DIFFUSION LEARNING IN NON-CONVEX ENVIRONMENTS

Stefan Vlaski^{*†} *and Ali H. Sayed*[†]

*Department of Electrical Engineering, University of California, Los Angeles †École Polytechnique Fédérale de Lausanne

ABSTRACT

Driven by the need to solve increasingly complex optimization problems in signal processing and machine learning, recent years have seen rising interest in the behavior of gradient-descent based algorithms in non-convex environments. Most of the works on distributed non-convex optimization focus on the deterministic setting, where exact gradients are available at each agent. In this work, we consider stochastic cost functions, where exact gradients are replaced by stochastic approximations and the resulting gradient noise persistently seeps into the dynamics of the algorithm. We establish that the diffusion algorithm continues to yield meaningful estimates in these more challenging, non-convex environments, in the sense that (a) despite the distributed implementation, restricted to local interactions, individual agents cluster in a small region around a common and well-defined vector, which will carry the interpretation of a network centroid, and (b) the network centroid inherits many properties of the centralized, stochastic gradient descent recursion, including the return of an $O(\mu)$ -mean-square-stationary point in at most $O(1/\mu^2)$ iterations.

Index Terms— Stochastic optimization, adaptation, non-convex, gradient noise, stationary points.

1. INTRODUCTION

The broad objective of distributed adaptation and learning is the solution of global, stochastic optimization problems over distributed networks through localized interactions and in the absence of information about the statistical properties of the data. When constant, rather than diminishing, step-sizes are employed, the resulting algorithms are adaptive in nature and able to adapt to drifts in the problem or data statistics. In this work, we consider a collection of Nagents, where each agent is equipped with a stochastic risk of the form $J_k(w) \triangleq \mathbb{E}_x Q_k(w; x)$ and the objective of the network is to seek the Pareto solution:

$$\min_{w} J(w), \quad \text{where } J(w) \triangleq \sum_{k=1}^{N} p_k J_k(w) \tag{1}$$

and p_k are positive weights that are normalized to add up to one and will be specified further below. In (1), the variable x represents data and expectations are computed relative to the distribution of this data. Algorithms for the solution of (1) have been studied extensively over recent years both with inexact [1–3] and exact [4–6] gradients. Here, we shall focus on the following Adapt-then-Combine variation of the diffusion algorithm [2]:

$$\boldsymbol{\phi}_{k,i} = \boldsymbol{w}_{k,i-1} - \mu \widehat{\nabla J}_k(\boldsymbol{w}_{k,i-1})$$
(2a)

$$w_{k,i} = \sum_{\ell=1}^{N} a_{\ell k} \phi_{\ell,i} \tag{2b}$$

where $\widehat{\nabla J}_k(\cdot)$ denotes a stochastic approximation to the true local gradient $\nabla J_k(\cdot)$ and $a_{\ell k}$ are convex combination weights satisfying:

$$a_{\ell k} \ge 0, \quad \sum_{\ell \in \mathcal{N}_k} a_{\ell k} = 1, \quad a_{\ell k} = 0 \text{ if } \ell \notin \mathcal{N}_k$$
 (3)

where \mathcal{N}_k denotes the set of neighbors of agent k.

Assumption 1 (Strongly-connected graph). We shall assume that the graph described by the weighted combination matrix $A = [a_{\ell k}]$ is strongly-connected.

It then follows from the Perron-Frobenius theorem [2,7,8], that A has a single eigenvalue at one while all other eigenvalues are inside the unit circle, so that $\rho(A) = 1$. Moreover, if we let p denote the right-eigenvector of A that is associated with the eigenvalue at one, and if we normalize the entries of p to add up to one, then it also holds that all entries of p are strictly positive, i.e.,

$$Ap = p, \quad \mathbb{1}^{\mathsf{T}}p = 1, \quad p_k > 0 \tag{4}$$

where the $\{p_k\}$ denote the individual entries of the Perron vector, p.

The performance of the diffusion algorithm (2a)–(2b) has been studied extensively in differentiable settings [3, 9], with extensions to multi-task [10], constrained [11] and non-differentiable [12, 13] environments. A common assumption in these works, along with others studying the behavior of distributed optimization algorithms, is that of convexity (or strong-convexity) of the aggregate risk J(w). While many problems of interest such as least-squares estimation [2], logistic regression [2], and support vector machines [14] are convex, there has been increased interest in the optimization of non-convex cost functions. Such problems appear frequently in the design of robust estimators [15] and the training of more complex machine learning architectures such as those involving dictionary learning [16] and artificial neural networks [17].

As a result, recent works have pursued the development and study of optimization algorithms for non-convex problems, both in the deterministic and stochastic setting [18,19]. Works on distributed non-convex optimization so far have largely focused on the *deterministic* setting, where exact gradients are available [20–22]. Other related works consider stochastic gradient approximations, but require a central master node [23], resulting in parallel, rather than a truly distributed implementation, or require a diminishing stepsize [24], rendering them inapplicable in scenarios that call for continuous adaptation. In contrast, we shall focus on implementations

This work was supported in part by NSF grant CCF-1524250. Emails:{stefan.vlaski, ali.sayed}@epfl.ch.

that employ stochastic gradient approximations and constant stepsizes, and do not require a central communication hub. This is motivated by two considerations. First, computation of the exact gradients $\nabla J_k(\cdot)$ may be infeasible in practice because (a) data may be streaming in, making it impossible to compute $\nabla \mathbb{E}_x Q_k(\cdot; \boldsymbol{x})$ in the absence of knowledge about the distribution of the data or (b) the data set, while available as a batch, may be so large that efficient computation of the full gradient is infeasible. Second, and perhaps surprisingly, there have been observations in the literature commenting on the benefit of introducing gradient perturbations into the optimization of non-convex functions [18, 19]. These observations are mainly motivated by the fact that gradient perturbations move iterates "out" of unstable saddle-points and towards more stable local minima, which are generally preferable [17]. Some works have pursued rigorous analytical justifications of this behavior for the class of strict saddle-points [18, 19].

We first establish that in non-convex environments, as is the case in the strongly-convex setting (see for example [3]), the evolution of the iterates $\boldsymbol{w}_{k,i}$ continues to be well-described by the evolution of the weighted centroid vector $\sum_{k=1}^{N} p_k \boldsymbol{w}_{k,i}$, in the sense that the iterates cluster around the network centroid after sufficient iterations. This observation allows us to establish an approximate descent relation for the network centroid, which mirrors that for the centralized stochastic gradient descent up to constants (see for example [18]), and use this relation to establish that the network centroid, around which all agents cluster, will reach an $O(\mu)$ -mean-square-stationary point after at most $O(1/\mu^2)$ iterations.

2. EVOLUTION ANALYSIS

We perform the analysis under the following common assumptions on the gradients and their approximations.

Assumption 2 (Lipschitz gradients). For each k, the gradient $\nabla J_k(\cdot)$ is Lipschitz, namely, for any $x, y \in \mathbb{R}^M$:

$$\|\nabla J_k(x) - \nabla J_k(y)\| \le \delta \|x - y\|$$
(5)

In light of (1) and Jensen's inequality, this implies for the aggregate cost:

$$\|\nabla J(x) - \nabla J(y)\| \le \delta \|x - y\| \tag{6}$$

Assumption 3 (Bounded gradients). For each k, the gradient $\nabla J_k(\cdot)$ is bounded, namely, for any $x \in \mathbb{R}^M$:

$$\|\nabla J_k(x)\| \le G \tag{7}$$

Assumption 4 (Gradient noise process). For each k, the gradient noise process is defined as

$$\boldsymbol{s}_{k,i}(\boldsymbol{w}_{k,i-1}) = \widehat{\nabla J}_k(\boldsymbol{w}_{k,i-1}) - \nabla J_k(\boldsymbol{w}_{k,i-1})$$
(8)

and satisfies

$$\mathbb{E}\left[\boldsymbol{s}_{k,i}(\boldsymbol{w}_{k,i-1})|\boldsymbol{w}_{k,i-1}\right] = 0 \tag{9a}$$

$$\mathbb{E}\left[\|\boldsymbol{s}_{k,i}(\boldsymbol{w}_{k,i-1})\|^{2}|\boldsymbol{w}_{k,i-1}\right] \leq \beta^{2} \|\nabla J_{k}(\boldsymbol{w}_{k,i-1})\|^{2} + \sigma^{2} \quad (9b)$$

for some non-negative constants
$$\{\beta^2, \sigma^2\}$$
.

2.1. Centroid recursion

In analyzing the recursions, it turns out to be useful to introduce the following extended quantities, which collect variables from across the network:

$$\boldsymbol{w}_i \triangleq \operatorname{col} \left\{ \boldsymbol{w}_{1,i}, \dots, \boldsymbol{w}_{N,i} \right\}$$
(10)

$$\mathcal{A} \triangleq A \otimes I_M \tag{11}$$

$$\widehat{\boldsymbol{g}}(\boldsymbol{w}_i) \triangleq \operatorname{col}\left\{\widehat{\nabla J_1}(\boldsymbol{w}_{1,i}), \dots, \widehat{\nabla J_N}(\boldsymbol{w}_{N,i})\right\}$$
 (12)

We can then write the diffusion recursion (2a)-(2b) compactly as

$$\boldsymbol{w}_{i} = \boldsymbol{\mathcal{A}}^{\mathsf{T}} \left(\boldsymbol{w}_{i-1} - \boldsymbol{\mu} \widehat{\boldsymbol{g}}(\boldsymbol{w}_{i-1}) \right)$$
(13)

Multiplying both sides of (13) by $(p^{\mathsf{T}} \otimes I)$ from the left, we obtain in light of (4):

$$(p^{\mathsf{T}} \otimes I) \mathbf{w}_{i} = (p^{\mathsf{T}} \otimes I) \mathbf{w}_{i-1} - \mu (p^{\mathsf{T}} \otimes I) \widehat{g}(\mathbf{w}_{i-1})$$
 (14)

Letting $\boldsymbol{w}_{c,i} = \sum_{k=1}^{K} p_k \boldsymbol{w}_{k,i} = (p^{\mathsf{T}} \otimes I) \boldsymbol{w}_i$ and exploiting the block-structure of the gradient term, we find:

$$\boldsymbol{w}_{c,i} = \boldsymbol{w}_{c,i-1} - \mu \sum_{k=1}^{N} p_k \widehat{\nabla J}_k(\boldsymbol{w}_{k,i-1})$$
(15)

Note that $\boldsymbol{w}_{c,i}$ is essentially a convex combination of estimates across the network and can be viewed as a weighted centroid. The recursion for $\boldsymbol{w}_{c,i}$ is reminiscent of a stochastic gradient step associated with the aggregate cost $\sum_{k=1}^{N} p_k J_k(\boldsymbol{w})$ with the exact gradients $\nabla J_k(\cdot)$ replaced by stochastic approximations $\widehat{\nabla J}_k(\cdot)$ and the stochastic gradients evaluated at $\boldsymbol{w}_{k,i-1}$, rather than $\boldsymbol{w}_{c,i-1}$. In fact, we can write:

$$\boldsymbol{w}_{c,i} = \boldsymbol{w}_{c,i-1} - \mu \sum_{k=1}^{N} p_k \nabla J_k(\boldsymbol{w}_{c,i-1}) - \mu \boldsymbol{d}_{i-1} - \mu \, \boldsymbol{s}_i \quad (16)$$

where we defined the perturbation terms:

$$\boldsymbol{d}_{i-1} \triangleq \sum_{k=1}^{N} p_k \left(\nabla J_k(\boldsymbol{w}_{k,i-1}) - \nabla J_k(\boldsymbol{w}_{c,i-1}) \right)$$
(17)

$$\boldsymbol{s}_{i} \triangleq \sum_{k=1}^{N} p_{k} \left(\widehat{\nabla J}_{k}(\boldsymbol{w}_{k,i-1}) - \nabla J_{k}(\boldsymbol{w}_{k,i-1}) \right)$$
(18)

Observe that d_{i-1} arises from the disagreement within the network, and in particular that if each $w_{k,i-1}$ remains close to the network centroid $w_{c,i-1}$, the perturbation will be small in light of the Lipschitz condition (5) on the gradients. The second perturbation term s_i arises from the noise introduced by stochastic gradient approximations at each agent. The remainder of this work is dedicated to establishing that recursion (16) will continue to exhibit some of the desired properties of (centralized) gradient descent, despite the presence of persistent and coupled perturbation terms.

2.2. Network disagreement

To begin with, we study more closely the evolution of the individual estimates $w_{k,i}$ relative to the network centroid $w_{c,i}$. From (16), we have:

$$\boldsymbol{w}_{c,i} \triangleq \mathbb{1} \otimes \boldsymbol{w}_{c,i} = \left(\mathbb{1}p^{\mathsf{T}} \otimes I\right) \left(\boldsymbol{w}_{i-1} - \mu \widehat{\boldsymbol{g}}(\boldsymbol{w}_{i-1})\right)$$
 (19)

and hence:

$$\boldsymbol{w}_{i} - \boldsymbol{w}_{c,i} = \left(\boldsymbol{\mathcal{A}}^{\mathsf{T}} - \mathbf{1}\boldsymbol{p}^{\mathsf{T}} \otimes \boldsymbol{I}\right) \left(\boldsymbol{w}_{i-1} - \boldsymbol{\mu}\boldsymbol{\widehat{g}}(\boldsymbol{w}_{i-1})\right)$$

$$\stackrel{(a)}{=} \left(\boldsymbol{\mathcal{A}}^{\mathsf{T}} - \mathbf{1}\boldsymbol{p}^{\mathsf{T}} \otimes \boldsymbol{I}\right) \left(\boldsymbol{I} - \mathbf{1}\boldsymbol{p}^{\mathsf{T}} \otimes \boldsymbol{I}\right) \left(\boldsymbol{w}_{i-1} - \boldsymbol{\mu}\boldsymbol{\widehat{g}}(\boldsymbol{w}_{i-1})\right)$$

$$= \left(\boldsymbol{\mathcal{A}}^{\mathsf{T}} - \mathbf{1}\boldsymbol{p}^{\mathsf{T}} \otimes \boldsymbol{I}\right) \left(\boldsymbol{w}_{i-1} - \boldsymbol{w}_{c,i-1} - \boldsymbol{\mu}\boldsymbol{\widehat{g}}(\boldsymbol{w}_{i-1})\right)$$
(20)

where (a) follows from the equality:

$$\left(\mathcal{A}^{\mathsf{T}} - \mathbb{1}p^{\mathsf{T}} \otimes I\right) \left(I - \mathbb{1}p^{\mathsf{T}} \otimes I\right) = \mathcal{A}^{\mathsf{T}} - \mathbb{1}p^{\mathsf{T}} \otimes I \qquad (21)$$

The recursive relation for the network disagreement allows for the following result.

Lemma 1 (Network disagreement). Under assumptions 1-4 and for sufficiently small step-sizes μ , the network disagreement is bounded as:

$$\mathbb{E} \|\boldsymbol{w}_{i} - \boldsymbol{w}_{c,i}\|^{2} \leq \mu^{2} \frac{\lambda_{2}^{2}}{(1 - \lambda_{2})^{2}} \left(G^{2} + \beta^{2}G^{2} + \sigma^{2}\right), \text{ for } i \geq i_{o}.$$

$$(22)$$

$$where \lambda_{2} = \rho \left(\mathcal{A}^{\mathsf{T}} - \mathbb{1}p^{\mathsf{T}} \otimes I\right) + \tau < 1.$$

Proof. Omitted due to space limitations. \Box

This result allows us to bound the perturbation terms encountered in (16).

Lemma 2 (Perturbation bounds). Under assumptions 1–4 and for sufficiently small step-sizes μ , the perturbation terms are bounded as:

$$\mathbb{E} \|\boldsymbol{d}_{i-1}\|^{2} \leq \mu^{2} c_{\rho} \delta^{2} p_{\max} \frac{\lambda_{2}^{2}}{(1-\lambda_{2})^{2}} \left(G^{2} + \beta^{2} G^{2} + \sigma^{2} \right)$$
(23)

$$\mathbb{E} \left\| \boldsymbol{s}_{i-1} \right\|^2 \le \beta^2 G^2 + \sigma^2 \tag{24}$$

after sufficient iterations $i \ge i_0$.

Proof. Omitted due to space limitations.

2.3. Evolution of the network centroid

Having established that after sufficient iterations, all agents in the network will have contracted around the centroid in a small cluster for small step-sizes, we can now leverage $w_{c,i}$ as a proxy for all $w_{k,i}$. Assumption 2 implies the following bound:

$$J(\boldsymbol{w}_{c,i}) \leq J(\boldsymbol{w}_{c,i-1}) + \nabla J(\boldsymbol{w}_{c,i-1})^{\mathsf{T}} (\boldsymbol{w}_{c,i} - \boldsymbol{w}_{c,i-1}) \\ + \frac{\delta}{2} \|\boldsymbol{w}_{c,i} - \boldsymbol{w}_{c,i-1}\|^2$$
(25)

This relation, along with (16) and the results from Lemma 2, allow us to establish the following theorem.

Theorem 1 (Descent relation). Under assumptions 1–4 and for sufficiently small step-sizes, we have:

$$\mathbb{E}\left\{J(\boldsymbol{w}_{c,i}) | \boldsymbol{w}_{c,i-1}\right\} \le J(\boldsymbol{w}_{c,i-1}) - \mu c_1 \|\nabla J(\boldsymbol{w}_{c,i-1})\|^2 + \mu^2 c_2$$
(26)

where

$$c_1 \triangleq \frac{1 - 2\mu\delta}{2} = O(1)$$

$$c_2 \triangleq \frac{\delta}{2} \left(\beta^2 G^2 + \sigma^2\right) + O(\mu) = O(1)$$

Proof. Omitted due to space limitations.

The descent relation (26) has the same form (apart from constants) as the one obtained for centralized gradient descent (see for example [18]) and is useful because it establishes a sufficient condition under which the diffusion recursion (2a)–(2b) is a descent recursion for the network centroid. Specifically, we will have $\mathbb{E} \{J(\boldsymbol{w}_{c,i}) | \boldsymbol{w}_{c,i-1}\} \leq J(\boldsymbol{w}_{c,i-1})$ whenever:

$$\|\nabla J(\boldsymbol{w}_{c,i-1})\|^2 \ge \mu \frac{c_2}{c_1}$$
 (27)

As a corollary we obtain:

Corollary 1 (Stationary points). Suppose the aggregate cost is bounded from below, i.e., $J(w) \ge J^{\circ}$. Then, the centroid $w_{c,i}$ will reach an $O(\mu)$ -mean-square-stationary point in at most $O(1/\mu^2)$ iterations. Specifically for some time i^* , we have:

$$\mathbb{E} \left\| \nabla J(\boldsymbol{w}_{c,i^{\star}}) \right\|^2 \le 2\mu \frac{c_2}{c_1}$$
(28)

and

$$i^{\star} \le \left(\frac{J(w_0) - J^o}{c_2}\right) \frac{1}{\mu^2} \tag{29}$$

Proof. We prove the result by contradiction. First, assume that there is no time i^* with $\mathbb{E} \|\nabla J(\boldsymbol{w}_{c,i^*})\|^2 \leq 2\mu \frac{c_2}{c_1}$. Then, for all *i*:

$$\mathbb{E} \left\| \nabla J(\boldsymbol{w}_{c,i}) \right\|^2 \ge 2\mu \frac{c_2}{c_1}$$
(30)

Iterating (26), we obtain:

$$\mathbb{E} J(\boldsymbol{w}_{c,i}) \leq J(w_0) - \mu c_1 \sum_{k=1}^{i} \left(\mathbb{E} \left\| \nabla J(\boldsymbol{w}_{c,k-1}) \right\|^2 - \mu \frac{c_2}{c_1} \right)$$
$$\leq J(w_0) - \mu^2 c_2 i \tag{31}$$

and hence $\lim_{i\to\infty} \mathbb{E} J(\boldsymbol{w}_{c,i}) \leq -\infty$ which contradicts $J(w) \geq J^o$ for all w. Hence, we conclude that there is a finite, first moment in time i^* with $\mathbb{E} \|\nabla J(\boldsymbol{w}_{c,i^*})\|^2 \leq 2\mu \frac{c_2}{c_1}$. Since i^* is the first such time, we still have for $i < i^*$:

$$\mathbb{E} \left\| \nabla J(\boldsymbol{w}_c) \right\|^2 \ge 2\mu \frac{c_2}{c_1} \tag{32}$$

Iterating (26) up to time i^* , we similarly obtain:

$$J^{o} \leq \mathbb{E} J(\boldsymbol{w}_{c,i^{\star}}) \leq J(w_{0}) - \mu^{2} c_{2} i^{\star}$$
(33)

Rearranging yields the result.

We conclude that at some time $i^* \leq O(1/\mu^2)$, the network centroid will reach an $O(\mu)$ -mean-square-stationary point.

3. APPLICATION: ROBUST REGRESSION

Consider a scenario where each agent k in the network observes streaming realizations $\{\gamma(k, i), h_{k,i}\}$ from the linear model:

$$\boldsymbol{\gamma}(k) = \boldsymbol{h}_k^{\mathsf{T}} \boldsymbol{w}^o + \boldsymbol{v}(k) \tag{34}$$

where $\gamma(k)$ denote scalar observations and v(k) denotes measurement noise. The most commonly used approach for solving such a

problem in a distributed setting is via least-mean-square error estimation, resulting in the local cost functions:

$$J_{k}^{\mathrm{LS}}(w) = \mathbb{E} \left\| \boldsymbol{\gamma}(k) - \boldsymbol{h}_{k}^{\mathsf{T}} w \right\|^{2}$$
(35)

The resulting problem is convex and has been studied extensively in the literature. While effective under the assumption of Gaussian noise, and similar well-behaved noise conditions, this approach is susceptible to outliers caused by heavy-tailed distributions for v(k) [15]. This is caused by the fact that the quadratic risk penalizes errors proportionally to their squared norm, and as such has a tendency to over-correct to outliers, even if they are very rare. Several alternative robust cost functions have been suggested in the literature. We consider two in particular in order to illustrate the advantages of allowing for non-convex costs in the context of robust estimation, namely the Huber loss and Tukey's biweight loss [15]. For ease of notation, let $e(w) \triangleq \gamma(k) - h_k^T w$. Then:

$$Q_{k}^{\mathrm{H}}(w; \boldsymbol{x}) = \begin{cases} \frac{1}{2} |\boldsymbol{e}(w)|^{2}, & \text{for } |\boldsymbol{e}(w)| \leq c_{H} \\ c_{H} |\boldsymbol{e}(w)| - \frac{1}{2} c_{H}^{2}, & \text{for } |\boldsymbol{e}(w)| > c_{H}. \end{cases}$$

$$\left(c_{H}^{2} \left(1 - \left(1 - |\boldsymbol{e}(w)|^{2} \right)^{3} \right) - c_{H} + \left(1 - \left(1 - |\boldsymbol{e}(w)|^{2} \right)^{3} \right) - c_{H} + \left(1 - \left(1 - \left(1 - |\boldsymbol{e}(w)|^{2} \right)^{3} \right) - c_{H} + c_{H} \right) \right) \right)$$
(36)

$$Q_k^{\rm B}(w; \boldsymbol{x}) = \begin{cases} \frac{c_B^2}{6} \left(1 - \left(1 - \frac{|\boldsymbol{e}(w)|^2}{c_B^2} \right)^3 \right), & \text{for } |\boldsymbol{e}(w)| \le c_B\\ \frac{c_B^2}{6} & \text{otherwise.} \end{cases}$$
(37)

where c_H, c_B are tuning constants. The Huber cost is merely convex (and not strongly-convex), while the Tukey loss is non-convex. As such previous results on the performance of the diffusion algorithm (2a)–(2b) are not applicable. Both of them do however satisfy assumptions 1–4 imposed in this work, and as such the results hold. In particular, since the Huber loss $J_k^H(w)$ has a unique, local minimum, which also happens to be locally strongly-convex, we can conclude that the algorithm will approach the global minimum. The Tukey loss on the other hand, is non-convex, and as such a more challenging problem.

The setting for the simulation results is shown in Figure 1.



Fig. 1: Graph with N = 20 nodes (left) and regressor power Tr $(R_{h,k})$ at each agent (right).

We first show the performance of each strategy in the nominal scenario, where $v(k) \sim \mathcal{N}(0, \sigma_v^2)$. We observe that the distributed strategies outperform the non-cooperative ones, and that despite differences in the rate of convergence, there is negligible difference in the performance of the least-squares, Huber and Tukey variations.

Next, we illustrate the performance in the presence of outliers, modeled as a bimodal distribution with $v(k) \sim (1 - \epsilon)\mathcal{N}(0, \sigma_v^2) + \epsilon \mathcal{N}(10, \sigma_v^2)$ and $\epsilon = 0.1$. We observe that the performance of the least-squares solution dramatically deteriorates, as is to be expected in the presence of deviations from the nominal model.



Fig. 2: MSD performance in the absence of outliers.



Fig. 3: MSD performance in the presence of outliers.

4. CONCLUSION

We presented a framework that allows for the study of the behavior of a class of distributed algorithms in non-convex environments and the presence of persistent stochastic perturbations. In particular, we showed that the evolution of the network continues to be welldescribed by the evolution of the network centroid, in the sense that all iterates $w_{k,i}$ cluster within $O(\mu^2)$ of the network centroid $w_{c,i}$ in the mean-square sense after sufficient iterations. This insight was leveraged to establish a descent relation for the network centroid and to conclude that $w_{c,i}$, and hence $w_{k,i}$ will necessarily reach a point that is $O(\mu)$ -mean-square-stationary in at most $O(1/\mu^2)$ iterations (Corollary 1). Directions for future work include the more careful study of the behavior around saddle-points, and the exploration of further applications, particularly in dictionary learning and deep learning.

5. REFERENCES

- A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Trans. Automatic Control*, vol. 54, no. 1, pp. 48–61, Jan 2009.
- [2] A. H. Sayed, "Adaptation, learning, and optimization over net-

works," *Foundations and Trends in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, July 2014.

- [3] J. Chen and A. H. Sayed, "On the learning behavior of adaptive networks - Part I: Transient analysis," *IEEE Transactions on Information Theory*, vol. 61, no. 6, pp. 3487–3517, June 2015.
- [4] K. Yuan, Q. Ling, and W. Yin, "On the convergence of decentralized gradient descent," *SIAM Journal on Optimization*, vol. 26, no. 3, pp. 1835–1854, 2016.
- [5] W. Shi, Q. Ling, G. Wu, and W. Yin, "EXTRA: An exact first-order algorithm for decentralized consensus optimization," *SIAM Journal on Optimization*, vol. 25, no. 2, pp. 944– 966, 2015.
- [6] K. Yuan, B. Ying, X. Zhao, and A. H. Sayed, "Exact diffusion for distributed optimization and learning – Part II: Convergence analysis," *IEEE Transactions on Signal Processing*, vol. 67, no. 3, pp. 724–739, Feb 2019.
- [7] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, 2003.
- [8] S. U. Pillai, T. Suel, and S. Cha, "The Perron-Frobenius theorem: Some of its applications," *IEEE Signal Processing Magazine*, vol. 22, no. 2, pp. 62–75, March 2005.
- [9] A. H. Sayed, "Adaptive networks," *Proceedings of the IEEE*, vol. 102, no. 4, pp. 460–497, April 2014.
- [10] R. Nassif, C. Richard, A. Ferrari, and A. H. Sayed, "Proximal multitask learning over networks with sparsity-inducing coregularization," *IEEE Transactions on Signal Processing*, vol. 64, no. 23, pp. 6329–6344, Dec 2016.
- [11] Z. J. Towfic and A. H. Sayed, "Adaptive penalty-based distributed stochastic convex optimization," *IEEE Trans. on Signal Process.*, vol. 62, no. 15, pp. 3924–3938, Aug. 2014.
- [12] S. Vlaski, L. Vandenberghe, and A. H. Sayed, "Diffusion stochastic optimization with non-smooth regularizers," in *Proc. of IEEE ICASSP*, Shanghai, China, March 2016, pp. 4149–4153.
- [13] B. Ying and A. H. Sayed, "Performance limits of stochastic sub-gradient learning, Part II: Multi-agent case," *Signal Processing*, vol. 144, pp. 253 – 264, 2018.
- [14] M. A. Hearst, S. T. Dumais, E. Osuna, J. Platt, and B. Scholkopf, "Support vector machines," *IEEE Intelligent Systems and their Applications*, vol. 13, no. 4, pp. 18–28, July 1998.
- [15] A. M. Zoubir, V. Koivunen, E. Ollila, and M. Muma, *Robust Statistics for Signal Processing*, Cambridge University Press, 2018.
- [16] I. Tosic and P. Frossard, "Dictionary learning," *IEEE Signal Processing Magazine*, vol. 28, no. 2, pp. 27–38, March 2011.
- [17] A. Choromanska, M. Henaff, M. Mathieu, G. B. Arous, and Y. LeCun, "The Loss Surfaces of Multilayer Networks," in *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, San Diego, May 2015, pp. 192–204.
- [18] R. Ge, F. Huang, C. Jin, and Y. Yuan, "Escaping from saddle points—online stochastic gradient for tensor decomposition," in *Proc. of Conference on Learning Theory*, Paris, France, 2015, pp. 797–842.

- [19] C. Jin, R. Ge, P. Netrapalli, S. M. Kakade, and M. I. Jordan, "How to escape saddle points efficiently," in *Proc. of ICML*, Sidney, Australia, Aug. 2017.
- [20] P. Di Lorenzo and G. Scutari, "NEXT: in-network nonconvex optimization," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 2, no. 2, pp. 120–136, June 2016.
- [21] Y. Wang, W. Yin, and J. Zeng, "Global convergence of admm in nonconvex nonsmooth optimization," *Journal of Scientific Computing*, vol. 78, no. 1, pp. 29–63, Jan. 2019.
- [22] Davood Hajinezhad, Mingyi Hong, Tuo Zhao, and Zhaoran Wang, "NESTT: A nonconvex primal-dual splitting method for distributed and stochastic optimization," in *Proceedings of the 30th International Conference on Neural Information Processing Systems*, Barcelona, Spain, Dec 2016, pp. 3215–3223.
- [23] Xiangru Lian, Yijun Huang, Yuncheng Li, and Ji Liu, "Asynchronous parallel stochastic gradient for nonconvex optimization," in *Proceedings of the 28th International Conference* on Neural Information Processing Systems, Montreal, Canada, Dec 2015, pp. 2737–2745.
- [24] T. Tatarenko and B. Touri, "Non-convex distributed optimization," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3744–3757, Aug. 2017.