

Decentralized Riemannian Algorithm for Nonconvex Minimax Problems

Xidong Wu, Zhengmian Hu, Heng Huang

Electrical and Computer Engineering, University of Pittsburgh, Pittsburgh, PA, United States
 xidong_wu@outlook.com, huzhengmian@gmail.com, henghuanghh@gmail.com

Abstract

The minimax optimization over Riemannian manifolds (possibly nonconvex constraints) has been actively applied to solve many problems, such as robust dimensionality reduction and deep neural networks with orthogonal weights (Stiefel manifold). Although many optimization algorithms for minimax problems have been developed in the Euclidean setting, it is difficult to convert them into Riemannian cases, and algorithms for nonconvex minimax problems with nonconvex constraints are even rare. On the other hand, to address the big data challenges, decentralized (serverless) training techniques have recently been emerging since they can reduce communications overhead and avoid the bottleneck problem on the server node. Nonetheless, the algorithm for decentralized Riemannian minimax problems has not been studied. In this paper, we study the distributed nonconvex-strongly-concave minimax optimization problem over the Stiefel manifold and propose both deterministic and stochastic minimax methods. The Stiefel manifold is a non-convex set. The global function is represented as the finite sum of local functions. For the deterministic setting, we propose DRGDA and prove that our deterministic method achieves a gradient complexity of $O(\epsilon^{-2})$ under mild conditions. For the stochastic setting, we propose DRSGDA and prove that our stochastic method achieves a gradient complexity of $O(\epsilon^{-4})$. The DRGDA and DRSGDA are the first algorithms for distributed minimax optimization with nonconvex constraints with exact convergence. Extensive experimental results on the Deep Neural Networks (DNNs) training over the Stiefel manifold demonstrate the efficiency of our algorithms.

Introduction

The minimax problems arise throughout machine learning, such as generative adversarial networks (GANs) (Goodfellow et al. 2014), adversarial training (Wu, Zhang, and Huang 2022), and policy evaluation (Wai et al. 2019). In recent years, minimax problems over Riemannian manifolds (possibly nonconvex constraints) have attracted intensive interest because of their important applications in many scientific and engineering fields (Jordan, Lin, and Vlatakis-Gkaragkounis 2022), such as subspace robust Wasserstein distance (Paty and Cuturi 2019), distributionally robust principal component analysis (PCA) (Zhang, Zhang, and

Sra 2022), adversarial training over Riemannian manifolds (Huang and Gao 2020) and orthonormal generative adversarial networks (GANs) (Müller, Klein, and Weinmann 2019). Among them, optimization over the Stiefel manifold is popular because, from robust PCA to deep neural networks with orthogonal weights, a plethora of machine learning applications can be cast into it (Huang et al. 2018).

As sizes of model parameters and training datasets keep increasing, machine learning tasks heavily rely on distributed training (Bao et al. 2022; Lu et al. 2019; Wu et al. 2022). Given that computing capacity increases rapidly, communication overhead gradually becomes the bottleneck of training speed. As a result, research on communication efficient distributed optimization is very important. One important direction is decentralized distributed training, where each worker node updates the local model based on its local data and communicates exclusively with its neighbors, instead of a central server (Lian et al. 2017). In other words, when there are a lot of worker nodes, the decentralized distributed training strategy is more conducive to communication. In addition, due to data heterogeneity in the distributed setting, distributed robust training has wide applications. The distributed robustification of many ML tasks can also be formulated as decentralized minimax optimization problems (Jordan, Lin, and Vlatakis-Gkaragkounis 2022).

Although decentralized distributed training has gained popularity, algorithms for decentralized Riemannian minimax problems have not been studied. The existing minimax optimization methods over Riemannian manifolds cannot be applied to these problems (Jordan, Lin, and Vlatakis-Gkaragkounis 2022; Zhang, Zhang, and Sra 2022; Han et al. 2022; Huang and Gao 2020), since (1) they mainly focus on the geodesically strongly-convex-concave case or require PL condition, which might not be satisfied in the deep learning and nonconvex optimization (Han et al. 2022); (2) algorithms cannot be used in the distributed setting.

In this paper, we study the following distributed nonconvex-strongly-concave minimax problem on Stiefel manifolds:

$$\min_{x \in \text{St}(d,r)} \max_{y \in \mathcal{Y}} F(x, y) = \frac{1}{n} \sum_{i=1}^n f_i(x_i, y_i) \quad (1)$$

$$\text{s.t. } x_i = x_j, y_i = y_j, (i, j) \in \mathcal{E}$$

where the function $f_i(x, y) : \mathcal{M} \times \mathcal{Y} \rightarrow \mathbb{R}$ is nonconvex

over $x \in \mathcal{M}$ and is μ -strongly concave over $y \in \mathcal{Y} \subseteq \mathbb{R}^d$. Here $\mathcal{Y} \subseteq \mathbb{R}^d$ is nonempty compact convex sets but $\mathcal{M} := \text{St}(d, r) = \{x \in \mathbb{R}^{d \times r} : x^\top x = I_r\}$ is the Stiefel manifold. Unlike the Euclidean distributed setting, the Stiefel manifold is a non-convex set. \mathcal{E} is the communication link set in the decentralized communication network. We consider a multi-client network system $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, \dots, n\}$ is the collection of work nodes indices and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the collection of communication links (i, j) , $i, j \in \mathcal{V}$, that connect worker nodes. Each worker node i possesses local data and updates the local model.

In this paper, we also focus on the stochastic form of decentralized minimax problem (1), defined as

$$\min_{x \in \text{St}(d, r)} \max_{y \in \mathcal{Y}} F(x, y) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[f_i(x, y; \xi)] \quad (2)$$

s.t. $x_i = x_j, y_i = y_j, (i, j) \in \mathcal{E}$

where ξ is a random variable following an unknown distribution \mathcal{D}_i . Many important applications can be written in the form of (1) and (2), such as robust geometry-aware PCA (Zhang, Zhang, and Sra 2022), distributionally robust DNN training with orthogonal weights (Huang and Gao 2020). PCA has a wide range of applications, and it can be used K-means clustering (Zhang, Zhang, and Sra 2022) and study the functional diversity of cells (Baden et al. 2016). For deep learning, orthonormality on parameters has been proven to increase generalization precision and speed up and stabilize convergence of neural network models (Wang et al. 2020; Huang et al. 2018), which corresponds to optimization over the Stiefel manifold. Although above minimax problems (1) and (2) have been frequently appeared in many machine learning applications, (1) and (2) in distributed settings have not been theoretically studied. The existing explicitly decentralized minimax optimization methods such as the decentralized gradient descent ascent method only focus on the minimax problems in Euclidean space.

The change from Euclidean to Riemannian nonconvex optimization and from centralized to distributed settings pose several difficulties. For example, we cannot use the arithmetic mean as a consensus step on the Stiefel manifold because the average result might not be on the Stiefel manifold. Under the Riemannian PL condition or g -convex-concave (Jordan, Lin, and Vlatakis-Gkaragkounis 2022; Han et al. 2022), variables x and y can be embedded as $z = (x^\top, y^\top)$ and solve two variables together, which cannot be applied in our nonconvex case. Then there exists a natural question:

Can we design new decentralized algorithms with exact convergence for distributed minimax optimization on the Stiefel manifold for solving the Problems (1) and (2)?

In the paper, we give an affirmative answer to the above question and propose a class of gradient descent ascent methods for solving the Problems (1) and (2). Our main contributions are three-fold:

1) We propose the first decentralized gradient descent ascent algorithm (DRGDA) over the Stiefel manifold for

the deterministic minimax problem (1). Our method is retraction-based and vector transport-free. We prove that the DRGDA has a sample complexity of $O(\epsilon^{-2})$ for finding an ϵ -stationary point, which matches the best complexity result achieved by the centralized minimax method on Euclidean space.

- 2) We also propose a decentralized stochastic first-order algorithm (DRSGDA) over the Stiefel manifold for the stochastic problem (2), which has a sample complexity of $O(\epsilon^{-4})$, which is the first theoretical result of the stochastic method in the decentralized Riemannian minimax problems.
- 3) Extensive experimental results on the robust DNN training over the Stiefel manifold confirm the effectiveness of our proposed algorithm.

Related Works

Minimax Optimization and Decentralized Distributed Learning

The minimax optimization recently has been successfully applied in many machine learning problems such as robust federated learning, reinforcement learning (Wai et al. 2019), and adversarial training (Wu, Bojchevski, and Huang 2022). Many gradient-based minimax methods were proposed for solving these minimax optimization problems (Lin, Jin, and Jordan 2020; Lin et al. 2020). In addition, some works developed accelerated gradient descent ascent algorithms with the variance-reduced techniques (Luo, Ye, and Zhang 2020). More recently, Huang, Wu, and Huang (2021) proposed mirror descent ascent methods to solve the nonsmooth nonconvex-strongly-concave minimax problems based on dynamic mirror functions. Huang, Wu, and Hu (2023) used momentum-based variance-reduced technique to design the single-loop algorithm (Acc-MDA).

At the same time, many decentralized distributed methods are proposed because they are useful when the network has communication bottlenecks or limited bandwidth between certain nodes and the central node. Lian et al. (2017) firstly provides theoretical support for the potential advantage of a decentralized distributed algorithm. Subsequently, D-SPIDER-SFO (Pan, Liu, and Wang 2020), DEEPSTORM (Mancino-Ball et al. 2022) and many other methods are proposed for the decentralized training.

Recently, Xian et al. (2021) and Zhang et al. (2021) extended the minimax problem to a decentralized distributed setting and solve the nonconvex-strongly-concave minimax problems in the Euclidean spaces. Zhang et al. (2021) proposed both deterministic gradient-based decentralized minimax methods and stochastic gradient-based decentralized minimax methods.

Riemannian Manifold Optimization

Riemannian manifold optimization is an important problem given that they are related to many machine learning problems, such as low-rank matrix completion, dictionary learning and deep learning (Huang et al. 2018). Liu et al. (2017) proposed first-order gradient methods for geodesically convex functions. Furthermore, Zhang, J Reddi, and

Sra (2016) applied variance-reduced technology to Riemannian manifold optimization and proposed fast stochastic methods. More recently, various fast Riemannian gradient-based methods have been proposed for non-convex optimization based on the retraction and vector transport (Kasai, Sato, and Mishra 2018). Then some efficient Riemannian adaptive optimization algorithms are introduced (Kasai, Jawanpuria, and Mishra 2019).

Recently, Huang and Gao (2020) proposed a class of Riemannian gradient descent ascent algorithms to solve the nonconvex-strongly-concave minimax problems on Riemannian manifolds. Then Han et al. (2022) introduced Riemannian Hamiltonian function and proposed Riemannian Hamiltonian gradient methods under the Riemannian PL condition. Zhang, Zhang, and Sra (2022) and Jordan, Lin, and Vlatakis-Gkaragkounis (2022) proved that the Riemannian corrected extragradient (RCEG) method achieves convergence at a linear rate in the geodesically strongly-convex-concave problems. In addition, some methods are proposed for variational inequalities, which are the implicit minimax problems, on Riemannian manifolds (Huang and Gao 2020).

In addition, several algorithms were proposed for distributed optimization on Riemannian manifolds (Chen et al. 2021). In the decentralized distributed setting, the consensus step is usually done by calculating the Karcher mean on the manifold (Shah 2017), or calculating the minimizer of the sum of the square of the Euclidean distances in the corresponding sub-manifold case (Chen et al. 2021). Grammenos et al. (2020) firstly present the algorithm for federated learning on Riemannian manifolds and focus on the PCA problems (Stiefel manifold).

Notations

We use \mathbf{x} and \mathbf{y} to denote a collection of all local parameters x_i, y_i , i.e., $\mathbf{x}^\top = (x_1^\top, x_2^\top, \dots, x_N^\top)$ and $\mathbf{y}^\top = (y_1^\top, y_2^\top, \dots, y_N^\top)$. \otimes denotes Kronecker product. Define $\bar{\mathbf{x}} = \mathbf{1} \otimes \bar{x}$. $\|\cdot\|$ denotes the ℓ_2 norm for vectors and Frobenius norm for matrices, respectively. $\nabla_x f(x, y)$ and $\nabla_y f(x, y)$ denote the Euclidean partial derivatives w.r.t. variables x and y respectively, and the Riemannian gradient is $\text{grad}_x g(x, y)$. For \mathcal{M} , the n -fold Cartesian product of \mathcal{M} is denoted as $\mathcal{M}_n = \mathcal{M} \times \dots \times \mathcal{M}$. I_r denotes r -dimension identity matrix. Given the mini-batch samples $\mathcal{B} = \{\xi_i\}_{i=1}^q$, we let $\nabla f(x, y; \mathcal{B}) = \frac{1}{q} \sum_{i=1}^q \nabla f(x, y; \xi_i)$. In general, the manifold \mathcal{M} is endowed with a smooth inner product $\langle x, y \rangle_x : \mathbb{T}_x \mathcal{M} \times \mathbb{T}_x \mathcal{M} \rightarrow \mathbb{R}$ on tangent space $\mathbb{T}_x \mathcal{M}$ for every $x \in \mathcal{M}$.

The network system $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is represented by double stochastic matrix W , which is defined as follows: (1) if there exists a link between node i and node j , then $W_{ij} > 0$, otherwise $W_{ij} = 0$, (2) $W = W^\top$ and (3) $W\mathbf{1} = \mathbf{1}$ and $\mathbf{1}^\top W = \mathbf{1}^\top$. The second-largest eigenvalue of W as λ .

Preliminaries

Riemannian Gradients

Given that the metric on tangent space $\mathbb{T}_x \mathcal{M}$ is induced from the Euclidean inner product $\langle \cdot, \cdot \rangle$, the Riemannian

gradient $\text{grad}h(x)$ on $\text{St}(d, r)$ is defined as $\text{grad}h(x) = \mathcal{P}_{\mathbb{T}_x \mathcal{M}} \nabla h(x)$, where $\mathcal{P}_{\mathbb{T}_x \mathcal{M}}$ is the orthogonal projection onto $\mathbb{T}_x \mathcal{M}$. On the $\text{St}(d, r)$, we have

$$\mathcal{P}_{\mathbb{T}_x \mathcal{M}} y = y - \frac{1}{2} x x^\top y + y^\top x \quad (3)$$

Retraction

Given $x \in \mathcal{M}$ and $u \in \mathbb{T}_x \mathcal{M}$, retraction $R_x : \mathbb{T}_x \mathcal{M} \rightarrow \mathcal{M}$ is defined as mapping from tangent space $\mathbb{T}_x \mathcal{M}$ onto \mathcal{M} with a local rigidity condition that preserves the gradients at $x \in \mathcal{M}$. The retraction R_x satisfies: 1) $R_x(0) = x$. 2) $DR_x(0) = id_{\mathbb{T}_x \mathcal{M}}$.

In the centralized Riemannian gradient descent iteration, the update step is defined as

$$x_{t+1} = R_{x_t}(-\alpha \text{grad}h(x_t)), \quad (4)$$

namely, it updates parameters along a negative Riemannian gradient direction on the tangent space, then it uses an operation called retraction R_{x_t} to map onto \mathcal{M} . Following (Chen et al. 2021), we also assume the second-order boundedness of retraction R_x , namely,

$$R_x(u) = x + u + O(\|u\|^2). \quad (5)$$

Furthermore, following (Boumal, Absil, and Cartis 2019; Chen et al. 2021), we have the following lemma:

Lemma 1. (Boumal, Absil, and Cartis 2019) *Let R be a second-order retraction over $\text{St}(d, r)$, we have*

$$\|R_x(u) - (x + u)\| \leq M \|u\|^2 \quad \forall x \in \text{St}(d, r), \forall u \in \mathbb{T}_x \mathcal{M} \quad (6)$$

Moreover, if the retraction is the polar decomposition. For all $x \in \text{St}(d, r)$ and $u \in \mathbb{T}_x \mathcal{M}$, the following inequality holds :

$$\|R_x(u) - z\| \leq \|x + u - z\|, \forall z \in \text{St}(d, r) \quad (7)$$

Unless otherwise stated, the sequel uses the polar retraction to offer a concise analysis. More details on the polar retraction are provided in (Liu, So, and Wu 2019).

Mean and Consensus on the Stiefel Manifold

In the Euclidean decentralized distributed algorithms, the $\|x_i - \bar{x}\|^2$ is usually used as the consensus error, where \bar{x} is defined as the Euclidean average point of x_1, \dots, x_n by

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (8)$$

However, in the decentralized distributed Riemannian setting, such as $\text{St}(d, r)$, we cannot directly average local parameters as in (8) because the arithmetic mean result might not be on the manifold. Therefore, we follow (Sarlette and Sepulchre 2009) and use induced arithmetic mean (IAM) on $\text{St}(d, r)$, defined as

$$\hat{x} := \underset{z \in \text{St}(d, r)}{\text{argmin}} \sum_{i=1}^n \|z - x_i\|^2 = \mathcal{P}_{\text{St}}(\bar{x}) \quad (9)$$

Algorithm 1: DRGDA Algorithm

- 1: **Input:** T , tuning parameters $\{\alpha, \beta, \eta\}$;
 - 2: **initialize:** Initial input $x_0^i \in \text{St}(d, r)$, $y_0^i \in \mathcal{Y}$, and then compute $u_0^i = \text{grad}_x f_i(x_0^i, y_0^i)$ and $v_0^i = \nabla_y f_i(x_0^i, y_0^i)$, where $i \in [n]$;
 - 3: **for** $t = 0, 1, \dots, T$ **do**
 - 4: $x_{t+1}^i = R_{x_t^i}(\mathcal{P}_{T_{x_t^i} \mathcal{M}}(\alpha \sum_{j=1}^n \mathbf{W}_{ij}^k x_t^j) - \beta w_t^i)$ where $w_t^i = \mathcal{P}_{T_{x_t^i} \mathcal{M}}(u_t^i)$
 - 5: $y_{t+1}^i = \sum_{j=1}^n \mathbf{W}_{ij}^k y_t^j + \eta v_t^i$;
 - 6: $u_{t+1}^i = \sum_{j=1}^n \mathbf{W}_{ij}^k u_t^j + \text{grad}_x f_i(x_{t+1}^i, y_{t+1}^i) - \text{grad}_x f_i(x_t^i, y_t^i)$;
 - 7: $v_{t+1}^i = \sum_{j=1}^n \mathbf{W}_{ij}^k v_t^j + \nabla_y f_i(x_{t+1}^i, y_{t+1}^i) - \nabla_y f_i(x_t^i, y_t^i)$;
 - 8: **end for**
-

where $\mathcal{P}_{\text{St}(\cdot)}$ is the orthogonal projection onto $\text{St}(d, r)$. And similarly we have

$$\min_{z \in \text{St}(d, r)} \sum_{i=1}^n \frac{1}{n} \|z - x_i\|^2 = \frac{1}{n} \|\mathbf{x} - \hat{\mathbf{x}}\|^2 \quad (10)$$

Furthermore, we also define

$$\|\mathbf{x} - \hat{\mathbf{x}}\|_\infty = \max_{i \in [n]} \|x_i - \hat{x}\| \quad (11)$$

As for the consensus problem over the Stiefel manifold, following the (Chen et al. 2021), we use the minimizer of the sum of the square of the l_2 norm over $\text{St}(d, r)$. It means:

$$\begin{aligned} \min J^k(\mathbf{x}) &:= \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n W_{ij}^k \|x_i - x_j\|^2 \quad (12) \\ \text{s.t. } &x_i \in \mathcal{M}, \forall i \in [n] \end{aligned}$$

where the superscript $k \geq 1$ is an integer. It should be mentioned we use k to offer flexibility in algorithm design and analysis. And \mathbf{W}_{ij}^k corresponds to performing k steps of communication on the tangent space.

Decentralized Riemannian Gradient Descent Ascent Methods

In the section, we propose deterministic and stochastic methods for solving decentralized Riemannian minimax Problems in (1) and (2), respectively.

DRGDA Algorithms

In this subsection, we first study the decentralized Riemannian gradient descent ascent (DRGDA) method with gradient tracking technologies for solving the problem (1). We describe our algorithm in Algorithm 1.

At the beginning of Algorithm 1, one can simply initialize local model parameters $x \in \mathcal{M}$ and $y \in \mathcal{Y}$ for all worker nodes with the same points.

At the step 4 of Algorithm 1, we use the decentralized Riemannian gradient method to update variable \mathbf{x} based on

the gradient tracker u_t^i . We firstly need to project the gradient tracker u_t^i onto the tangent space $T_{x_t^i} \mathcal{M}$, which follows a retraction consensus update (Seen in the supplementary materials). We can regard the step 4 as applying Riemannian gradient method to solve the following problem:

$$\min_{\mathbf{x} \in \text{St}(d, r)} \beta F(\mathbf{x}) + \alpha J^k(\mathbf{x}) \quad (13)$$

where $J^k(\mathbf{x})$ is the consensus problem defined in (12) and more details seen in the supplementary materials. The consensus step and computation of Riemannian gradient can be done in parallel (Chen et al. 2021).

At the step 5 of Algorithm 1, we follow conventional decentralized gradient ascent steps, where we update local parameters with gradient tracker and then do the consensus step with double stochastic matrix \mathbf{W} .

At the step 6 of Algorithm 1, we use u_t^i to tracks the average Riemannian gradient $\text{grad}_x f_i(x_t^i, y_t^i)$. Gradient tracking is a popular technology in the decentralized distributed learning to guarantee the convergence of the global objective function. Note, we can regard $\text{grad}_x f_i(x_t^i, y_t^i)$ at the step 6 as the projected Euclidean gradient, and we do not need to project it on the tangent space $T_{x_t^i} \mathcal{M}$ to save the computation cost, and we only do the projection at the step 4.

At the step 7 of Algorithm 1, we use the similar operation. Each worker node calculates the local full gradients, and then we update the gradient tracker v_t^i for y .

DRSGDA Algorithms

We also consider stochastic method because Stochastic gradient descent are widely used in learning tasks (Sun et al. 2022). In the DRGDA algorithm, worker nodes need to compute local full gradients in each iteration, which result in a high sample complexity when the local datasets are large. It also cannot be applied in the online learning. This limitation motivates us to leverage the stochastic method and propose Riemannian decentralized stochastic gradient descent ascent (DRSGDA) method. Algorithm 2 show the algorithmic framework of DRSGDA method.

At the step 4 and step 5 of Algorithm 2, we use the similar step to update variable $\text{St}(d, r)$ and $y \in \mathcal{Y}$. The different is that we use the \underline{u}_t and \underline{v}_t , which is the gradient tracker for the unbiased gradient estimators. At the steps 7 and 8 of Algorithm 2, we calculate the stochastic Riemannian gradient for variable \mathbf{x} , and calculate the stochastic gradient for variable \mathbf{y} based on sampled data to update the \underline{u}_t and \underline{v}_t .

Convergence Analysis

Assumptions

In the subsection, we state assumptions for the problems.

Assumption 1. Each component function $f_i(x, y)$ is twice continuously differentiable in both $x \in \text{St}(d, r)$ and $y \in \mathcal{Y}$. For $i \in [n]$, we have

$$\begin{aligned} \|\nabla_x f_i(x_1, y) - \nabla_x f_i(x_2, y)\| &\leq L_{11} \|x_1 - x_2\|, \\ \|\nabla_x f_i(x, y_1) - \nabla_x f_i(x, y_2)\| &\leq L_{12} \|y_1 - y_2\|, \\ \|\nabla_y f_i(x_1, y) - \nabla_y f_i(x_2, y)\| &\leq L_{21} \|x_1 - x_2\|, \\ \|\nabla_y f_i(x, y_1) - \nabla_y f_i(x, y_2)\| &\leq L_{22} \|y_1 - y_2\|. \end{aligned}$$

Algorithm 2: DRSGDA Algorithm

- 1: **Input:** T , tuning parameters $\{\alpha, \beta, \eta\}$;
 - 2: **initialize:** Initial input $x_0^i \in \text{St}(d, r)$, $y_0^i \in \mathcal{Y}$, and then compute $u_0^i = \text{grad}_x f_i(x_0^i, y_0^i; \mathcal{B}_0^i)$ and $v_0^i = \nabla_y f_i(x_0^i, y_0^i; \mathcal{B}_0^i)$, where $i \in [n]$;
 - 3: **for** $t = 0, 1, \dots, T$ **do**
 - 4: $x_{t+1}^i = R_{x_t^i}(\mathcal{P}_{\mathbb{T}_{x_t^i} \mathcal{M}}(\alpha \sum_{j=1}^n \mathbf{W}_{ij}^k x_t^j) - \beta w_t^i)$ where $w_t^i = \mathcal{P}_{\mathbb{T}_{x_t^i} \mathcal{M}}(\underline{u}_t^i)$
 - 5: $y_{t+1}^i = \sum_{j=1}^n \mathbf{W}_{ij}^k y_t^j + \eta v_t^i$
 - 6: Draw a mini-batch i.i.d. samples $\mathcal{B}_{t+1}^i = \{\xi_j^i\}_{j=1}^q$, and then compute
 - 7: $\underline{u}_{t+1}^i = \sum_{j=1}^n \mathbf{W}_{ij}^k u_t^j + \text{grad}_x f_i(x_{t+1}^i, y_{t+1}^i; \mathcal{B}_{t+1}^i) - \text{grad}_x f_i(x_t^i, y_t^i; \mathcal{B}_t^i)$;
 - 8: $\underline{v}_{t+1}^i = \sum_{j=1}^n \mathbf{W}_{ij}^k v_t^j + \nabla_y f_i(x_{t+1}^i, y_{t+1}^i; \mathcal{B}_{t+1}^i) - \nabla_y f_i(x_t^i, y_t^i; \mathcal{B}_t^i)$;
 - 9: **end for**
-

and we also assume the uniform upper bound of $\|\text{grad}_x f_i(x_t^i, y_t^i)\|$ is D .

The L -smooth assumption is widely used in the minimax problems in the Euclidean space. Since the Stiefel manifold is embedded in Euclidean space, we are free to do so, and detailed comparison is provided by (Chen et al. 2021).

One could also consider the following Lipschitz inequality (Huang and Gao 2020),

$$\|\mathbb{P}_{x \rightarrow y} \text{grad} f(x) - \text{grad} f(y)\|_F \leq L'_g d_g(x, y) \quad (14)$$

Because it involves vector transport and geodesic distance, which add computational burden, we choose to use assumption 1 for simplicity. And $L_{11}, L_{12}, L_{21}, L_{22}$ are all constants and we set $L = \max\{L_{11}, L_{12}, L_{21}, L_{22}\}$. Since global function $F(x, y)$ is a finite sum of local functions, it is also the L -smooth.

We also bound the partial derivative with respect to variable x as (Chen et al. 2021). It is reasonable since the Stiefel manifold is compact and often satisfied in practice. One common example is the finite-sum form of function.

Assumption 2. The function $f_i(x, y)$ is μ -strongly concave in $y \in \mathcal{Y}$, i.e., for all $x \in \text{St}(d, r)$ and $y_1, y_2 \in \mathcal{Y}$, then the following inequality holds

$$f_i(x, y_1) \leq f_i(x, y_2) + \langle \nabla_y f_i(x, y_2), y_1 - y_2 \rangle - \frac{\mu}{2} \|y_1 - y_2\|^2. \quad (15)$$

Then we could easily get $F(x, y)$ is also μ -strongly concave in $y \in \mathcal{Y}$. Then, there exists a unique solution to the problem $\max_{y \in \mathcal{Y}} F(x, y)$ for any x . Here we let $y^*(x) = \arg \max_{y \in \mathcal{Y}} F(x, y)$ and $\Phi(x) = F(x, y^*(x)) = \max_{y \in \mathcal{Y}} F(x, y)$.

Assumption 3. The function $\Phi(x)$ is bounded below in $\text{St}(d, r)$, i.e., $\Phi^* = \inf_{x \in \text{St}(d, r)} \Phi(x) > -\infty$.

A Useful Convergence Metric

Since finding the global optimal solution for the non-convex min-max problem is NP-hard in general (Lin, Jin, and Jordan

2020), we firstly introduce a useful convergence metric to measure convergence of our algorithms for both deterministic and stochastic settings. Given the variable sequence generated from our algorithms, we define a convergence metric as follows:

$$\mathfrak{M}_t = \|\text{grad}_x F(x_t, y_t)\| + \frac{1}{n} \|x_t - \hat{x}_t\| + \frac{L}{n} \|\bar{y}_t - y^*(\hat{x}_t)\|, \quad (16)$$

where the first and second terms of \mathfrak{M}_t measure the optimality gap of non-convex minimax problems, while the last term measures the convergence of \bar{y}_t to the unique maximizer of $F(\hat{x}_t, \cdot)$

Convergence Analysis of DRGDA Algorithms

In the subsection, we study the convergence properties of our DRGDA algorithm.

Lemma 2. Under the above assumptions, the gradient of function $\Phi(x) = \max_{y \in \mathcal{Y}} F(x, y)$ is L_Φ -Lipschitz with respect to retraction, and the mapping or function $y^*(x) = \arg \max_{y \in \mathcal{Y}} f(x, y)$ is κ -Lipschitz with respect to retraction. $\forall x_1, x_2 \in \mathcal{M}$ and $u \in \mathbb{T}_{x_1} \mathcal{M}$, we have

$$\|y^*(x_1) - y^*(x_2)\| \leq \kappa \|x_1 - x_2\|$$

$$\|\text{grad} \Phi(x_1) - \text{grad} \Phi(x_2)\| \leq L_\Phi \|x_1 - x_2\|$$

where $L_\Phi = \kappa L_{12} + L_{11}$, and $\kappa = L_{21}/\mu$ denotes the number condition.

The lemma 2 guarantees that $\Phi(x)$ is L_Φ -Lipschitz and $y^*(\cdot)$ is κ -Lipschitz, and it provides the theoretical support to use $\Phi(x)$ in the proof.

Theorem 1. Under Assumptions (1, 2 and 3) Let $k \geq \lceil \log_{\lambda_2}(\frac{1}{2\sqrt{n}}) \rceil$, $\alpha \leq \frac{1}{M}$, $0 < \beta \leq \min\{\beta_0, \frac{1}{L+9C_2}, \frac{\alpha\delta_1}{10D}\}$ and $\eta \leq \frac{1}{L}$. And suppose the sequence generated from Algorithm 1 and the initial parameters start from the same points, we have

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathfrak{M}_t \leq \frac{10(H_0 - H_T + C_4)}{T\beta} + \frac{10\beta(H_0 - H_T + C_4)C_3}{T} \quad (17)$$

where β_0 and Q is introduced in the proof of Theorem 1, and we let

$$H_t = \Phi(\hat{x}_t) + \frac{8L}{\mu\eta} \|\bar{y}_{t+1} - y^*(\hat{x}_{t+1})\|^2 + \frac{Q}{n} \|\mathbf{y}_t - \bar{\mathbf{y}}_t\|^2 + \frac{\eta Q}{n} \|\mathbf{v}_t - \bar{\mathbf{v}}_t\|^2$$

and $C_1 = \frac{1}{(1-\rho_k)^2}$; $C_2 = [\frac{1}{2\eta} + \frac{L_\Phi}{L} + \frac{40\kappa^2 L}{\mu^2 \eta^2} + (1 + \frac{1}{c_0})6L^2\eta Q]$; $C_3 = \frac{2}{(1-\rho_k)^2}$; $C_4 = \frac{C_3 C_7 \delta_1^2 \alpha^2}{10}$; $C_5 = (2MD + L + C_2(5\delta_1^2 + 5))$; $C_6 = \frac{100rD^2 C_1}{L} + 2\eta L^2 + \frac{67L^2}{\mu^2} + 6L^2\eta Q + C_2 101\alpha^4 C_1 \delta_1^2 + 8MD\alpha^2 + 2D\alpha + L$; $C_7 = 2L^2\eta Q + \frac{100rD^2 C_1}{L}$; $C_8 = \frac{2}{(1-\sigma_2^2)}$

Remark 1. To achieve an ϵ^2 -stationary solution, i.e., $\frac{1}{T} \sum_{t=0}^{T-1} \mathfrak{M}_t \leq \epsilon^2$, the total samples evaluated across the network system is on the order of $O(m\epsilon^{-2})$, where m is the total number of data points.

Convergence Analysis of DRSFDA Algorithms

Assumption 4. Each component function $f_i(x, y; \xi)$ has an unbiased stochastic gradient with bounded variance σ^2 , i.e., $\forall i \in [n], x \in \text{St}(d, r), y \in \mathcal{Y}$

$$\begin{aligned} \mathbb{E}[\nabla f_i(x, y; \xi)] &= \nabla f_i(x, y), \\ \mathbb{E}\|\text{grad}_x f_i(x, y; \xi) - \text{grad}_x f_i(x, y)\|^2 &\leq \sigma_1^2 \\ \mathbb{E}\|\nabla_y f_i(x, y; \xi) - \nabla_y f_i(x, y)\|^2 &\leq \sigma_2^2. \end{aligned}$$

and we define $\sigma = \max\{\sigma_1, \sigma_2\}$. In addition, we also assume a uniform upper bound of $\|\text{grad}_x f_i(x_t^i, y_t^i; \mathcal{B})\|$ is D .

Assumption 4 imposes the bounded variance of stochastic (Riemannian) gradients, which is commonly used in the stochastic optimization (Lin, Jin, and Jordan 2020; Huang and Gao 2020).

Theorem 2. Under Assumptions (1, 2, 3 and 4) Let $k \geq \lceil \log_{\lambda_2}(\frac{1}{2\sqrt{n}}) \rceil$, $\alpha \leq \frac{1}{M}$, $0 < \beta \leq \min\{\beta_0, \frac{1}{L+9C_2}, \frac{\alpha\delta_1}{2(\bar{\gamma}\sigma+5D)}\}$ and $\eta \leq \frac{1}{L}$. And suppose the sequence generated from Algorithm 2

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \mathfrak{M}_t &\leq 10 \frac{\mathbb{E}H_0 - \mathbb{E}H_T + C_4}{T\beta} \\ &+ \frac{20\beta(\mathbb{E}H_0 - \mathbb{E}H_T + C_4)C_3}{T} + \frac{C_9\sigma^2}{B} \end{aligned} \quad (18)$$

where β_0 and Q is introduced in the proof of Theorem 2 and we let

$$\begin{aligned} H_t &= \Phi(\hat{x}_t) + \frac{8L}{\mu\eta} \|\bar{y}_{t+1} - y^*(\hat{x}_{t+1})\|^2 + \frac{Q}{n} \|\mathbf{y}_t - \bar{\mathbf{y}}_t\|^2 \\ &+ \frac{\eta Q}{n} \|\mathbf{v}_t - \bar{\mathbf{v}}_t\|^2 \end{aligned}$$

and $C_1 = \frac{1}{(1-\rho_k)^2}$; $C_2 = [\frac{1}{2\eta} + \frac{L_\Phi}{2} + \frac{40\kappa^2 L}{\mu^2 \eta^2} + (1 + \frac{1}{c_0})18L^2\eta Q]$; $C_3 = \frac{2}{(1-\rho_k)^2}$; $C_4 = \frac{C_3 C_7 \delta_1^2 \alpha^2}{4}$; $C_5 = (2MD + L + C_2(5\delta_1^2 + 7))$; $C_6 = \frac{100rD^2 C_1}{L} + 2\eta L^2 + \frac{67L^2}{\mu^2} + 6L^2\eta Q + C_2 101\alpha^4 C_1 \delta_1^2 + 8MD\alpha^2 + 2D\alpha + L$; $C_7 = 2L^2\eta Q + \frac{100rD^2 C_1}{L}$; $C_8 = \frac{2}{(1-\sigma_2^t)^2}$; $C_9 = [2C_3 + 4(C_5 + C_6 C_3 + C_7 C_3) + (48\alpha^2 C_3 C_8 + 12C_8)L^2]\beta\bar{\gamma} + 180\eta Q(1 + \frac{1}{c_0})\frac{\sigma^2}{B\beta} + \frac{195C_2\beta\bar{\gamma}^2\sigma^2}{B}$

Remark 2. To achieve an ϵ^2 -stationary solution, i.e., $\frac{1}{T} \sum_{t=0}^{T-1} \mathfrak{M}_t \leq \epsilon^2$, Let $B = T$, and we choose $T \geq \epsilon^{-2}$, the total samples evaluated across the network system is on the order of $O(\epsilon^{-4})$.

Numerical Experiments

We conducted numerical experiments to validate the efficiency of our algorithms on two tasks: 1) Orthonormal fair

classification networks and 2) distributionally robust optimization with orthonormal weights. In the deterministic setting, we use the GT-GDA (Zhang et al. 2021) as the baseline to compare with DRDGA. In the stochastic setting, we use GNSDA which is motivated by the GNSD algorithm (Lu et al. 2019), DM-HSGD (Xian et al. 2021), GT-SRVR (Zhang et al. 2021) as the comparison baselines for solving nonconvex-strongly-concave minimax problems. Since these methods were not designed for optimization on the Stiefel manifold, we add the retraction operation (projection-like) when we do the test experiments. The experiments are conducted using computers with 2.3 GHz Intel Core i9 CPUs and NVIDIA Tesla P40 GPUs. We set the number of worker nodes as $n = 20$ and use the ring-based topology as the communication network (Xian et al. 2021).

Orthonormal Fair Classification Networks

In the first task, we train orthonormal fair classification networks by minimizing the maximum loss over different categories. Orthonormality on parameters corresponds to optimization over the Stiefel manifold. We use Convolutional Neural Networks (CNNs) as classifiers, whose architectures are given in supplementary materials. In the experiment, we use the MNIST, Fashion-MNIST, and CIFAR-10 datasets as in (Huang, Wu, and Huang 2021). Following (Huang, Wu, and Huang 2021), we mainly focus on three categories in each dataset: digital numbers $\{0, 2, 3\}$ in the MNIST dataset, and T-shirt/top, Coat, and Shirt categories in the Fashion-MNIST dataset, and airplane, automobile, and bird in the CIFAR10 dataset. The datasets are evenly divided into disjoint sets across all worker nodes. Then we train this fair classifier by solving the following minimax problem:

$$\min_{w \in \text{St}(d, r)} \max\{\mathcal{L}_1(w), \mathcal{L}_2(w), \mathcal{L}_3(w)\}, \quad (19)$$

where $\mathcal{L}_1, \mathcal{L}_2$, and \mathcal{L}_3 are the cross-entropy loss functions corresponding to the samples in three different categories and w denotes the CNN model parameters. Problem (19) can be re-written as the following:

$$\begin{aligned} \min_{w \in \text{St}(d, r)} \max_{u \in \mathcal{U}} \left\{ \sum_{i=1}^3 u_i \mathcal{L}_i(w) - \rho \|u\|^2 \right\} \quad (20) \\ \text{s.t. } \mathcal{U} = \left\{ u \mid u_i \geq 0, \sum_{i=1}^3 u_i = 1 \right\}, \end{aligned}$$

where $\rho > 0$ is the tuning parameter, and u is a weight vector for different loss functions.

In the experiment, we use Xavier normal initialization to CNN layer. The grid search is used to tune parameters for all methods. For all datasets, we choose the $\{\alpha, \beta, \eta\}$ from the set $\{0.0001, 0.001, 0.005, 0.01\}$ for DRGDA and DRS-GDA. For other methods, we tune the learning rates from the set $\{0.0001, 0.001, 0.005, 0.01\}$. For DM-HSGD, we also set $\{\beta_x, \beta_y\}$ from the set $\{0.1, 0.9\}$. The batch sizes for MNIST and Fashion-MNIST are 100 while that for CIFAR-10 is 50. The initial batch sizes for GT-SRVR and DM-HSGD are set as 300.

The training progress for orthonormal fair classification networks is reported in Figures 1 and 2. We can see that our

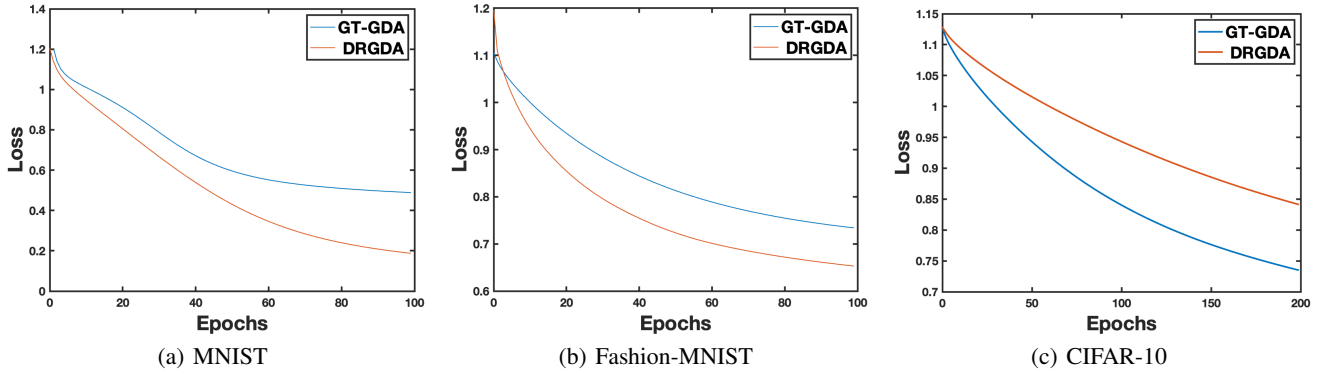


Figure 1: Results of different deterministic methods on the orthonormal fair classification networks task.

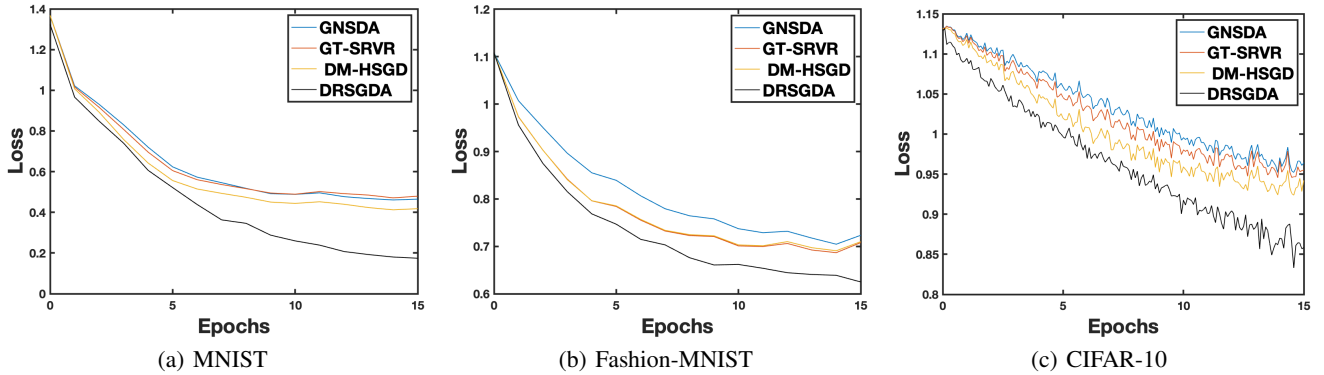


Figure 2: Results of different stochastic methods on the orthonormal fair classification networks task.

methods (DRGDA and DRSGDA) converge faster than the other comparison baselines on all datasets.

Distributionally Robust Optimization

In the second task, we focus on distributionally robust optimization with orthonormal weights which is defined as the following:

$$\min_{w \in \text{St}(d,r)} \max_{\mathbf{p} \in \Delta_n} \left\{ \sum_{i=1}^n p_i \ell_i(w; \xi_i) - \left\| \mathbf{p} - \frac{1}{n} \right\|^2 \right\} \quad (21)$$

where $\mathbf{p} = (p_1, \dots, p_n)$ is the i -th component of variable \mathbf{y} . Δ_n represents the simplex in \mathbb{R}^n . $\ell_i(w; \xi)$ denotes the loss function over the Stiefel manifold.

In this task, we also use the datasets MNIST, Fashion-MNIST, and CIFAR-10 datasets with the same DNN architecture provided in the supplementary materials. Similarly, the datasets are evenly divided by all worker nodes. We also use the grid search to tune parameters. For all methods, we tune the step size from the set $\{0.001, 0.01, 0.1, 0.5, 1\}$. For DM-HSGD, we set $\{\beta_x, \beta_y\}$ from the set $\{0.1, 0.9\}$. The batch size for all methods is set as 100. The initial batch size for GT-SRVR and DM-HSGD is set as 300.

The experimental results are reported in supplementary materials. From the results, we can see that our method DRSGDA achieves the best test accuracy and converges fastest compared with other baseline methods.

Conclusion

In the paper, we proposed decentralized Riemannian gradient descent ascent methods for solving the decentralized Riemannian minimax problems (1) and (2), respectively. We focus on non-convex strong-concave and $x \in \text{St}(d, r)$. In particular, our methods are the first algorithms for distributed minimax optimization with nonconvex constraints with exact convergence. For the deterministic setting, we introduced the decentralized gradient descent ascent algorithm (DRGDA) over the Stiefel manifold. We proved that the DRGDA has a sample complexity of $O(\epsilon^{-2})$ for finding an ϵ -stationary point, which matches the best complexity result achieved by the centralized minimax method on the Euclidean space. We also propose a decentralized stochastic first-order algorithm (DRSGDA) over the Stiefel manifold for the stochastic setting, which has a sample complexity of $O(\epsilon^{-4})$. This is the first theoretical result of the stochastic method in the decentralized Riemannian minimax problems.

Acknowledgements

This work was partially supported by NSF IIS 1838627, 1837956, 1956002, 2211492, CNS 2213701, CCF 2217003, DBI 2225775.

References

- Baden, T.; Berens, P.; Franke, K.; Román Rosón, M.; Bethge, M.; and Euler, T. 2016. The functional diversity of retinal ganglion cells in the mouse. *Nature*, 529(7586): 345–350.
- Bao, R.; Wu, X.; Xian, W.; and Huang, H. 2022. Doubly sparse asynchronous learning for stochastic composite optimization. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI*, 1916–1922.
- Boumal, N.; Absil, P.-A.; and Cartis, C. 2019. Global rates of convergence for nonconvex optimization on manifolds. *IMA Journal of Numerical Analysis*, 39(1): 1–33.
- Chen, S.; Garcia, A.; Hong, M.; and Shahrampour, S. 2021. Decentralized riemannian gradient descent on the stiefel manifold. In *International Conference on Machine Learning*, 1594–1605. PMLR.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Grammenos, A.; Mendoza Smith, R.; Crowcroft, J.; and Mascolo, C. 2020. Federated principal component analysis. *Advances in Neural Information Processing Systems*, 33: 6453–6464.
- Han, A.; Mishra, B.; Jawanpuria, P.; Kumar, P.; and Gao, J. 2022. Riemannian Hamiltonian methods for min-max optimization on manifolds. *arXiv preprint arXiv:2204.11418*.
- Huang, F.; and Gao, S. 2020. On Riemannian Gradient-Based Methods for Minimax Problems. *arXiv preprint arXiv:2010.06097*.
- Huang, F.; Wu, X.; and Hu, Z. 2023. AdaGDA: Faster Adaptive Gradient Descent Ascent Methods for Minimax Optimization. In *International Conference on Artificial Intelligence and Statistics*. PMLR.
- Huang, F.; Wu, X.; and Huang, H. 2021. Efficient mirror descent ascent methods for nonsmooth minimax problems. *Advances in Neural Information Processing Systems*, 34: 10431–10443.
- Huang, L.; Liu, X.; Lang, B.; Yu, A.; Wang, Y.; and Li, B. 2018. Orthogonal weight normalization: Solution to optimization over multiple dependent stiefel manifolds in deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Jordan, M. I.; Lin, T.; and Vlatakis-Gkaragkounis, E.-V. 2022. First-Order Algorithms for Min-Max Optimization in Geodesic Metric Spaces. *arXiv preprint arXiv:2206.02041*.
- Kasai, H.; Jawanpuria, P.; and Mishra, B. 2019. Riemannian adaptive stochastic gradient algorithms on matrix manifolds. In *International Conference on Machine Learning*, 3262–3271. PMLR.
- Kasai, H.; Sato, H.; and Mishra, B. 2018. Riemannian stochastic recursive gradient algorithm. In *International Conference on Machine Learning*, 2516–2524. PMLR.
- Lian, X.; Zhang, C.; Zhang, H.; Hsieh, C.-J.; Zhang, W.; and Liu, J. 2017. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. *Advances in Neural Information Processing Systems*, 30.
- Lin, T.; Jin, C.; and Jordan, M. 2020. On gradient descent ascent for nonconvex-concave minimax problems. In *International Conference on Machine Learning*, 6083–6093. PMLR.
- Lin, T.; Jin, C.; Jordan, M.; et al. 2020. Near-optimal algorithms for minimax optimization. *arXiv preprint arXiv:2002.02417*.
- Liu, H.; So, A. M.-C.; and Wu, W. 2019. Quadratic optimization with orthogonality constraint: explicit Łojasiewicz exponent and linear convergence of retraction-based line-search and stochastic variance-reduced gradient methods. *Mathematical Programming*, 178(1): 215–262.
- Liu, Y.; Shang, F.; Cheng, J.; Cheng, H.; and Jiao, L. 2017. Accelerated first-order methods for geodesically convex optimization on Riemannian manifolds. *Advances in Neural Information Processing Systems*, 30.
- Lu, S.; Zhang, X.; Sun, H.; and Hong, M. 2019. GNSD: A gradient-tracking based nonconvex stochastic algorithm for decentralized optimization. In *2019 IEEE Data Science Workshop (DSW)*, 315–321. IEEE.
- Luo, L.; Ye, H.; and Zhang, T. 2020. Stochastic Recursive Gradient Descent Ascent for Stochastic Nonconvex-Strongly-Concave Minimax Problems. *arXiv preprint arXiv:2001.03724*.
- Mancino-Ball, G.; Miao, S.; Xu, Y.; and Chen, J. 2022. Proximal Stochastic Recursive Momentum Methods for Nonconvex Composite Decentralized Optimization. *arXiv preprint arXiv:2211.11954*.
- Müller, J.; Klein, R.; and Weinmann, M. 2019. Orthogonal wasserstein gans. *arXiv preprint arXiv:1911.13060*.
- Pan, T.; Liu, J.; and Wang, J. 2020. D-SPIDER-SFO: A decentralized optimization algorithm with faster convergence rate for nonconvex problems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 1619–1626.
- Paty, F.-P.; and Cuturi, M. 2019. Subspace robust Wasserstein distances. In *International conference on machine learning*, 5072–5081. PMLR.
- Sarlette, A.; and Sepulchre, R. 2009. Consensus optimization on manifolds. *SIAM journal on Control and Optimization*, 48(1): 56–76.
- Shah, S. M. 2017. Distributed Optimization on Riemannian Manifolds for multi-agent networks. *arXiv preprint arXiv:1711.11196*.
- Sun, J.; Huai, M.; Jha, K.; and Zhang, A. 2022. Demystify Hyperparameters for Stochastic Optimization with Transferable Representations. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1706–1716.

- Wai, H.-T.; Hong, M.; Yang, Z.; Wang, Z.; and Tang, K. 2019. Variance reduced policy evaluation with smooth function approximation. *Advances in Neural Information Processing Systems*, 32: 5784–5795.
- Wang, J.; Chen, Y.; Chakraborty, R.; and Yu, S. X. 2020. Orthogonal convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11505–11515.
- Wu, X.; Huang, F.; Hu, Z.; and Huang, H. 2022. Faster Adaptive Federated Learning. *arXiv preprint arXiv:2212.00974*.
- Wu, Y.; Bojchevski, A.; and Huang, H. 2022. Adversarial Weight Perturbation Improves Generalization in Graph Neural Network. *arXiv preprint arXiv:2212.04983*.
- Wu, Y.; Zhang, H.; and Huang, H. 2022. RetrievalGuard: Provably Robust 1-Nearest Neighbor Image Retrieval. In *International Conference on Machine Learning*, 24266–24279. PMLR.
- Xian, W.; Huang, F.; Zhang, Y.; and Huang, H. 2021. A faster decentralized algorithm for nonconvex minimax problems. *Advances in Neural Information Processing Systems*, 34: 25865–25877.
- Zhang, H.; J Reddi, S.; and Sra, S. 2016. Riemannian SVRG: Fast stochastic optimization on Riemannian manifolds. *Advances in Neural Information Processing Systems*, 29.
- Zhang, P.; Zhang, J.; and Sra, S. 2022. Minimax in Geodesic Metric Spaces: Sion’s Theorem and Algorithms. *arXiv preprint arXiv:2202.06950*.
- Zhang, X.; Liu, Z.; Liu, J.; Zhu, Z.; and Lu, S. 2021. Taming Communication and Sample Complexities in Decentralized Policy Evaluation for Cooperative Multi-Agent Reinforcement Learning. *Advances in Neural Information Processing Systems*, 34: 18825–18838.