assumption	Noise Seeting	Iteration	Rate	
Pointwise Learning	SGD	Lei (2020)	C, L -S	\searrow	$T \asymp n$	$O(1/\sqrt{n})$	
				$F(\mathbf{w}^*) = 0$		$O(1/n)$	
			C, G -Lip, L -S	\searrow		$O(1/\sqrt{n})$	
			SC, G -Lip, L -S			$O(1/n\sigma)$	
Pairwise Learning		Lei (2021)	C, L -S	$F(\mathbf{w}^*) = O(1/n)$			$O(1/\sqrt{n})$
			C, G -Lip	\searrow		$T \asymp n^2$	$O(1/\sqrt{n})$
			SC, L -S			$T \asymp n$	$O(1/n\sigma)$
			SC, G -Lip			$T \asymp n^2$	$O(1/n\sigma)$
	RCD	This Work	C, L -S, Lip-grad	$F(\mathbf{w}^*) = O(1/n)$	$T \asymp \sqrt{n}$	$O(1/\sqrt{n})$	
					$T \asymp n$	$O(1/n)$	
			SC, L -S, Lip-grad	\searrow	$T \asymp \log(n)$	$O(\sqrt{\log(n)}/n)$	

Table 1: All the above convergence rates are based on excess risk bounds in expectation. C means the convexity as Assumption 4 and SC means the strong convexity as Assumption 5. G -Lip refers to Assumption 1, L -S refers to Assumption 2 and Lip-grad refers to Assumption 3. Furthermore, Lei (2020) refers to Lei and Ying (2020), Lei (2021) refers to Lei, Liu, and Ying (2021).

develop sharper generalization bounds under the convex case. This motivates the excess risk bound $O(1/n)$, which matches the approximate optimal rate under the strongly convex case. However, we should note that the approximate optimal rate is accessible with a faster computing $T \asymp \log(n)$ for strongly convex empirical risks.

The main work is organized according to the convexity of the empirical risk. We consider the on-average argument stability and develop the corresponding excess risk bounds. The early-stopping strategy is useful for balancing optimization and estimation, by which we present the optimal convergence rate. Furthermore, there are two key points about our proof in comparison with the pointwise case (Wang, Wu, and Lei 2021): One is applying the coercivity property to bound the expansiveness of RCD updates since the expectation of randomized coordinates leads to the gradient descent operator. The other is following the optimization error bounds for pointwise learning directly since they both use unbiased gradient estimations.

Related Work

In this section, we review the related work on RCD and generalization analysis for pairwise learning.

Randomized Coordinate Descent (RCD). The real-world performance of RCD has demonstrated its significant efficiency in many large-scale optimization tasks, including regularized risk minimization (Chang, Hsieh, and Lin 2008; Shalev-Shwartz and Tewari 2009), low-rank matrix completion and learning (Hu and Kwok 2019; Callahan, Vu, and Raich 2024), and optimal transport problems (Xie, Wang, and Zhang 2024). The convergence analysis of RCD and its accelerated variant was first proposed by Nesterov (2012), where global estimates of the convergence rate were considered. Then the strategies to accelerate RCD were further explored (Richtárik and Takáč 2014), for which the corresponding convergence properties were established for structural optimization problems (Zhao et al. 2014; Lu and Xiao 2015, 2017). RCD was also studied under various settings including nonconvex optimization (Beck and Teboulle 2021;

Chen, Li, and Lu 2023), volume sampling (Rodomanov and Kropotov 2020) and differential privacy (Damaskinos et al. 2021). The above study mainly considered the empirical behavior of RCD. However, the aim of this paper is to quantify the generalization performance of machine learning models trained by RCD.

Generalization for Pairwise Learning. The generalization ability shows how models based on training datasets will adapt to testing datasets. It serves as an important indicator for the enhancement of models and algorithms in the view of statistical learning theory (SLT). To investigate the generalization performance for pairwise learning, methods of uniform convergence analysis and stability analysis have been applied under this wide learning framework. More details are described below.

The uniform convergence approach considers the connection between generalization errors and U-statistics, from which generalization bounds via corresponding U-processes are developed sufficiently. Complexity measures including VC dimension (Vapnik, Levin, and Le Cun 1994), covering numbers (Zhou 2002) and Rademacher complexities (Bartlett and Mendelson 2001) for the hypothesis space play a key role in this approach. For pairwise learning, these measures have been used for studying the generalization of specific tasks such as ranking (Cléménçon, Lugosi, and Vayatis 2008; Rejchel 2012) and metric learning (Cao, Guo, and Ying 2016; Ye, Zhan, and Jiang 2019; Dong et al. 2020). Recently, some work also explored the generalization of deep networks with these tasks (Huang et al. 2023; Zhou, Wang, and Zhou 2024). Furthermore, generalizations for the pairwise learning framework were studied under various settings, including PL condition (Lei, Liu, and Ying 2021), regularized risk minimization (Lei, Ledent, and Kloft 2020) and online learning (Wang et al. 2012; Kar et al. 2013). As compared to the stability analysis, the complexity analysis enjoys the ability of yielding generalization bounds for non-convex objective functions (Mei, Bai, and Montanari 2018; Davis and Drusvyatskiy 2022). However, generalization bounds yielded by the uniform convergence approach are inevitably associated with input dimensions (Agarwal

and Niyogi 2009; Feldman 2016; Schliserman, Sherman, and Koren 2024), which can be avoided in the stability analysis.

Algorithmic stability serves as an important concept in SLT, which is closely related to learnability and consistency (Feldman 2016; Rakhlin, Mukherjee, and Poggio 2005). The basic framework for stability analysis was proposed by Bousquet and Elisseeff (2002), where the concept of uniform stability was introduced and then extended to study randomized algorithms (Elisseeff et al. 2005). The power of algorithmic stability for generalization analysis further inspired several other stability measures including uniform argument stability (Liu et al. 2017), on-average loss stability (Shalev-Shwartz et al. 2010; Lei, Ledent, and Kloft 2020; Lei, Liu, and Ying 2021), on-average argument stability (Lei and Ying 2020; Deora et al. 2024), locally elastic stability (Deng, He, and Su 2021; Lei, Sun, and Liu 2023) and Bayes stability (Li, Luo, and Qiao 2020). While various stability measures were useful for deriving generalization bounds in expectation, applications of uniform stability implied elegant high-probability generalization bounds (Feldman and Vondrak 2019; Bousquet, Klochkov, and Zhivotovskiy 2020; Klochkov and Zhivotovskiy 2021). Furthermore, the stability analysis promoted the study for the generalization of stochastic gradient descent (SGD) effectively (Deng et al. 2023), which was considered under the paradigm of pairwise learning (Lei, Ledent, and Kloft 2020; Lei, Liu, and Ying 2021) or pointwise and pairwise learning (Wang et al. 2023; Chen et al. 2023). In contrast to SGD, a more sufficient generalization analysis of RCD is needed under the framework of pairwise learning. It provides us guidelines to apply RCD in large-scale optimization problems for pairwise learning.

Other than the approach based on uniform convergence or algorithmic stability, the generalization for pairwise learning was also studied from the perspective of algorithmic robustness (Bellet and Habrard 2015; Christmann and Zhou 2016), convex analysis (Ying and Zhou 2016), integral operators (Fan et al. 2016; Guo et al. 2017) and information theoretical analysis (Dong et al. 2024).

Preliminaries

Let $S = \{z_1, \dots, z_n\}$ be a set drawn independently from a probability measure ρ defined over a sample space $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, where \mathcal{X} is an input space and $\mathcal{Y} \subset \mathbb{R}$ is an output space. For pairwise learning, our aim is to build a model $h : \mathcal{X} \mapsto \mathbb{R}$ or $h : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$ to simulate the potential mapping lying on ρ . We further assume that the model is parameterized as $h_{\mathbf{w}}$ and the vector \mathbf{w} belongs to a parameter space $\mathcal{W} \subseteq \mathbb{R}^d$. As the essential feature of pairwise learning, the nonnegative loss function takes the form of $f : \mathcal{W} \times \mathcal{Z} \times \mathcal{Z} \mapsto \mathbb{R}$. Since ranking and metric learning are the most popular applications of pairwise learning, we take them as examples here to show how the learning framework involves various learning tasks. Besides, we present details of AUC maximization below, which is used as the experimental validation for our results.

Example 1. (Ranking). Ranking models usually take the form of $h_{\mathbf{w}} : \mathcal{X} \mapsto \mathbb{R}$. Given two instances $z = (x, y), z' =$

(x', y') , we adopt the ordering of $h_{\mathbf{w}}(x), h_{\mathbf{w}}(x')$ as the prediction of the ordering for y, y' . As a result, the prediction $h_{\mathbf{w}}(x) - h_{\mathbf{w}}(x')$ and the true ordering $\text{sgn}(y - y')$ jointly formulate the approach to measure the performance of models. The loss function in this problem is further defined as the pairwise formulation of $f(\mathbf{w}; z, z') = \phi(\text{sgn}(y - y')(h_{\mathbf{w}}(x) - h_{\mathbf{w}}(x')))$. Here we can choose the logistic loss $\phi(t) = \log(1 + \exp(-t))$ or the hinge loss $\phi(t) = \max\{1 - t, 0\}$.

Example 2. (Supervised metric learning). For this problem with output space as $\mathcal{Y} = \{+1, -1\}$, the most usual aim is to learn a Mahalanobis metric $d_{\mathbf{w}}(x, x') = (x - x')^\top \mathbf{w}(x - x')$. Under the parameter $\mathbf{w} \in \mathbb{R}^{d \times d}$ and the corresponding metric, we hope that the distance metric between two instances is consistent with the similarity of labels. Let ϕ be the logistic or the hinge loss defined in Example 1. We can formulate this metric learning problem under the framework of pairwise learning by the loss function as $f(\mathbf{w}; z, z') = \phi(\tau(y, y')d_{\mathbf{w}}(x, x'))$, where $\tau(y, y') = 1$ if $y = y'$ and $\tau(y, y') = -1$ if $y \neq y'$.

Example 3. (AUC Maximization). AUC score is widely applied to measure the performance of classification models for imbalanced data. With the binary output space $\mathcal{Y} = \{+1, -1\}$, it shows the probability that the model $h_{\mathbf{w}} : \mathcal{X} \mapsto \mathbb{R}$ scores a positive instance higher than a negative instance. Therefore, the loss function for AUC maximization usually takes the form of $f(\mathbf{w}; z, z') = g(\mathbf{w}^\top(x - x'))\mathbb{I}_{[y=1, y'=-1]}$, where g can be chosen in the same way as ϕ in Example 1 and \mathbb{I} denotes the indicator function. This demonstrates that AUC maximization also falls into the framework of pairwise learning.

With the pairwise loss function, the population risk is defined as the following

$$F(\mathbf{w}) = \mathbb{E}_{z_i, z_j \sim \rho} [f(\mathbf{w}; z_i, z_j)],$$

which can measure the performance of $h_{\mathbf{w}}$ in real applications. Since ρ is unknown, we consider the empirical risk

$$F_S(\mathbf{w}) = \frac{1}{n(n-1)} \sum_{i, j \in [n]: i \neq j} f(\mathbf{w}; z_i, z_j),$$

where $[n] := \{1, \dots, n\}$. Let $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} F(\mathbf{w})$ and $\mathbf{w}_S = \arg \min_{\mathbf{w} \in \mathcal{W}} F_S(\mathbf{w})$. To approximate the best model $h_{\mathbf{w}^*}$, we apply a randomized algorithm A to the training dataset S and get a corresponding output model. We then use $A(S)$ to denote the parameter of the output model.

Comparing the acquired parameter $A(S)$ and the best parameter \mathbf{w}^* , the excess risk $F(A(S)) - F(\mathbf{w}^*)$ can quantify the performance of $A(S)$ appropriately. We are interested in bounding the excess risk to provide theoretical supports for the practice of learning tasks. To study the risk adequately, we introduce the following decomposition

$$F(A(S)) - F(\mathbf{w}^*) = [F(A(S)) - F(\mathbf{w}^*)] - [F_S(A(S)) - F_S(\mathbf{w}^*)] + [F_S(A(S)) - F_S(\mathbf{w}^*)]. \quad (1)$$

Taking expectation on both sides of the above equation and noting $\mathbb{E}_S [F_S(\mathbf{w}^*)] = F(\mathbf{w}^*)$, we further decompose the

excess risk as

$$\mathbb{E}_{S,A} [F(A(S)) - F(\mathbf{w}^*)] = \mathbb{E}_{S,A} [F(A(S)) - F_S(A(S))] + \mathbb{E}_{S,A} [F_S(A(S)) - F_S(\mathbf{w}^*)]. \quad (2)$$

The first and the second term on the right-hand side are referred to as estimation error (generalization gap) and optimization error respectively. We incorporate SLT and optimization theory to control the two errors, respectively.

In this paper, we consider the learning framework below, which combines RCD and pairwise learning.

Definition 1. (RCD for pairwise learning). Let $\mathbf{w}_1 \in \mathcal{W}$ be the initial point and $\{\eta_t\}$ be a nonnegative stepsize sequence. At the t -th iteration, we first draw i_t from the discrete uniform distribution over $\{1, \dots, d\}$ and then update along the i_t -th coordinate as

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta_t \nabla_{i_t} F_S(\mathbf{w}_t) \mathbf{e}_{i_t}, \quad (3)$$

where $\nabla_{i_t} F_S(\mathbf{w}_t)$ denotes the gradient of the empirical risk w.r.t. to the i_t -th coordinate and \mathbf{e}_{i_t} is a vector with the i_t -th coordinate being 1 and other coordinates being 0.

Considering the generalization for the above paradigm, we leverage the concept of algorithmic stability to handle the estimation error. Algorithmic stability shows how algorithms react to perturbations of training datasets. Various stability measures have been proposed to study the generalization gap in SLT, including uniform stability, argument stability and on-average stability. Here we introduce the uniform stability and the on-average argument stability, with the latter being particularly useful for generalization analysis in this paper. It is notable that we follow Lei and Ying (2020) in the definition of ℓ_1 and ℓ_2 on-average argument stabilities. The ℓ_1 on-average argument stability refers to the ℓ_1 -norm of the vector $(\|A(S) - A(S_1)\|_2, \dots, \|A(S) - A(S_n)\|_2)$, while the ℓ_2 on-average argument stability refers to the ℓ_2 -norm of this vector.

Definition 2. (Algorithmic Stability). Drawing independently from ρ , we get the following two datasets

$$S = \{z_1, \dots, z_n\} \quad \text{and} \quad S' = \{z'_1, \dots, z'_n\}.$$

We then replace z_i in S with z'_i for any $i \in [n]$ and have

$$S_i = \{z_1, \dots, z_{i-1}, z'_i, z_{i+1}, \dots, z_n\}.$$

Let $\mathbf{x} \in \mathbb{R}^d$ be a vector of dimension d . Then we denote the p -norm $\|\mathbf{x}\|_p = (\sum_{i=1}^d |\mathbf{x}_i|^p)^{1/p}$ and show several stability measures below.

(a) Randomized algorithm A is ϵ -uniformly stable if for any $S, S_i \in \mathcal{Z}^n$ the following inequality holds

$$\sup_{z, \tilde{z}} [f(A(S), z, \tilde{z}) - f(A(S_i), z, \tilde{z})] \leq \epsilon.$$

(b) We say A is ℓ_1 on-average argument ϵ -stable if

$$\mathbb{E}_{S, S', A} \left[\frac{1}{n} \sum_{i=1}^n \|A(S) - A(S_i)\|_2 \right] \leq \epsilon.$$

(c) We say A is ℓ_2 on-average argument ϵ -stable if

$$\mathbb{E}_{S, S', A} \left[\frac{1}{n} \sum_{i=1}^n \|A(S) - A(S_i)\|_2^2 \right] \leq \epsilon^2.$$

As indicated below, We prepare several necessary assumptions so that relative generalization bounds can be derived effectively. Assumption 1 and Assumption 2 are useful for bounding the on-average argument stability. Assumption 3 is mainly applied in the proof of the optimization error. The other two assumptions show the convexity of the empirical risk, which is the basic condition for the establishment of relative theorems.

Assumption 1. For all $(z, z') \in \mathcal{Z} \times \mathcal{Z}$ and $\mathbf{w} \in \mathcal{W}$, the loss function satisfies the G -Lipschitz continuity condition as $\|\nabla f(\mathbf{w}, z, z')\|_2 \leq G$.

Assumption 2. For all $(z, z') \in \mathcal{Z} \times \mathcal{Z}$ and $\mathbf{w}, \mathbf{w}' \in \mathcal{W}$, the loss function is L -smooth as $\|\nabla f(\mathbf{w}; z, z') - \nabla f(\mathbf{w}'; z, z')\|_2 \leq L\|\mathbf{w} - \mathbf{w}'\|_2$.

Assumption 3. For any S , F_S has coordinate-wise Lipschitz continuous gradients with parameter $\tilde{L} > 0$, i.e., we have the following inequality for all $\alpha \in \mathbb{R}$, $\mathbf{w} \in \mathcal{W}$, $i \in [d]$

$$F_S(\mathbf{w} + \alpha \mathbf{e}_i) \leq F_S(\mathbf{w}) + \alpha \nabla_i F_S(\mathbf{w}) + \tilde{L} \alpha^2 / 2.$$

Assumption 4. F_S is convex for any S , i.e., $F_S(\mathbf{w}) - F_S(\mathbf{w}') \geq \langle \mathbf{w} - \mathbf{w}', \nabla F_S(\mathbf{w}') \rangle$ holds for all $\mathbf{w}, \mathbf{w}' \in \mathcal{W}$.

Assumption 5. F_S is σ -strongly convex for any S , i.e., the following inequality holds for all $\mathbf{w}, \mathbf{w}' \in \mathcal{W}$

$$F_S(\mathbf{w}) - F_S(\mathbf{w}') \geq \langle \mathbf{w} - \mathbf{w}', \nabla F_S(\mathbf{w}') \rangle + \sigma \|\mathbf{w} - \mathbf{w}'\|_2^2 / 2,$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product of two vectors.

With Definition 1 and Definition 2, we can further quantify stabilities of RCD for pairwise learning. Then we show connections between the estimation error and stability measures by the following lemma, which is the key to apply algorithmic stability effectively in generalization analysis. While part (a) of Lemma 1 is motivated by the case of point-wise learning (Hardt, Recht, and Singer 2016) and derived with the technique similar to Lei, Liu, and Ying (2021), part (b) and part (c) are introduced from Lei, Liu, and Ying (2021) and Lei, Ledent, and Kloft (2020) respectively. In part (c), the base of the natural logarithm takes the symbol as e and $\lceil \alpha \rceil$ means rounding up for α .

Lemma 1. Let S, S_i be constructed as Definition 2. Then we bound estimation errors with stability measures below.

(a) Let Assumption 1 hold. Then the estimation error can be bounded by the ℓ_1 on-average argument stability below

$$\begin{aligned} \mathbb{E}_{S,A} [F(A(S)) - F_S(A(S))] \\ \leq \frac{2G}{n} \sum_{i=1}^n \mathbb{E}_{S, S', A} [\|A(S_i) - A(S)\|_2]. \end{aligned} \quad (4)$$

(b) Let Assumption 2 hold. Then for any $\gamma > 0$ we have the following estimation error bound with the ℓ_2 on-average argument stability

$$\begin{aligned} \mathbb{E}_{S,A} [F(A(S)) - F_S(A(S))] &\leq \frac{L}{\gamma} \mathbb{E}_{S,A} [F_S(A(S))] \\ &+ \frac{2(L + \gamma)}{n} \sum_{i=1}^n \mathbb{E}_{S, S', A} [\|A(S_i) - A(S)\|_2^2]. \end{aligned} \quad (5)$$

(c) Let n denote the sample size of S . Assume for any S and $(z, z') \in \mathcal{Z} \times \mathcal{Z}$, $|f(A(S); z, z')| \leq R$ holds for $R > 0$. Suppose that A is ϵ -uniformly-stable and $\delta \in (0, 1/e)$, then the following inequality holds with probability at least $1 - \delta$

$$\begin{aligned} |F(A(S)) - F_S(A(S))| &\leq 4\epsilon + e \left(12\sqrt{2}R \sqrt{\frac{\log(e/\delta)}{n-1}} \right. \\ &\quad \left. + 48\sqrt{6}\epsilon \lceil \log_2(n-1) \rceil \log(e/\delta) \right). \end{aligned} \quad (6)$$

Remark 1. While estimation error bound (4) is established under the Lipschitz continuity condition, (5) remove this condition based on the ℓ_2 on-average argument stability measure. Inequality (5) holds with the L -smoothness of the loss function, which replaces the Lipschitz constant in (4) by the empirical risk. Furthermore, if A is ℓ_2 on-average argument ϵ -stable, we can take $\gamma = \sqrt{L\mathbb{E}_{S,A}[F_S(A(S))]} / (\sqrt{2}\epsilon)$ in part (b) and get $\mathbb{E}_{S,A}[F(A(S)) - F_S(A(S))] \leq \sqrt{2L\mathbb{E}_{S,A}[F_S(A(S))]} \epsilon + 2L\epsilon^2$. If the empirical risk $\mathbb{E}_{S,A}[F_S(A(S))] = O(1/n)$, then we further know $\mathbb{E}_{S,A}[F(A(S)) - F_S(A(S))] = O(\epsilon^2 + \epsilon/\sqrt{n})$, which means the estimation error bound is well dependent on the stability measure ϵ via the small risk of the output model (Hardt, Recht, and Singer 2016). Other than the generalization error in expectation, the link in high probability (6) presents the convergence rate of $O(n^{-\frac{1}{2}} + \epsilon\sqrt{\log_2(n)})$ for ϵ -uniformly stable algorithm. This result is achieved by combining a concentration inequality from Bousquet, Klochkov, and Zhivotovskiy (2020) and the decoupling technique in Lei and Ying (2020).

Besides the estimation error, we need to tackle the optimization error to achieve complete excess risk bounds. The optimization error analysis for pointwise learning can be directly extended to pairwise learning since they both use unbiased gradient estimations. Since pointwise learning and pairwise learning mainly differ in terms of loss structure, Lemma 2 from pointwise learning also works for pairwise learning.

Lemma 2. Let $\{\mathbf{w}_t\}$ be produced by RCD (3) with nonincreasing step sizes $\eta_t \leq 1/\tilde{L}$. Let Assumptions 3,4 hold, then the following two inequalities holds for any $\mathbf{w} \in \mathcal{W}$

$$\mathbb{E}_A[F_S(\mathbf{w}_t) - F_S(\mathbf{w})] \leq \frac{d(\|\mathbf{w}_1 - \mathbf{w}\|_2^2 + 2\eta_1 F_S(\mathbf{w}_1))}{2\sum_{j=1}^t \eta_j} \quad (7)$$

and

$$\begin{aligned} 2\sum_{j=1}^t \eta_j^2 \mathbb{E}_A[F_S(\mathbf{w}_j) - F_S(\mathbf{w})] \\ \leq d\eta_1 \|\mathbf{w}_1 - \mathbf{w}\|_2^2 + 2d\eta_1^2 F_S(\mathbf{w}_1). \end{aligned} \quad (8)$$

Let Assumption 5 hold and $\mathbf{w}_S = \arg \min_{\mathbf{w} \in \mathcal{W}} F_S(\mathbf{w})$, then we have the following inequality

$$\begin{aligned} \mathbb{E}_A[F_S(\mathbf{w}_{t+1}) - F_S(\mathbf{w}_S)] \\ \leq (1 - \eta_t \sigma/d) \mathbb{E}_A[F_S(\mathbf{w}_t) - F_S(\mathbf{w}_S)]. \end{aligned} \quad (9)$$

In the arXiv version, Appendix B restates the above two lemmas and prepares some other lemmas. The proof for part (a) of Lemma 1 is given in Appendix B.1. Considering the stability analysis, we introduce the coercivity property of the gradient descent operator in Appendix B.3 (Hardt, Recht, and Singer 2016). Then we show the self-bounding property of L -smooth functions in Appendix B.4 (Srebro, Sridharan, and Tewari 2010), which plays a key role in introducing empirical risks into the ℓ_2 on-average argument stability.

Main Results

In this section, we show our results on generalization analysis of RCD for pairwise learning. For both convex and strongly convex cases, we derive the on-average argument stability bounds and as well as the corresponding excess risk bounds. Results are organized according to the convexity of the empirical risk.

Generalization for Convex Case

This subsection describes the ℓ_2 on-average argument stabilities for the convex empirical risk. Based on stability analysis, we consider generalization bounds in expectation under the setting that applies RCD for pairwise learning.

If the empirical risk is convex and L -smooth, then the gradient descent operator enjoys the coercivity property according to Hardt, Recht, and Singer (2016). Since taking expectations for the coordinate descent operator yields the gradient descent operator, the coercivity property is useful to bound the expansiveness of RCD updates in the stability analysis. With the coercivity property of the coordinate descent operator in expectation, we further incorporate the self-bounding property of L -smooth functions to measure the ℓ_2 on-average argument stability. Then we handle the estimation error by plugging the stability measure into part (b) of Lemma 1. We finally introduce the optimization error and derive the corresponding excess risk bound. The proof is given in Appendix C of the arXiv version.

Theorem 3. Let Assumptions 2, 3, 4 hold. Let $\{\mathbf{w}_t\}, \{\mathbf{w}_t^{(i)}\}$ be produced by (3) with $\eta_t \leq 1/L$ based on S and S_i respectively. Then the ℓ_2 on-average argument stability satisfies

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{S,S',A} [\|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^{(i)}\|_2^2] \\ \leq \frac{128L}{n^2 d} \left(\frac{t}{d} + 1 \right) \sum_{j=1}^t \eta_j^2 \mathbb{E}_{S,A} [F_S(\mathbf{w}_j)]. \end{aligned} \quad (10)$$

Assume that the nonincreasing step size sequence $\{\eta_t\}$ satisfies $\eta_t \leq 1/\tilde{L}$. Then, for any $\gamma \geq 0$, we have

$$\begin{aligned} \mathbb{E}_{S,A} [F(\mathbf{w}_T) - F(\mathbf{w}^*)] \\ = O \left(\frac{d(1 + L\gamma^{-1})}{\sum_{t=1}^T \eta_t} + \frac{L(L + \gamma)(T + d)}{n^2 d} \right) \\ + O \left(\frac{L}{\gamma} + \frac{L(L + \gamma)(T + d)}{n^2 d^2} \sum_{t=1}^T \eta_t^2 \right) \times F(\mathbf{w}^*). \end{aligned} \quad (11)$$

Furthermore, for a constant step size as $\eta_t \equiv \eta$, we choose $T \asymp n^{\frac{1}{2}} d L^{-\frac{1}{2}}$ and get

$$\mathbb{E}_{S,A} [F(\mathbf{w}_T) - F(\mathbf{w}^*)] = O\left(\sqrt{\frac{L}{n}}\right). \quad (12)$$

Assuming that $F(\mathbf{w}^*) = O(Ln^{-1})$, we choose $T \asymp ndL^{-1}$ to give

$$\mathbb{E}_{S,A} [F(\mathbf{w}_T) - F(\mathbf{w}^*)] = O\left(\frac{L}{n}\right). \quad (13)$$

Remark 2. For pairwise learning, Eq. (10) shows that RCD enjoys the ℓ_2 on-average argument stability is of the order of $O(L(t+d) \sum_{j=1}^t \eta_j^2 \mathbb{E}_{S,A} [F_S(\mathbf{w}_j)] / (n^2 d^2))$. This bound means that the output model of RCD becomes more and more stable with the sample size increasing or the number of iterations decreasing. In further detail, since the estimation error can be bounded by the stability bound according to (5), decreasing the number of iterations is beneficial to controlling the estimation error. However, increasing the number of iterations corresponds to the optimization process, which is the key to control the optimization error. As a result, the early stopping strategy is adopted to balance the estimation and optimization for a good generalization.

Remark 3. Fixing $F(\mathbf{w}^*)$, we choose an appropriate number of iterations for the excess risk bound (12). Besides, we incorporate $F(\mathbf{w}^*)$ into the excess risk bound and get the convergence rate (13). It is obvious that (13) exploits the low noise setting to yield the optimistic bound (Srebro, Sridharan, and Tewari 2010). Furthermore, since RCD updates in expectation is closely related to the gradient descent operator, we consider the batch setting (Nikolakakis, Karbasi, and Kalogerias 2023) and find that results here are identical to those of full-batch GD (Nikolakakis et al. 2023). Turning to SGD for pairwise learning (Lei, Liu, and Ying 2021), the ℓ_2 on-average argument stability takes a slower rate as $O(1/n)$ under the same setting. The excess risk bound can achieve the rate of $(1/\sqrt{n})$ in a general setting with $\eta_t = \eta \asymp 1/\sqrt{T}$ and $T \asymp n$. With the low noise setting $F(\mathbf{w}^*) = O(n^{-1})$, the optimistic bound $O(1/n)$ is also derived.

Generalization for Strongly Convex Case

This subsection presents generalization analysis of RCD for pairwise learning in a strongly convex setting. In the strongly convex case of pointwise learning (Wang, Wu, and Lei 2021), the ℓ_2 on-average argument stability and the corresponding generalization bound were not taken into consideration. Therefore, we not only measure the stability here but also derive the excess risk bound for the strongly convex empirical risk. We show the proof in Appendix D of the arXiv version.

Theorem 4. *Let Assumptions 2, 3, 5 hold. Let $\{\mathbf{w}_t\}$, $\{\mathbf{w}_t^{(i)}\}$ be produced by (3) with $\eta_t \leq \beta/L$ for any $\beta \in (0, 1)$ based on S and S_i , respectively. Then the ℓ_2 on-average argument*

stability is

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{S,S',A} \left[\|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^{(i)}\|_2^2 \right] \\ & \leq \frac{128L}{n^2 d} \sum_{j=1}^t \left(\frac{t}{d} \prod_{k=j+1}^t \left(1 - \frac{2\eta_k(1-\beta)(n-2)\sigma}{nd} \right)^2 \right. \\ & \quad \left. + \prod_{k=j+1}^t \left(1 - \frac{2\eta_k(1-\beta)(n-2)\sigma}{nd} \right) \right) \eta_j^2 \mathbb{E}_{S,A} [F_S(\mathbf{w}_j)]. \end{aligned} \quad (14)$$

Let step sizes be fixed as $\eta_t \equiv \eta \leq 1/\tilde{L}$. For any $\gamma \geq 0$, we develop the excess risk bound as

$$\begin{aligned} & \mathbb{E}_{S,A} [F(\mathbf{w}_{T+1}) - F_S(\mathbf{w}_S)] \\ & = O\left(\left(1 + \frac{L}{\gamma} \right) (1 - \eta\sigma/d)^T + \frac{L(d+T)(L+\gamma)}{(n-2)^2 \sigma^2 (1-\beta)^2} \right) \\ & \quad + O\left(\frac{L}{\gamma} + \frac{L(d+T)(L+\gamma)}{(n-2)^2 \sigma^2 (1-\beta)^2} \right) \times \mathbb{E}[F_S(\mathbf{w}_S)]. \end{aligned} \quad (15)$$

Choosing $T \asymp d\sigma^{-1} \log(n\sigma L^{-1})$ yields

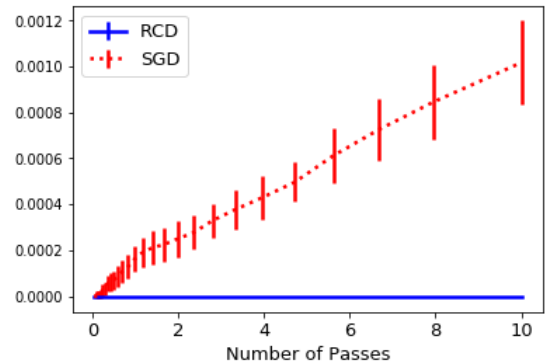
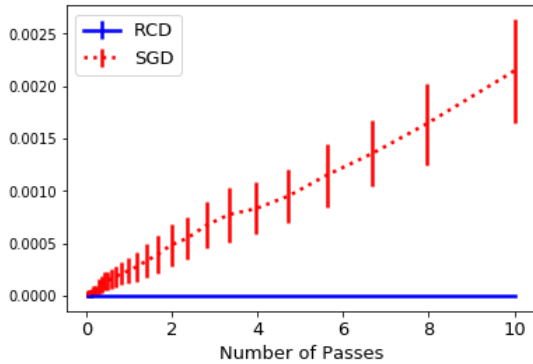


$$\mathbb{E}_{S,A} [F(\mathbf{w}_{T+1}) - F_S(\mathbf{w}_S)] = O\left(\frac{Ld^{\frac{1}{2}}}{n\sigma^{\frac{3}{2}}} \sqrt{\log\left(\frac{n\sigma}{L}\right)} \right). \quad (16)$$

Remark 4. As shown in (14), the stability measure involves a weighted sum of empirical risks. This demonstrates that low risks of output models can improve the stability along the training process. The measure also shows information including the convexity parameter σ and learning rates η_j which are closely associated with the interplay between RCD and training datasets. Furthermore, the strong convexity of the empirical risk obviously leads to a better stability as compared to the convex case (10).

Remark 5. The convergence rate (16) gives the choice of T to balance the estimation and optimization. Indeed, the optimal convergence rate lies between $O(1/(n\sigma))$ and $O(\sqrt{\log(n\sigma)}/(n\sigma^{\frac{3}{2}}))$, for which the corresponding choices of T are smaller than that we give. It is notable that the approximate optimal rate here almost matches the optimistic bound (13). Besides, the strong convexity promotes the fast computing $T \asymp \log(n)$ as compared to $T \asymp n$ under the convex case. Results here are the same as those for full-batch GD (Nikolakakis et al. 2023), which can verify the theorem since the expectation for RCD leads to the gradient descent operator. Considering SGD, Lei, Liu, and Ying (2021) present the generalization bounds for pairwise learning. Under the same setting of smoothness and strong convexity, SGD achieves the excess risk bound $O(1/(n\sigma))$. However, the convergence rate of SGD requires the number of iterations as $T \asymp O(n/\sigma)$ and a small $F(\mathbf{w}^*)$.

Experimental Verification

Here we choose the example of AUC maximization to verify the theoretical results on stability measures. Figure 1 shows



(a) stability results for hinge loss

(b) stability results for logistic loss

Figure 1: Euclidean distance Δ_t as a function of the number of passes for the hinge loss.

results based on the usps dataset from LIBSVM (Chang and Lin 2011). More sufficient experiments about other datasets are presented in Appendix E of the arXiv version.

We follow the settings of SGD for pairwise learning (Lei, Liu, and Ying 2021) and compare the results of RCD and SGD. In each experiment, we randomly choose 80 percents of the dataset as the training set S . Then we perturb a signal example of S to construct the neighboring dataset S' . We apply RCD or SGD to S, S' and get two iterate sequences, with which we plot the Euclidean distance $\Delta_t = \|\mathbf{w}_t - \mathbf{w}'_t\|_2$ in each iteration. While the learning rates are set as $\eta_t = \eta/\sqrt{T}$ with $\eta \in \{0.05, 0.25, 1, 4\}$ for RCD, we only compare RCD and SGD under the setting of $\eta = 0.05$. Letting n be the sample size, we report Δ_t as a function of T/n (the number of passes). We repeat the experiments 100 times, and consider the average and the standard deviation.

Considering the comparison between SGD and RCD, while the term $(T/n)^2$ dominates the convergence rates of stability bounds for SGD according to Lei, Liu, and Ying (2021), the on-average argument stability bound (10) for RCD takes the order of $O(T/n^2)$. The results for the comparison are consistent with the theoretical stability bounds. Furthermore, the Euclidean distance under the logistic loss is significantly smaller than that under the hinge loss, which is consistent with the discussions of Lei, Liu, and Ying (2021).

Conclusion

In this paper, we study the generalization performance of RCD for pairwise learning. We measure the on-average argument stability develop the corresponding excess risk bound. Results for the convex empirical risk show us how the early-stopping strategy can balance estimation and optimization. The excess risk bounds enjoy the convergence rates of $O(1/\sqrt{n})$ and $O(\sqrt{\log(n)}/n)$ under the convex and strongly convex cases, respectively. Furthermore, we incorporate the low noise setting $F(\mathbf{w}^*) = O(1/n)$ to explore better generalizations under the convex case.

There remain several questions for further investigation. Explorations under the nonparametric or the non-convex case are important for extending the applications of RCD. RCD for the specific large-scale matrix optimization also deserves a fine-grained generalization analysis.

Acknowledgments

The work of L. Wu is supported by the National Natural Science Foundation of China (72431008, 61903309) and the Sichuan Science and Technology Program (2023NSF-SC1355). The work of Y. Lei is supported by the Research Grants Council of Hong Kong [Project No. 22303723]. We are also grateful to the anonymous AAAI reviewers for their insightful and constructive comments.

References

- Agarwal, S.; and Niyogi, P. 2009. Generalization bounds for ranking algorithms via algorithmic stability. *Mathematics of Operations Research*, 10(2): 441–474.
- Bartlett, P. L.; and Mendelson, S. 2001. Rademacher and Gaussian complexities: Risk bounds and structural results. In *International Conference on Computational Learning Theory*, 224–240.
- Beck, A.; and Teboulle, M. 2021. Dual randomized coordinate descent method for solving a class of nonconvex problems. *SIAM Journal on Optimization*, 31(3): 1877–1896.
- Bellet, A.; and Habrard, A. 2015. Robustness and generalization for metric learning. *Journal of Machine Learning Research*, 151: 259–267.
- Bousquet, O.; and Elisseeff, A. 2002. Stability and generalization. *Journal of Machine Learning Research*, 2(3): 499–526.
- Bousquet, O.; Klochkov, Y.; and Zhivotovskiy, N. 2020. Sharper bounds for uniformly stable algorithms. In *Conference on Learning Theory*, 610–626.
- Callahan, M.; Vu, T.; and Raich, R. 2024. Provable randomized coordinate descent for matrix completion. In *International Conference on Acoustics, Speech and Signal Processing*, 9421–9425.
- Cao, Q.; Guo, Z.-C.; and Ying, Y. 2016. Generalization bounds for metric and similarity learning. *Machine Learning*, 102(1): 115–132.
- Chang, C.-C.; and Lin, C.-J. 2011. LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3): 1–27.
- Chang, K.-W.; Hsieh, C.-J.; and Lin, C.-J. 2008. Coordinate descent method for large-scale ℓ_2 -loss linear support vector machines. *Journal of Machine Learning Research*, 9(7): 1369–1398.
- Chen, J.; Chen, H.; Jiang, X.; Gu, B.; Li, W.; Gong, T.; and Zheng, F. 2023. On the stability and generalization of triplet learning. In *AAAI Conference on Artificial Intelligence*, volume 37, 7033–7041.
- Chen, Z.; Li, Y.; and Lu, J. 2023. On the global convergence of randomized coordinate gradient descent for non-convex optimization. *SIAM Journal on Optimization*, 33(2): 713–738.
- Christmann, A.; and Zhou, D.-X. 2016. On the robustness of regularized pairwise learning methods based on kernels. *Journal of Complexity*, 37: 1–33.
- Cléménçon, S.; Lugosi, G.; and Vayatis, N. 2008. Ranking and empirical minimization of U-statistics. *The Annals of Statistics*, 36(2): 844–974.
- Cortes, C.; and Mohri, M. 2003. AUC optimization vs. error rate minimization. In *Advances in Neural Information Processing Systems*, 313–320.
- Damaskinos, G.; Mendler-Dünner, C.; Guerraoui, R.; Papandreou, N.; and Parnell, T. 2021. Differentially private stochastic coordinate descent. In *AAAI Conference on Artificial Intelligence*, 7176–7184.
- Davis, D.; and Drusvyatskiy, D. 2022. Graphical convergence of subgradients in nonconvex optimization and learning. *Mathematics of Operations Research*, 47(1): 209–231.
- Deng, X.; Sun, T.; Li, S.; and Li, D. 2023. Stability-based generalization analysis of the asynchronous decentralized S-GD. In *AAAI Conference on Artificial Intelligence*, 7340–7348.
- Deng, Z.; He, H.; and Su, W. 2021. Toward better generalization bounds with locally elastic stability. In *International Conference on Machine Learning*, 2590–2600.
- Deora, P.; Ghaderi, R.; Taheri, H.; and Thrampoulidis, C. 2024. On the Optimization and Generalization of Multi-head Attention. *Transactions on Machine Learning Research*.
- Dong, M.; Yang, X.; Zhu, R.; Wang, Y.; and Xue, J.-H. 2020. Generalization bound of gradient descent for non-convex metric learning. In *Advances in Neural Information Processing Systems*, 9794–9805.
- Dong, Y.; Gong, T.; Chen, H.; He, Z.; Li, M.; Song, S.; and Li, C. 2024. Towards Generalization beyond Pointwise Learning: A Unified Information-theoretic Perspective. In *International Conference on Machine Learning*.
- Elisseeff, A.; Evgeniou, T.; Pontil, M.; and Kaelbling, L. P. 2005. Stability of randomized learning algorithms. *Journal of Machine Learning Research*, 6(1): 55–79.
- Fan, J.; Hu, T.; Wu, Q.; and Zhou, D.-X. 2016. Consistency analysis of an empirical minimum error entropy algorithm. *Applied and Computational Harmonic Analysis*, 41(1): 164–189.
- Feldman, V. 2016. Generalization of erm in stochastic convex optimization: The dimension strikes back. In *Advances in Neural Information Processing Systems*, 3576–3584.
- Feldman, V.; and Vondrak, J. 2019. High probability generalization bounds for uniformly stable algorithms with nearly optimal rate. In *Conference on Learning Theory*, 1270–1279.
- Gao, W.; Jin, R.; Zhu, S.; and Zhou, Z.-H. 2013. One-pass AUC optimization. In *International Conference on Machine Learning*, 906–914.
- Guo, Z.-C.; Xiang, D.-H.; Guo, X.; and Zhou, D.-X. 2017. Thresholded spectral algorithms for sparse approximations. *Analysis and Applications*, 15(3): 433–455.
- Hardt, M.; Recht, B.; and Singer, Y. 2016. Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine Learning*, 1225–1234.
- Hu, E.-L.; and Kwok, J. T. 2019. Low-rank matrix learning using biconvex surrogate minimization. *IEEE Transactions on Neural Networks and Learning Systems*, 30(11): 3517–3527.
- Hu, T.; Fan, J.; Wu, Q.; and Zhou, D.-X. 2015. Regularization schemes for minimum error entropy principle. *Analysis and Applications*, 13(4): 437–455.
- Huang, S.; Zhou, J.; Feng, H.; and Zhou, D.-X. 2023. Generalization analysis of pairwise learning for ranking with deep neural networks. *Neural Computation*, 35(6): 1135–1158.

- Kar, P.; Sriperumbudur, B.; Jain, P.; and Karnick, H. 2013. On the generalization ability of online learning algorithms for pairwise loss functions. In *International Conference on Machine Learning*, 441–449.
- Klochkov, Y.; and Zhivotovskiy, N. 2021. Stability and deviation optimal risk bounds with convergence rate $O(1/n)$. In *Advances in Neural Information Processing Systems*, 5065–5076.
- Lei, Y.; Ledent, A.; and Kloft, M. 2020. Sharper generalization bounds for pairwise learning. In *Advances in Neural Information Processing Systems*, 21236–21246.
- Lei, Y.; Liu, M.; and Ying, Y. 2021. Generalization guarantee of SGD for pairwise learning. In *Advances in Neural Information Processing Systems*, 21216–21228.
- Lei, Y.; Sun, T.; and Liu, M. 2023. Stability and Generalization for Minibatch SGD and Local SGD. *arXiv preprint arXiv:2310.01139*.
- Lei, Y.; and Ying, Y. 2020. Fine-grained analysis of stability and generalization for stochastic gradient descent. In *International Conference on Machine Learning*, 5809–5819.
- Li, J.; Luo, X.; and Qiao, M. 2020. On generalization error bounds of noisy gradient methods for non-convex learning. In *International Conference on Learning Representations*.
- Liu, M.; Zhang, X.; Chen, Z.; Wang, X.; and Yang, T. 2018. Fast stochastic AUC maximization with $o(1/n)$ -convergence rate. In *International Conference on Machine Learning*, 3189–3197.
- Liu, T.; Lugosi, G.; Neu, G.; and Tao, D. 2017. Algorithmic stability and hypothesis complexity. In *International Conference on Machine Learning*, 2159–2167.
- Lu, Z.; and Xiao, L. 2015. On the complexity analysis of randomized block-coordinate descent methods. *Mathematical Programming*, 152(1): 615–642.
- Lu, Z.; and Xiao, L. 2017. Efficiency of the accelerated coordinate descent method on structured optimization problems. *SIAM Journal on Optimization*, 27(1): 110–123.
- Mei, S.; Bai, Y.; and Montanari, A. 2018. The landscape of empirical risk for nonconvex losses. *The Annals of Statistics*, 46: 2747–2774.
- Nesterov, Y. 2012. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2): 341–362.
- Nikolakakis, K. E.; Haddadpour, F.; Karbasi, A.; and Kalogerias, D. S. 2023. Beyond lipschitz: Sharp generalization and excess risk bounds for full-batch gd. In *International Conference on Learning Representations*.
- Nikolakakis, K. E.; Karbasi, A.; and Kalogerias, D. 2023. Select without fear: Almost all mini-batch schedules generalize optimally. *arXiv:2305.02247*.
- Rakhlin, A.; Mukherjee, S.; and Poggio, T. 2005. Stability results in learning theory. *Analysis and Applications*, 3(4): 397–417.
- Rejchel, W. 2012. On ranking and generalization bounds. *Journal of Machine Learning Research*, 13(1): 1373–1392.
- Richtárik, P.; and Takáč, M. 2014. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Mathematical Programming*, 144(1): 1–38.
- Rodomanov, A.; and Kropotov, D. 2020. A randomized coordinate descent method with volume sampling. *SIAM Journal on Optimization*, 30(3): 1878–1904.
- Schliserman, M.; Sherman, U.; and Koren, T. 2024. The Dimension Strikes Back with Gradients: Generalization of Gradient Methods in Stochastic Convex Optimization. *arXiv preprint arXiv:2401.12058*.
- Shalev-Shwartz, S.; Shamir, O.; Srebro, N.; and Sridharan, K. 2010. Learnability, stability and uniform convergence. *Journal of Machine Learning Research*, 11(10): 2635–2670.
- Shalev-Shwartz, S.; and Tewari, A. 2009. Stochastic methods for l_1 regularized loss minimization. In *International Conference on Machine Learning*, 929–936.
- Srebro, N.; Sridharan, K.; and Tewari, A. 2010. Smoothness, low noise and fast rates. In *Advances in Neural Information Processing Systems*, 2199–2207.
- Vapnik, V.; Levin, E.; and Le Cun, Y. 1994. Measuring the VC-dimension of a learning machine. *Neural Computation*, 6(5): 851–876.
- Wang, J.; Chen, J.; Chen, H.; Gu, B.; Li, W.; and Tang, X. 2023. Stability-based generalization analysis for mixtures of pointwise and pairwise learning. In *AAAI Conference on Artificial Intelligence*, 10113–10121.
- Wang, P.; Wu, L.; and Lei, Y. 2021. Stability and generalization for randomized coordinate descent. In *International Joint Conferences on Artificial Intelligence*, 3104–3110.
- Wang, Y.; Khardon, R.; Pechyony, D.; and Jones, R. 2012. Generalization bounds for online learning algorithms with pairwise loss functions. In *Conference on Learning Theory*, 1–22.
- Xie, Y.; Wang, Z.; and Zhang, Z. 2024. Randomized methods for computing optimal transport without regularization and their convergence analysis. *Journal of Scientific Computing*, 100(11): 37–65.
- Ye, H.-J.; Zhan, D.-C.; and Jiang, Y. 2019. Fast generalization rates for distance metric learning. *Machine Learning*, 108(2): 267–295.
- Ying, Y.; Wen, L.; and Lyu, S. 2016. Stochastic online AUC maximization. In *Advances in Neural Information Processing Systems*, 451–459.
- Ying, Y.; and Zhou, D.-X. 2016. Online pairwise learning algorithms. *Neural Computation*, 28(4): 743–777.
- Zhao, T.; Yu, M.; Wang, Y.; Arora, R.; and Liu, H. 2014. Accelerated mini-batch randomized block coordinate descent method. In *Advances in Neural Information Processing Systems*, 3329–3337.
- Zhou, D.-X. 2002. The covering number in learning theory. *Journal of Complexity*, 18(3): 739–767.
- Zhou, J.; Wang, P.; and Zhou, D.-X. 2024. Generalization analysis with deep ReLU networks for metric and similarity learning. *arXiv:2405.06415*.