DIFFUSION OPTIMISTIC LEARNING FOR MIN-MAX OPTIMIZATION

H. Cai*, S. A. Alghunaim•, A. H. Sayed*

*École Polytechnique Fédérale de Lausanne, Switzerland
•Kuwait University, Kuwait

ABSTRACT

This work introduces and studies the convergence of a stochastic diffusion-optimistic learning (DOL) strategy for solving distributed nonconvex (NC) and Polyak–Lojasiewicz (PL) min-max optimization problems. Problems of this type are of interest due to a wide range of applications, including in generative adversarial networks (GANs), adversarial machine learning, and reinforcement learning. We prove that the DOL algorithm approaches an ε-stationary point through cooperation among agents following a left-stochastic communication protocol. The good performance of the proposed algorithm is illustrated by means of computer simulations.

***Index Terms*—** Minimax optimization, nonconvex-PL, optimistic algorithm, diffusion strategy.

1. INTRODUCTION

In recent years, there has been a surge of interest in studying minimax optimization problems due to their wide range of applications, ranging from generative adversarial networks (GANs) [1], to reinforcement learning [2], and adversarial learning [3].

There exists an extensive body of work on minimax optimization problems in the single agent scenario where the processing of the data is carried out at a *single* processing unit [4–19]. Several variations have been considered for this purpose. For example, some works studied strongly-convex strongly-concave (SC-SC) formulations [5, 6], while other works focused on convex-concave (C-C) formulations [7–9]. In addition, some results established the intractability of finding stationary points in the general nonconvex-nonconcave setting, even under smoothness conditions [4]. For this reason, many recent investigations have been concentrating on solving minimax problems under more tractable conditions. For instance, some works now assume a nonconvex-(strongly) concave setting [10–12] or a (PL)nonconvex-PL setting [16,17,20], while others consider the structured nonconvex-nonconcave setting under the assumption that the (weak) Minty variational inequality holds [13–15].

For *multi-agent* minimax optimization problems, on the other hand, the parameters/data are spread across multiple computing nodes, which cooperate with each other to solve the optimization problem [21–28]. Distributed multi-agent learning is advantageous to deal with the growing demand for training large-scale models, as well as to enable privacy-preserving and scalability features. In the multi-agent scenario, some works have already addressed (S)C-(S)C problems [21, 22]. Moreover, some works have also approached the problem by relying on variational inequalities [27, 28].

Nevertheless, the important scenarios involving nonconvexity in the primal variable and SC/PL in the dual variable over multi-agent systems remain largely unexplored. To our knowledge, there exist only a few distributed multi-agent works under this setting [23–26]. We summarize the distinction between our work and these earlier contributions as follows: i) We establish an $\mathcal{O}(\frac{1}{T^{1/2}})$ convergence rate for the primal objective using a double-call variant of optimistic gradient *without* the use of large batch sizes. Many existing results require large batch sizes where the stochastic gradient is computed and averaged over large data samples to guarantee convergence, see, e.g., [10, 12, 27, 28]. Moreover, other works require the strong assumption that the *stochastic loss function* is smooth (see remarks under Assumption 2) [23–26]; ii) we establish convergence for both the primal and dual objectives, while works such as [23–26] only demonstrate convergence for the primal objective; iii) we consider left-stochastic combination matrices, which are more general than the doubly-stochastic setting considered in [23–28], and enable a wider flexibility for the communication protocol; and iv) our theoretical results are established without the strong assumption that the *stochastic* loss functions should be smooth, thus broadening the scope of applicability beyond [23–26]. We further note that only the recent work [26] considers nonconvex-PL formulations.

Therefore, in this work, we study the stochastic multi-agent minimax problem, which assumes nonconvexity over the primal variable and PL over the dual variable. This setting encompasses nonconvex strong-concavity as a subproblem. The PL setting has attracted wide interest recently because it has been shown to hold in the neighborhood of the minima for over-parameterized neural networks [29]. We propose a decentralized optimistic diffusion method, which generalizes the optimistic gradient algorithm from game theory [30]. We show that the proposed method converges to a stationary point for both primal and dual objectives.

The remainder of this work is organized as follows. We present the problem and algorithm in Section 2, and we demonstrate the convergence results in Section 3. The perfor-

Emails: haoyuan.cai@epfl.ch, sulaiman.alghunaim@ku.edu.kw, ali.sayed@epfl.ch

mance of the algorithm is simulated in Section 4. Section 5 summarizes the conclusion of this work.

2. PROBLEM FORMULATION

We consider a network of K agents connected by a graph topology and used to solve the following minimax problem:

$$\min_{x\in\mathbb{R}^{M_1}}\max_{y\in\mathbb{R}^{M_2}} J(x,y)=\sum_{k=1}^{K}p_k J_k(x,y) \tag{1}$$

$$\text{where}\quad J_k(x,y)\triangleq \mathbb{E}_{\boldsymbol{\xi}_k}Q_k(x,y;\boldsymbol{\xi}_k) \tag{2}$$

Here, x,y are the global parameters, and each agent k has access to its local data set $\{\boldsymbol{\xi}_k\}$. The agents approximate their local risk functions $J_k:\mathbb{R}^{M_1}\times\mathbb{R}^{M_2}\to\mathbb{R}$ by using the loss functions $Q_k(x,y;\boldsymbol{\xi}_k)$. The scalars p_k are positive weights satisfying $\sum_{k=1}^{K}p_k=1$. In this work, we study the case where each $J_k(\cdot,y)$ is nonconvex and smooth in x, while the global objective $-J(x,\cdot)$ is ν-PL in y.

2.1. Diffusion Optimistic Learning (DOL)

Single agent scenario: We consider first the optimistic gradient method; this technique can be derived from the forward-reflected-backward framework [31], and it admits the following update rule at iteration i [30]:

$$x_i=\overline{x}_i-\mu_1\nabla_x J(x_{i-1},y_{i-1}) \tag{2a}$$

$$y_i=\overline{y}_i+\mu_2\nabla_y J(x_{i-1},y_{i-1}) \tag{2b}$$

$$\overline{x}_{i+1}=\overline{x}_i-\mu_1\nabla_x J(x_i,y_i) \tag{2c}$$

$$\overline{y}_{i+1}=\overline{y}_i+\mu_2\nabla_y J(x_i,y_i) \tag{2d}$$

where μ_1,μ_2 are step-size or learning rates, set by the designer. The algorithm starts from $i=0$, say, with initial conditions $\bar{x}_0=\bar{y}_0=x_{-1}=y_{-1}$. In the small step-size regime, the variables $\{x_{i+1},y_{i+1}\}$ will be close to $\{\overline{x}_{i+1},\overline{y}_{i+1}\}$. We can simplify (2a)-(2d) and rewrite them in terms of the variables x_i and y_i as follows [12, 18]:

$$x_{i+1}=x_i-\mu_1\Big(2\nabla_x J(x_i,y_i)-\nabla_x J(x_{i-1},y_{i-1})\Big) \tag{3a}$$

$$y_{i+1}=y_i+\mu_2\Big(2\nabla_y J(x_i,y_i)-\nabla_y J(x_{i-1},y_{i-1})\Big) \tag{3b}$$

Multi-agent scenario: In the distributed stochastic environment, each agent k utilizes a local sample $\{\boldsymbol{\xi}_k\}$ to approximate the true local gradient of the risk function by using its loss value. By integrating the adapt-then-combine (ATC) diffusion strategy [32, 33] into the update rules (3a)-(3b), we arrive at the **DOL** strategy listed in **Algorithm 1**. In this description, $\mathcal{N}_k$ denotes the set of neighbors of agent k, and $a_{\ell k}$ is the scaling weight for the information flowing from agent ℓ to agent k.

To study the convergence of the algorithm, we introduce the following network quantities:

$$\boldsymbol{x}_{c,i}\triangleq\sum_{k=1}^{K}p_k\boldsymbol{x}_{k,i}\in\mathbb{R}^{M\times 1}\quad\text{(centroid)} \tag{4a}$$

$$\boldsymbol{\mathcal{X}}_i\triangleq\text{col}\{\boldsymbol{x}_{1,i},\dots,\boldsymbol{x}_{K,i}\}\in\mathbb{R}^{MK\times 1} \tag{4b}$$

$$\boldsymbol{\mathcal{X}}_{c,i}\triangleq\text{col}\{\boldsymbol{x}_{c,i},\dots,\boldsymbol{x}_{c,i}\}\in\mathbb{R}^{MK\times 1} \tag{4c}$$

$$\boldsymbol{\mathcal{G}}_{x,i}\triangleq\text{col}\{\boldsymbol{g}_{x,1,i},\dots,\boldsymbol{g}_{x,K,i}\}\in\mathbb{R}^{MK\times 1} \tag{4d}$$

and similarly for $\boldsymbol{y}_{c,i}\in\mathbb{R}^{M\times 1},\boldsymbol{\mathcal{Y}}_i,\boldsymbol{\mathcal{Y}}_{c,i},\boldsymbol{\mathcal{G}}_{y,i}\in\mathbb{R}^{MK\times 1}$.

Algorithm 1 Diffusion Optimistic Learning (**DOL**)

Initialize: $\boldsymbol{x}_{k,i},\boldsymbol{y}_{k,i},\boldsymbol{g}_{x,k,i},\boldsymbol{g}_{y,k,i}$, $(i{=}-1,-2)$, step sizes μ_1,μ_2.

for $i=1,2,\dots$ **do**
  for each agent k **do**
    Compute stochastic gradient with sample $\boldsymbol{\xi}_{x,k,i},\boldsymbol{\xi}_{y,k,i}$
    $\boldsymbol{g}_{x,k,i-1}=2\nabla_x Q_k(\boldsymbol{x}_{k,i-1},\boldsymbol{y}_{k,i-1};\boldsymbol{\xi}_{x,k,i})-\nabla_x Q_k(\boldsymbol{x}_{k,i-2},\boldsymbol{y}_{k,i-2};\boldsymbol{\xi}_{x,k,i})$
    $\boldsymbol{g}_{y,k,i-1}=2\nabla_y Q_k(\boldsymbol{x}_{k,i-1},\boldsymbol{y}_{k,i-1};\boldsymbol{\xi}_{y,k,i})-\nabla_y Q_k(\boldsymbol{x}_{k,i-2},\boldsymbol{y}_{k,i-2};\boldsymbol{\xi}_{y,k,i})$
    Adaptation
    $\boldsymbol{\phi}_{k,i}=\boldsymbol{x}_{k,i-1}-\mu_1\boldsymbol{g}_{x,k,i-1}$
    $\boldsymbol{\psi}_{k,i}=\boldsymbol{y}_{k,i-1}+\mu_2\boldsymbol{g}_{y,k,i-1}$
    Combination
    $\boldsymbol{x}_{k,i}=\sum_{\ell\in\mathcal{N}_k}a_{\ell k}\boldsymbol{\phi}_{\ell,i},\quad \boldsymbol{y}_{k,i}=\sum_{\ell\in\mathcal{N}_k}a_{\ell k}\boldsymbol{\psi}_{\ell,i}$
  end for
end for

If we let $A=[a_{\ell k}]$ and $\mathcal{A}\triangleq A\otimes I$, where $\otimes$ stands for the Kronecker product operator, then **DOL** can be expressed using the network recursion:

$$\begin{aligned}\boldsymbol{\mathcal{X}}_i&=\mathcal{A}^\top\{\boldsymbol{\mathcal{X}}_{i-1}-\mu_1\boldsymbol{\mathcal{G}}_{x,i-1}\}\\ \boldsymbol{\mathcal{Y}}_i&=\mathcal{A}^\top\{\boldsymbol{\mathcal{Y}}_{i-1}+\mu_2\boldsymbol{\mathcal{G}}_{y,i-1}\}\end{aligned} \tag{5}$$

3. CONVERGENCE RESULTS

In this section, we list the assumptions used in our analysis and state the main convergence results. Due to space constraints, detailed proofs are omitted and will appear elsewhere e.g., in an arXiv preprint.

3.1. Assumptions

Assumption 1. *We assume each local risk function $J_k(x,y)$ is nonconvex in x while $-J(x,y)$ is ν-PL in y, i.e., for any x and y in the domain of $J(x,y)$, it holds that*

$$\|\nabla_y J(x,y)\|^2\geq 2\nu\Big\{\max_y J(x,y)-J(x,y)\Big\} \tag{6}$$

where ν is a strictly positive constant. □

Assumption 2. *We assume the gradients associated with each local risk function are L_f-Lipschitz, i.e.,*

$$\|\nabla_w J_k(x_1,y_1)-\nabla_w J_k(x_2,y_2)\|\quad (w=x\text{ or }y) \tag{7}$$

$$\leq L_f\Big(\|x_1-x_2\|+\|y_1-y_2\|\Big)$$ □

Remark 1. *We assume* (7) *holds for the local risk $J_k(x,y)$. This is in contrast to the works [23–26], where smoothness of the local loss $Q_k(x,y;\boldsymbol{\xi}_k)$ is also assumed.*

Since the goal of the primal problem is to minimize over the variable x, we introduce the primal objective as [23–26]:

$$P(x)\triangleq\max_y J(x,y) \tag{8}$$

Assumption 3. *We assume the primal function $P(x)$ is lower bounded, i.e., $P^\star=\inf_x P(x)>-\infty$.* □

Assumption 4. *The graph is strongly connected and the matrix $A=[a_{\ell,k}]$ is left-stochastic and primitive. As a result, it holds that $\mathbb{1}^\top A=\mathbb{1}^\top$ and A^r has strictly positive entries for some integer r [32, 33].* □

The last assumption ensures that A has a single maximum eigenvalue at 1 with Perron eigenvector p with positive entries and such that $Ap = p, \mathbb{1}^\top p = 1$. Furthermore, the combination matrix A admits a Jordan decomposition of the following form $A = VJV^{-1}$ where [32, 33]:

$$V = \left[p, V_R\right], J = \begin{bmatrix} 1 & 0 \\ 0 & J_\gamma \end{bmatrix}, V^{-1} = \begin{bmatrix} \mathbb{1}^\top \\ V_L^\top \end{bmatrix} \tag{9}$$

Here, the submatrix J_γ consists of Jordan blocks with arbitrarily small values γ on the lower diagonal, and where $\mathbb{1}^\top V_R = 0$, $V_L^\top V_R = I$, $\|J_\gamma\| < 1$. For convenience of the analysis, we let $\mathcal{J}_\gamma \triangleq J_\gamma \otimes I, \mathcal{V}_R \triangleq V_R \otimes I, \mathcal{V}_L \triangleq V_L \otimes I$.

Assumption 5. *We denote the filtration generated by the random processes by $\boldsymbol{\mathcal{F}}_i = \{(\boldsymbol{x}_{k,j}, \boldsymbol{y}_{k,j}) \mid k = 1, \ldots, K, j = -2, -1, \ldots, i\}$, and assume the stochastic gradients are unbiased with bounded variance:*

$$\mathbb{E}\Big\{\nabla_w Q_k(\boldsymbol{x}_{k,i}, \boldsymbol{y}_{k,i}; \boldsymbol{\xi}_{w,k,i}) \mid \boldsymbol{\mathcal{F}}_i\Big\} = \nabla_w J_k(\boldsymbol{x}_{k,i}, \boldsymbol{y}_{k,i})$$
$$\mathbb{E}\Big\{\|\nabla_w Q_k(\boldsymbol{x}_{k,i}, \boldsymbol{y}_{k,i}; \boldsymbol{\xi}_{w,k,i}) - \nabla_w J_k(\boldsymbol{x}_{k,i}, \boldsymbol{y}_{k,i})\|^2 \mid \boldsymbol{\mathcal{F}}_i\Big\}$$
$$\leq \sigma_k^2 \quad (w = x \text{ or } y) \tag{10}$$

Moreover, we assume the data samples $\boldsymbol{\xi}_{w,k,i}$ are independent of each other for all k, i and w. □

Assumption 6. *We assume the disagreement between the local and global gradients is bounded, i.e.,*

$$\|\nabla_w J_k(x, y) - \nabla_w J(x, y)\|^2 \leq G^2 \ (w = x \text{ or } y) \tag{11}$$

□

3.2. Main Results

We adopt the ε-approximate stationary condition as the convergence criterion, which is the standard criterion for this setting [10–12, 23–25]. The majority of single(multi)-agent minimax works [10–12, 23–25] focus on studying the convergence of the primal objective, i.e., on finding a point $\boldsymbol{x}^\star_{c,T_0}$ that satisfies

$$\mathbb{E}\|\nabla P(\boldsymbol{x}^\star_{c,T_0})\|^2 \leq \varepsilon^2 \tag{12}$$

for an arbitrarily small constant ε and after some iterations T_0. In contrast, we seek to find the network centroid $(\boldsymbol{x}^\star_{c,T_0}, \boldsymbol{y}^\star_{c,T_0+T_1})$ that satisfies the following first-order stationary conditions:

$$\mathbb{E}\|\nabla_x J(\boldsymbol{x}^\star_{c,T_0}, \boldsymbol{y}^\star_{c,T_0+T_1})\|^2 \leq \varepsilon^2 \tag{13a}$$
$$\mathbb{E}\|\nabla_y J(\boldsymbol{x}^\star_{c,T_0}, \boldsymbol{y}^\star_{c,T_0+T_1})\|^2 \leq \varepsilon^2 \tag{13b}$$

for an arbitrarily small constant ε and after some iterations $T_0 + T_1$. Note that the convergence condition (12) only guarantees the solution for the primal variable $\boldsymbol{x}$, since the primal objective $P(\boldsymbol{x})$ does not involve the direct computation of the dual variable $\boldsymbol{y}$. Nonetheless, both primal and dual solutions are essential, and we therefore need to establish the more challenging property (13a)-(13b).

We present the following main results.

Lemma 1. *Under **Assumptions** 2-6, the averaged network deviation is bounded as*

$$\begin{aligned}&\frac{1}{T}\sum_{i=0}^{T-1}\Big(\mathbb{E}\|\boldsymbol{\mathcal{X}}_i - \boldsymbol{\mathcal{X}}_{c,i}\|^2 + \mathbb{E}\|\boldsymbol{\mathcal{Y}}_i - \boldsymbol{\mathcal{Y}}_{c,i}\|^2\Big)\\ &\leq \mathcal{O}\Big(\frac{\eta^2(4KG^2 + \mu_1\sigma^2)\mu_2}{1 - \|\mathcal{J}_\gamma^\top\|^2}\Big)\end{aligned} \tag{14}$$

for sufficiently small step sizes μ_1, μ_2 and $\mu_1 \leq \mu_2$, where $\sigma^2 = \sum_{k=1}^K \sigma_k^2$ is the sum of the local gradient noise variances, and $\eta^2 = \|\mathcal{V}_R^\top\|^2\|\mathcal{V}_L^\top\|^2\|\mathcal{J}_\gamma^\top\|^2$ is a constant.

Theorem 1. *Under **Assumptions** 1-6, the expected gradient norm of the primal objective satisfies*

$$\begin{aligned}&\frac{1}{T}\sum_{i=0}^{T-1}\mathbb{E}\|\nabla P(\boldsymbol{x}_{c,i})\|^2\\ &\leq \mathcal{O}\Big(\frac{1}{\mu_1 T}\Big) + \mathcal{O}\Big(a^2\mu_2 \max_k p_k^2\Big) + \mathcal{O}\Big(e^2\mu_2^2 \max_k p_k^2\Big)\\ &\quad + \mathcal{O}\Big(\frac{d^2KG^2\mu_2}{1 - \|\mathcal{J}_\gamma^\top\|^2}\Big)\end{aligned} \tag{15}$$

for sufficiently large number of iterations T and sufficiently small step sizes $\mu_1 \ll \mu_2 < 1$, and where

$$\begin{aligned}&a^2 = \frac{L_f^3\sigma^2}{\nu^2}, \quad e^2 = K\sigma^2 L_f^2\Big(\frac{8L_f^2}{\nu^2} + 4\Big)\\ &d^2 = \Big(\frac{8L_f^2}{\nu^2} + 4\Big)L_f^2\|\mathcal{V}_R^\top\|^2\|\mathcal{V}_L\|^2\|\mathcal{J}_\gamma^\top\|^2\end{aligned} \tag{16}$$

are constants. We choose $\mu_1 = \mathcal{O}\Big(\frac{1 - \|\mathcal{J}_\gamma^\top\|}{T^{1/2}}\Big)$ and $\mu_2 = C\mu_1$ for sufficiently large constant C, and get the non-asymptotic bound for sufficiently large T:

$$\frac{1}{T}\sum_{i=0}^{T-1}\mathbb{E}\|\nabla P(\boldsymbol{x}_{c,i})\|^2 \leq \mathcal{O}\Big(\frac{1}{T^{1/2}}\Big) + \mathcal{O}\Big(\frac{1}{T}\Big) \tag{17}$$

*That is, **DOL** outputs an ε-stationary point after $T_0 = \mathcal{O}(\varepsilon^{-4})$ iterations and gradient evaluation complexity, i.e.,*

$$\begin{aligned}\mathbb{E}\|\nabla P(\boldsymbol{x}^\star_{c,T_0})\|^2 &= \inf_{i=0,\ldots,T_0-1}\mathbb{E}\|\nabla P(\boldsymbol{x}_{c,i})\|^2\\ &\leq \frac{1}{T_0}\sum_{i=0}^{T_0-1}\mathbb{E}\|\nabla P(\boldsymbol{x}_{c,i})\|^2 \leq \mathcal{O}(\varepsilon^2)\end{aligned} \tag{18}$$

Theorem 2. *Under **Assumptions** 1-6, we can find an ε-stationary point $(\boldsymbol{x}^\star_{c,T_0}, \boldsymbol{y}^\star_{c,T_0+T_1})$ satisfying* (13a)-(13b) *for small ε and after large number of iterations $T_0 + T_1 = \mathcal{O}(\varepsilon^{-4}) + \mathcal{O}(\varepsilon^{-4})$ with the parameter choices for μ_1, μ_2 shown in **Theorem** 1.*

Proof. We use the sequential approach to achieve (13a)-(13b), i.e., we first show that **DOL** finds a global primal variable $\boldsymbol{x}^\star_{c,T_0}$ that satisfies the following stationary condition:

$$\begin{aligned}\mathbb{E}\|\nabla P(\boldsymbol{x}^\star_{c,T_0})\|^2 &= \mathbb{E}\|\nabla_x J(\boldsymbol{x}^\star_{c,T_0}, \boldsymbol{y}^o(\boldsymbol{x}^\star_{c,T_0}))\|^2\\ &\leq \frac{\varepsilon^2}{4} = \mathcal{O}(\varepsilon^2)\end{aligned} \tag{19}$$

after $T_0 = \mathcal{O}(\varepsilon^{-4})$ iterations using the parameters considered in **Theorem 1**, where $\boldsymbol{y}^o(\boldsymbol{x}^\star_{c,T_0})$ denotes the optimal dual variable when fixing $\boldsymbol{x}^\star_{c,T_0}$. On the other hand, by halting the update of $\boldsymbol{x}^\star_{c,T_0}$, we can find a point $\boldsymbol{y}^\star_{c,T_0+T_1}$ that satisfies the

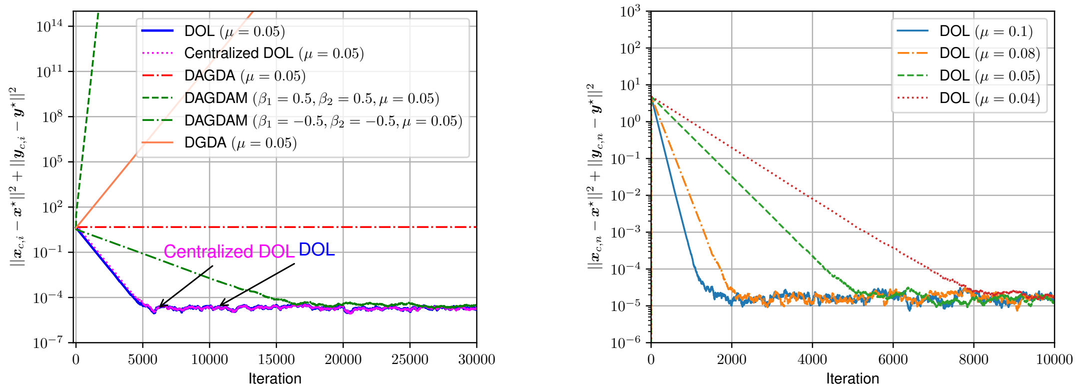

(a) **DOL** versus different methods.

(b) **DOL** under different step sizes.

Fig. 1. Performance comparison of **DOL** with diffusion gradient descent ascent (DGDA), diffusion alternating gradient descent ascent (DAGDA), and diffusion alternating gradient descent ascent with momentum (DAGDAM) for stochastic bilinear game.

second stationary condition (13b) after $T_1 = \mathcal{O}(\varepsilon^{-4})$ iterations. The first stationary condition (13a) will also be satisfied under condition (19). □

4. COMPUTER SIMULATIONS

The experiment utilizes a stochastic bilinear function to illustrate the performance of **DOL**. In this example, the agents collaborate to solve the following global objective:

$$\min_x \max_y J(x, y) = \sum_{k=1}^{K} p_k J_k(x, y)$$
$$J_k(x, y) \triangleq \mathbb{E}_{\boldsymbol{\xi}_k} \left[x^\top A(\boldsymbol{\xi}_k) y + x^\top g_x(\boldsymbol{\xi}_k) + g_y^\top(\boldsymbol{\xi}_k) y \right] \tag{20}$$

where $A(\boldsymbol{\xi}_k)$, $g_x(\boldsymbol{\xi}_k)$, $g_y(\boldsymbol{\xi}_k)$ are random quantities with dependence on the random sample $\boldsymbol{\xi}_k$. We consider the models in [19] to generate the data samples for each agent k:

$$A(\boldsymbol{\xi}_k) \triangleq A_k + N(\boldsymbol{\xi}_k) \in \mathbb{R}^{M \times M}$$
$$[N(\boldsymbol{\xi}_k)]_{ij} \sim \mathcal{N}(m_{A_k}, \sigma_{A_k}^2) \tag{21}$$

where A_k is a constant matrix. Each entry of $N(\boldsymbol{\xi}_k)$ is a Gaussian random variable; all variables are independent of each other:

$$g_x(\boldsymbol{\xi}_k) \in \mathbb{R}^{M \times 1} \sim \mathcal{N}(m_{x_k} \mathbb{1}_M, \sigma_{x_k}^2 I_M)$$
$$g_y(\boldsymbol{\xi}_k) \in \mathbb{R}^{M \times 1} \sim \mathcal{N}(m_{y_k} \mathbb{1}_M, \sigma_{y_k}^2 I_M) \tag{22}$$

The Nash equilibrium of (20) can be verified by setting the gradients of $J(x, y)$ relative to x and y to 0, respectively.

4.1. Experimental Setting

We generate the random samples for each agent k as follows: the mean values are set to zero, i.e., $m_{A_k} = m_{x_k} = m_{y_k} = 0$, the variances are set to $\sigma_{A_k}^2 = \sigma_{x_k}^2 = \sigma_{y_k}^2 = 1\text{e}{-4}$, the constant matrices are set to $A_k = I_M$, and $M = 20$. Under this setting, the Nash equilibrium is given by $(x^o, y^o) = (0_{M\times 1}, 0_{M\times 1})$. We then generate a strongly-connected network with $K = 10$ agents and use the averaging rule [33, 34]. We also implement **DOL** under a fully-connected network topology. We further compare **DOL** with diffusion gradient descent ascent, diffusion alternating gradient descent ascent (DAGDA) (which updates the primal and dual variables in a sequential manner) [17], and the DAGDA with momentum (DAGDAM) [35] by extending the corresponding algorithms to the multi-agent scenario.

4.2. Simulation Result

Figure 1(a) illustrates the mean-square-error between the network centroid $(\boldsymbol{x}_{c,i}, \boldsymbol{y}_{c,i})$ obtained by different algorithms and the Nash equilibrium (x^o, y^o). It is observed from this figure that DGDA diverges, while the DAGDA trajectory remains flat, neither converging nor diverging. On the other hand, DAGDAM with conventional positive momentum ($\beta_1 > 0, \beta_2 > 0$) has unstable effect in this application where it diverges fast, while DAGDAM with negative momentum ($\beta_1 < 0, \beta_2 < 0$) is able to converge, however, to the less accurate neighbourhood of the Nash equilibrium. In comparison, the **DOL** method converges to the neighbourhood of the Nash equilibrium and also matches the centralized solution.

Figure 1(b) illustrates the convergence result of **DOL** for different step-sizes. It is observed that **DOL** converges for a wide range of step sizes.

5. CONCLUSION

In this work, we introduced and studied the convergence behavior of diffusion optimistic learning (**DOL**) for fully distributed stochastic nonconvex-PL minimax optimization problems. Convergence towards an ε-stationary point is guaranteed under suitable parameters. We simulated the algorithm by considering a stochastic bilinear game, and observed that **DOL** outperforms DAGDA and DAGDAM and converges to the neighbourhood of the Nash equilibrium.

6. ACKNOWLEDGEMENT

The authors wish to thank Dr. Flavio Pavan for feedback on an earlier draft of this manuscript.

7. REFERENCES

[1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, 2014, pp. 1–9.

[2] B. Dai, A. Shaw, L. Li, L. Xiao, N. He, Z. Liu, J. Chen, and L. Song, "Sbeed: Convergent reinforcement learning with nonlinear function approximation," in *International Conference on Machine Learning*. PMLR, 2018, pp. 1125–1134.

[3] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," *arXiv:1706.06083*, 2017.

[4] C. Daskalakis, S. Skoulakis, and M. Zampetakis, "The complexity of constrained min-max optimization," in Proc. *ACM SIGACT Symposium on Theory of Computing*, 2021, pp. 1466–1478.

[5] G. Zhang, Y. Wang, L. Lessard, and R. B. Grosse, "Near-optimal local convergence of alternating gradient descent-ascent for minimax optimization," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022, pp. 7659–7679.

[6] D. Kovalev and A. Gasnikov, "The first optimal algorithm for smooth and strongly-convex-strongly-concave minimax optimization," in *Advances in Neural Information Processing Systems*, 2022, pp. 14691–14703.

[7] D. Goktas and A. Greenwald, "Convex-concave min-max stackelberg games," in *Advances in Neural Information Processing Systems*, 2021, pp. 2991–3003.

[8] N. Golowich, S. Pattathil, C. Daskalakis, and A. Ozdaglar, "Last iterate is slower than averaged iterate in smooth convex-concave saddle point problems," in *Conference on Learning Theory*. PMLR, 2020, pp. 1758–1784.

[9] H. He, S. Zhao, Y. Xi, and J. Ho, "GDA-AM: On the effectiveness of solving min-imax optimization via anderson mixing," in *International Conference on Learning Representations*, 2022.

[10] T. Lin, C. Jin, and M. Jordan, "On gradient descent ascent for nonconvex-concave minimax problems," in *International Conference on Machine Learning*. PMLR, 2020, pp. 6083–6093.

[11] L. Luo, H. Ye, Z. Huang, and T. Zhang, "Stochastic recursive gradient descent ascent for stochastic nonconvex-strongly-concave minimax problems," in *Advances in Neural Information Processing Systems*, 2020, pp. 20566–20577.

[12] P. Mahdavinia, Y. Deng, H. Li, and M. Mahdavi, "Tight analysis of extra-gradient and optimistic gradient methods for nonconvex minimax problems," in *Advances in Neural Information Processing Systems*, 2022, pp. 31213–31225.

[13] J. Diakonikolas, C. Daskalakis, and M. I. Jordan, "Efficient methods for structured nonconvex-nonconcave min-max optimization," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 2746–2754.

[14] T. Pethick, O. Fercoq, P. Latafat, P. Patrinos, and V. Cevher, "Solving stochastic weak minty variational inequalities without increasing batch size," *arXiv:2302.09029*, 2023.

[15] A. Bohm, "Solving nonconvex-nonconcave min-max problems exhibiting weak minty solutions," *Transactions on Machine Learning Research*, pp. 1–21, 2023.

[16] J. Yang, A. Orvieto, A. Lucchi, and N. He, "Faster single-loop algorithms for minimax optimization without strong concavity," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022, pp. 5485–5517.

[17] J. Yang, N. Kiyavash, and N. He, "Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems," in *Advances in Neural Information Processing Systems*, 2020, pp. 1153–1165.

[18] E. Gorbunov, A. Taylor, and G. Gidel, "Last-iterate convergence of optimistic gradient method for monotone variational inequalities," in *Advances in Neural Information Processing Systems*, 2022, pp. 21858–21870.

[19] C. J. Li, Y. Yu, N. Loizou, G. Gidel, Y. Ma, N. L. Roux, and M. Jordan, "On the convergence of stochastic extragradient for bilinear games using restarted iteration averaging," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022, pp. 9793–9826.

[20] M. Nouiehed, M. Sanjabi, T. Huang, J. D. Lee, and M. Razaviyayn, "Solving a class of non-convex min-max games using iterative first order methods," in *Advances in Neural Information Processing Systems*, 2019, pp. 1–9.

[21] D. Mateos-Núnez and J. Cortés, "Distributed saddle-point subgradient algorithms with laplacian averaging," *IEEE Transactions on Automatic Control*, , no. 6, pp. 2720–2735, 2016.

[22] M. I. Qureshi and U. A. Khan, "Distributed saddle point problems for strongly concave-convex functions," *arXiv:2202.05812*, 2022.

[23] W. Xian, F. Huang, Y. Zhang, and H. Huang, "A faster decentralized algorithm for nonconvex minimax problems," in *Advances in Neural Information Processing Systems*, 2021, pp. 25865–25877.

[24] H. Gao, "Decentralized stochastic gradient descent ascent for finite-sum minimax problems," *arXiv:2212.02724*, 2022.

[25] L. Chen, H. Ye, and L. Luo, "A simple and efficient stochastic algorithm for decentralized nonconvex-strongly-concave minimax optimization," *arXiv:2212.02387*, 2022.

[26] F. Huang and S. Chen, "Near-optimal decentralized momentum method for nonconvex-PL minimax problems," *arXiv:2304.10902*, 2023.

[27] A. Beznosikov, P. Dvurechenskii, A. Koloskova, V. Samokhin, S. U. Stich, and A. Gasnikov, "Decentralized local stochastic extra-gradient for variational inequalities," in *Advances in Neural Information Processing Systems*, 2022, pp. 38116–38133.

[28] M. Liu, W. Zhang, Y. Mroueh, X. Cui, J. Ross, T. Yang, and P. Das, "A decentralized parallel algorithm for training generative adversarial nets," in *Advances in Neural Information Processing Systems*, 2020, pp. 11056–11070.

[29] C. Liu, L. Zhu, and M. Belkin, "Loss landscapes and optimization in over-parameterized non-linear systems and neural networks," *Applied and Computational Harmonic Analysis*, pp. 85–116, 2022.

[30] L. D. Popov, "A modification of the arrow-hurwicz method for search of saddle points," *Mathematical Notes of the Academy of Sciences of the USSR*, pp. 845–848, 1980.

[31] Y. Malitsky and M. K. Tam, "A forward-backward splitting method for monotone inclusions without cocoercivity," *SIAM Journal on Optimization*, pp. 1451–1472, 2020.

[32] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends® in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, 2014.

[33] A. H. Sayed, *Inference and Learning from Data*, vol. 3, Cambridge University Press, 2022.

[34] A. H. Sayed, "Adaptive networks," *Proceedings of the IEEE*, vol. 102, no. 4, pp. 460–497, 2014.

[35] G. Gidel, R. A. Hemmat, M. Pezeshki, R. L. Priol, G. Huang, S. Lacoste-Julien, and I. Mitliagkas, "Negative momentum for improved game dynamics," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2019, pp. 1802–1811.