

On Robustness of the Normalized Subgradient Method with Randomly Corrupted Subgradients

Berkay Turan

César A. Uribe

Hoi-To Wai

Mahnoosh Alizadeh

Abstract—Numerous modern optimization and machine learning algorithms rely on subgradient information being trustworthy and hence, they may fail to converge when such information is corrupted. In this paper, we consider the setting where subgradient information may be arbitrarily corrupted (with a given probability) and study the robustness properties of the normalized subgradient method. Under the probabilistic corruption scenario, we prove that the normalized subgradient method, whose updates rely solely on directional information of the subgradient, converges to a minimizer for convex, strongly convex, and weakly-pseudo convex functions satisfying certain conditions. Numerical evidence on linear regression and logistic classification problems supports our results.

I. INTRODUCTION

Gradient-based methods are the most commonly used optimization algorithms in many modern applications including machine learning, control, and signal processing. Owing to their convenience for theoretical analysis and implementation, extensive research focusing on accelerating gradient descent has emerged in the literature. Existing algorithms aim to accelerate gradient descent by choosing adaptive stepsizes [1]–[7]. In addition to their practical success, these methods are supported by strong theoretical guarantees such as fast convergence rates. Furthermore, distributed implementations of optimization algorithms, where the data is stored in worker machines, are gaining traction to preserve privacy. In these schemes, the gradients are computed at the worker machines and communicated to a master machine [8].

However, the above algorithms rely on the assumption that the gradients are *trustworthy*. In reality, the gradients might be erroneous due to computational errors at the machines or data corruption caused by an adversary. Moreover, distributed implementations of these algorithms require reliable communication between the machines. Such schemes are susceptible to man-in-the-middle attacks, where an adversary can take over network sub-systems and arbitrarily alter the information communicated between the machines to prevent convergence to the optimal solution, i.e., a Byzantine attack [9]. In these situations, gradient-based methods that use the exact gradient information may fail to converge, as the erroneous gradient can have an arbitrarily large effect on the algorithm.

This paper studies the robustness property of the normalized subgradient method (NSM) [10], which performs

a subgradient update with an adaptive step size scaled as the reciprocal of the norm of subgradient. We consider a probabilistic corruption scenario where the subgradients received are arbitrarily corrupted at random. As the NSM only uses the directional information of the subgradient, the effects of corrupted subgradients will be limited. We show that when the subgradient corruption probability is below a certain threshold, the NSM converges to an optimal solution for convex and weakly pseudo-convex problems. In particular, let t be the iteration number and let the stepsize decay at $\mathcal{O}(1/t)$, then the NSM finds an optimal solution at the rate of $\mathcal{O}(\log(t)/t)$.

Related Work: There is an extensive literature on developing and analyzing the convergence of optimization algorithms, including those that consider adversarial manipulation. Limited by space, we review categories of prior art that we believe are most related to this work.

1) *Normalized (sub)gradient method:* NSM is a well-studied algorithm for optimization and is supported by theoretical convergence guarantees [10], [11]. When restricted to optimizing differentiable functions, the NSM has been shown to evade saddle points quickly in non-convex optimization problems [12], [13]. Moreover, [14] studies a stochastic version of NSM and shows convergence to a global minimum for quasi-convex functions. A recent work [15] studies an online learning problem and establishes regret bounds for normalized gradient methods for weakly pseudo-convex functions. Scholars have also used normalization of the gradient in first-order optimization algorithms to accelerate the convergence [16], [17].

2) *Learning under adversarial corruption:* The task of learning under adversarial corruption has been a popular research topic. The authors of [18] provide a comprehensive survey on effective learning techniques against an adversarial opponent and give a taxonomy for classifying attacks. Commonly studied tasks for adversarial machine learning have been classification [19]–[22] and linear regression [23]–[26], where the adversarial corruption is usually due to data manipulation. Closest to our setup would be [27], [28], which consider a probabilistic corruption scenario called p -tampering. However, in these works the adversary is restricted to choosing valid tampered data with correct labels instead of the arbitrary corruption we study in this paper.

3) *Robust distributed optimization:* This line of work [29]–[37] focuses on developing robust algorithms for distributed optimization/federated learning by means of robust aggregation techniques. In these studies, there is assumed to be a

B. Turan and M. Alizadeh are with Dept. of ECE, UCSB, Santa Barbara, CA, USA. C. A. Uribe is with Dept. of ECE, Rice University, TX, USA. H. T. Wai is with Dept. of SEEM, CUHK, Shatin, Hong Kong. This work is supported by UCOP Grant LFR-18-548175, NSF grant #1847096, and CUHK Direct Grant #4055113. E-mails: bturan@ucsb.edu, cauribe@rice.edu, htwai@se.cuhk.edu.hk, alizadeh@ucsb.edu

constant α fraction of Byzantine worker machines that are being controlled by an adversary, who alters the messages transmitted from workers to the master machine in order to prevent convergence of distributed implementation of optimization algorithms. The popular solution idea is to filter out the adversarial messages by robust mean estimation of the messages received from all the worker machines.

There is no existing work on robust optimization in a probabilistically and arbitrarily corrupted subgradient setting to the best of our knowledge. In this paper, we study the robustness of the NSM in this context.

Paper Organization: The remainder of the paper is organized as follows. In Section II, we formalize the problem setup and demonstrate how faulty subgradients can affect the convergence. In Section III, we describe the normalized subgradient method (Algorithm 1) and analyze its convergence in a randomly corrupted subgradient setting for convex, strongly convex, and weakly pseudo-convex functions. In Section IV, we provide numerical experiments demonstrating the robustness of the normalized subgradient method.

Notations. Unless otherwise specified, $\|\cdot\|$ denotes the standard Euclidean norm. Given a vector x_t , $x_{t,i}$ indicates the i 'th entry of x_t . The abbreviation *a.s.* indicates almost sure convergence.

II. PROBLEM SETUP

In this section, we formally set up our problem and introduce key concepts and definitions that will be used in the paper. Let $f(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}$ be a cost function of a parameter vector $x \in \mathcal{X} \subseteq \mathbb{R}^d$, where \mathcal{X} is the parameter set that is assumed to be convex and compact with diameter R , i.e., $\|x_1 - x_2\| \leq R, \forall x_1, x_2 \in \mathcal{X}$. Our goal is to find the parameter that minimizes the cost function:

$$x^* = \arg \min_{x \in \mathcal{X}} f(x). \quad (1)$$

(Sub)gradient-based methods are effective for iteratively solving (1). At each iteration t , the iterate x_t is updated upon observing the feedback $g(x_t)$, where $g(x_t) \in \partial f(x_t)$ such that $\partial f(x_t)$ is the subdifferential of f at x_t ; note that if f is differentiable, then $\partial f(x_t) = \{\nabla f(x_t)\}$ (We use the notation $g(x_t) = \nabla f(x_t)$ even if f is differentiable but non-convex). For simplicity of exposition, we will use subdifferential and gradient interchangeably depending on whether the function is differentiable or not.

Gradient-based methods however rely on the feedback received being a trustworthy subgradient information. In this paper, we consider the case where at each iteration t , the feedback received is corrupt with probability p , potentially due to an adversarial attack. Therefore at each iteration t , the feedback is determined as:

$$h_t = \begin{cases} g(x_t) & \text{with probability } 1 - p, \\ b_t & \text{with probability } p. \end{cases} \quad (2)$$

We let the corrupt feedback b_t to be chosen potentially by an adversary, who is assumed to have full knowledge of the problem. The adversary's goal is to prevent the iterates x_t from converging to the optimal solution. We note that this

Algorithm 1 Normalized Subgradient Method (NSM)

Input: Initialize $x_1 \in \mathcal{X}$, step size γ_t , and T

for $t = 1$ to T **do**

Given x_t , receive feedback h_t according to (2).

if $\|h_t\| > 0$ **then**

$x_{t+1} = \Pi_{\mathcal{X}} \left(x_t - \gamma_t \frac{h_t}{\|h_t\|} \right)$

else

$x_{t+1} = x_t$

end if

end for

model encompasses all the cases where the feedback can become corrupt (e.g., communication errors, computational errors, noise) since we set no restrictions on b_t . It is important to mention that because b_t is assumed to be arbitrary, existing subgradient methods under stochastic errors, e.g., [38]–[40], are not applicable in this setting. This literature models the error on the gradients as additive noise with bounded mean and the variance, whereas the adversarial manipulation we consider in this work can be unbounded and arbitrary.

A toy example: In order to make the problem setup clearer and to demonstrate how corrupt subgradients can affect the convergence, let us study a simple example. Let $x \in \mathcal{X} \subset \mathbb{R}^d$, where $\mathcal{X} = \{x | x_0 = x_1 = \dots = x_{d-1}; |x_i| \leq R\}$, and $f(x) = x_{d-1}^4$. The optimal solution is $x^* = [0, \dots, 0]^T$. The function f is differentiable, the gradient at iteration t is:

$$g(x_t) = [0, 0, \dots, 0, 4x_{t,0}^3]^T. \quad (3)$$

Suppose that $x_{t,0} > 0$ and we are applying the regular gradient descent algorithm. Given stepsize γ_t , if $h_t = g(x_t)$, the next iterate after projection is $x_{t+1} = [x_{t,0} - 4x_{t,0}^3\gamma_t/d, \dots, x_{t,0} - 4x_{t,0}^3\gamma_t/d]^T$. Suppose that at each iteration t , the adversary has full knowledge and sets the corrupt gradient as

$$b_t = [-c, -c, -c, \dots, -c, -c]^T, \quad (4)$$

and therefore if $h_t = b_t$, the next iterate is $x_{t+1} = \min\{[R, \dots, R]^T, [x_{t,0} + \gamma_t c, \dots, x_{t,0} + \gamma_t c]^T\}$. We can observe that if the adversary sets $c > (1/p - 1)(4R^3/d)$, then $x_{t+1,i} \geq x_{t,i}$ in expectation and the adversary will prevent convergence (since $x_t \leq R$).

Let us study how the normalized subgradient method outlined in Algorithm 1 would potentially solve this problem. At iteration t , the normalized gradient is:

$$\frac{g(x_t)}{\|g(x_t)\|} = [0, 0, \dots, 0, 1]^T. \quad (5)$$

The normalized corrupt gradient is:

$$\frac{b_t}{\|b_t\|} = \frac{[-1, -1, -1, \dots, -1, -1]^T}{\sqrt{d}} \quad (6)$$

If $h_t = g(x_t)$, $x_{t+1} = [x_{t,0} - \gamma_t/d, \dots, x_{t,0} - \gamma_t/d]^T$. If $h_t = b_t$, $x_{t+1} = [x_{t,0} + \gamma_t/\sqrt{d}, \dots, x_{t,0} + \gamma_t/\sqrt{d}]^T$ assuming that $x_{t,0} + \gamma_t/\sqrt{d} \leq R$. Therefore, in expectation:

$$\mathbb{E}[x_{t+1}|x_t] = [x_{t,0} - (1-p)\gamma_t/d + p\gamma_t/\sqrt{d}, \dots]^T. \quad (7)$$

Hence, if $p < 1/(1 + \sqrt{d})$, the iterates will converge; if $p > 1/(1 + \sqrt{d})$ they will diverge in expectation. This is not the case for the regular gradient descent, where the adversary can set the magnitude of the corrupt gradient large enough to prevent convergence even for low p . Furthermore, the threshold probability p , below which the normalized subgradient method succeeds, is an inherent feature of the cost function that we will discuss later in Section IV-C.

In the next section, we will analyze the convergence of the NSM in the previously described randomly corrupted subgradient setting for (a) convex, (b) strongly convex, and (c) weakly pseudo-convex functions. We conclude this section defining the class of functions studied in this paper.

Definition 1. A differentiable function $f(\cdot)$ is said to be β -smooth if there exists $\beta > 0$ such that

$$\|g(x_1) - g(x_2)\| \leq \beta \|x_1 - x_2\| \quad (8)$$

holds for all $x_1, x_2 \in \mathcal{X}$.

Definition 2. A differentiable function $f(\cdot)$ is said to be μ -strongly convex if there exists $\mu > 0$ such that

$$\langle g(x_1) - g(x_2), x_1 - x_2 \rangle \geq \mu \|x_1 - x_2\|^2 \quad (9)$$

holds for all $x_1, x_2 \in \mathcal{X}$.

Definition 3. A function $f(\cdot)$ is said to satisfy the acute angle condition if there exists some ϕ satisfying $0 \leq \phi < \frac{\pi}{2}$ such that

$$\langle g(x), x - x^*(x) \rangle \geq \cos \phi \|g(x)\| \|x - x^*(x)\| \quad (10)$$

holds for all $x \in \mathcal{X}$ and $g(x) \in \partial f(x)$, where $x^*(x)$ is the point in the set of minima of f that is nearest to x .

As an example, the cost function studied in the simple example $f(x) = x_{d-1}^4$ satisfies the acute angle condition with $\cos \phi = 1/\sqrt{d}$. Another example is $f(x) = |x_1| + |x_2|$, which satisfies the acute angle condition with $\cos \phi = 1/\sqrt{2}$.

Definition 4. A differentiable function $f(\cdot)$ is said to be weakly pseudo-convex if there exists $K > 0$ such that

$$f(x) - f(x^*) \leq K \frac{\langle g(x), x - x^*(x) \rangle}{\|g(x)\|} \quad (11)$$

holds for all $x \in \mathcal{X}$, with the convention that $g(x)/\|g(x)\| = 0$ if $g(x) = 0$, where $x^*(x)$ is the point in the set of minima of f that is nearest to x .

Some examples of weakly pseudo-convex functions are:

- Differentiable, Lipschitz continuous, and pseudo-convex functions [10];
- Star-convex [41] and smooth functions (e.g., $f(x) = |x|(1 - e^{-|x|})$);
- Functions with bounded gradient that satisfy the acute angle condition;
- Functions with bounded gradient that satisfy the α -homogeneity with respect to its minimum; ([15], Proposition 3), i.e., there exists $\alpha > 0$ such that

$$f(t(x - x^*) + x^*) - f(x^*) = t^\alpha (f(x) - f(x^*)) \quad (12)$$

holds for all $x \in \mathcal{X}$ and $t \geq 0$. For instance, $f(x) = (x_1^2 + x_2^2)^2 + 10(x_1^2 - x_2^2)^2$ satisfies this condition and is not convex.

III. ROBUSTNESS ANALYSIS

In this section, we study the normalized subgradient method and show that it can be used to solve the problem defined in Section II. The idea behind is that by normalization, the feedback is restricted to contain only a directional information. This allows us to limit the adversary's attack potential by not allowing arbitrarily large updates, which would have been possible without normalization.

We summarize the normalized subgradient method in Algorithm 1. At each iteration t , the algorithm uses the feedback h_t to compute the normalized vector $h_t/\|h_t\|$ as the update direction. Similar to a standard gradient descent method, it moves the iterate along that direction with stepsize γ_t and then projects the point back to the decision set \mathcal{X} .

In order to prove the convergence of Algorithm 1 in a randomly corrupted subgradient setting, we will use the Robbins-Siegmund Theorem [42] as an auxiliary result:

Theorem 1 (Robbins-Siegmund). Let $(V_t)_{t \geq 1}$, $(\alpha_t)_{t \geq 1}$, $(\chi_t)_{t \geq 1}$, $(\eta_t)_{t \geq 1}$ be four nonnegative $(\mathcal{F})_{t \geq 1}$ -adapted processes such that $\sum_t \alpha_t < \infty$ and $\sum_t \chi_t < \infty$ almost surely. If for each $t \in \mathbb{N}$,

$$\mathbb{E}[V_{t+1} | \mathcal{F}_t] \leq V_t(1 + \alpha_t) + \chi_t - \eta_t \quad (13)$$

then $(V_t)_{t \geq 1}$ converges almost surely to a random variable V_∞ and $\sum_t \eta_t$ is finite almost surely.

We can now state the main result on the convergence of the normalized subgradient method for convex (but possibly non-smooth) functions that meet the acute angle condition (i.e., functions that meet Definition 3) in a randomly corrupted subgradient setting.

Theorem 2. Let $f(\cdot)$ be a convex function defined on \mathcal{X} satisfying the acute angle condition for some $\phi \in [0, \pi/2)$. Suppose that the corruption probability p satisfies

$$p < \frac{\cos \phi}{1 + \cos \phi}.$$

Let $\gamma = \frac{R}{2((1-q)\cos \phi - q)}$ for some $q \in [p, \frac{\cos \phi}{1 + \cos \phi})$ and set the step size as $\gamma_t = \gamma/t$. Then, the iterates generated by Algorithm 1 have the following properties:

$$\mathbb{E}[\|x_{T+1} - x^*(x_{T+1})\|^2] \leq \frac{\gamma^2(1 + \log T)}{T} \quad (14)$$

and

$$\lim_{t \rightarrow \infty} \|x_t - x^*(x_t)\|^2 = 0, \quad a.s. \quad (15)$$

Proof: Let Y_t be a Bernoulli(p) random variable, which

indicates whether the subgradient is trustworthy or corrupt:

$$\|x_{t+1} - x^*(x_{t+1})\|^2 \stackrel{(a)}{\leq} \|x_{t+1} - x^*(x_t)\|^2 \quad (16)$$

$$= \|\Pi_{\mathcal{X}}(x_t - \gamma_t \frac{h_t}{\|h_t\|}) - x^*(x_t)\|^2 \quad (17)$$

$$\stackrel{(b)}{\leq} \|x_t - x^*(x_t) - \gamma_t \frac{h_t}{\|h_t\|}\|^2 \quad (18)$$

$$= \|x_t - x^*(x_t)\|^2 + \gamma_t^2 - 2\gamma_t \langle x_t - x^*(x_t), \frac{h_t}{\|h_t\|} \rangle \quad (19)$$

$$= \|x_t - x^*(x_t)\|^2 + \gamma_t^2 - 2\gamma_t Y_t \langle x_t - x^*(x_t), \frac{b_t}{\|b_t\|} \rangle \quad (20)$$

$$- 2\gamma_t(1 - Y_t) \langle x_t - x^*(x_t), \frac{g(x_t)}{\|g(x_t)\|} \rangle$$

$$\stackrel{(c)}{\leq} \|x_t - x^*(x_t)\|^2 + \gamma_t^2 + 2\gamma_t Y_t \|x_t - x^*(x_t)\| \quad (21)$$

$$- 2\gamma_t(1 - Y_t) \langle x_t - x^*(x_t), \frac{g(x_t)}{\|g(x_t)\|} \rangle$$

$$\stackrel{(d)}{\leq} \|x_t - x^*(x_t)\|^2 + \gamma_t^2 + 2\gamma_t Y_t \|x_t - x^*(x_t)\| \quad (22)$$

$$- \frac{2\gamma_t(1 - Y_t)}{\|g(x_t)\|} \cos \phi \|x_t - x^*(x_t)\| \|g(x_t)\|$$

$$= \|x_t - x^*(x_t)\|^2 + \gamma_t^2 \quad (23)$$

$$- 2\gamma_t \|x_t - x^*(x_t)\| (\cos \phi(1 - Y_t) - Y_t),$$

where (a) follows by definition of $x^*(x)$, (b) is due to nonexpansiveness Euclidean projection, (c) is due to Cauchy–Schwarz inequality, and (d) uses (10). Taking expectation of both sides and noting that x_t and Y_t are independent:

$$\mathbb{E} [\|x_{t+1} - x^*(x_{t+1})\|^2] \leq \mathbb{E} [x_t - x^*(x_t)]^2 + \gamma_t^2 \quad (24)$$

$$- 2\gamma_t (\cos \phi(1 - p) - p) \mathbb{E} [\|x_t - x^*(x_t)\|].$$

Since $p < \frac{\cos \phi}{1 + \cos \phi}$; $(\cos \phi(1 - p) - p) > 0$. Therefore, to upper bound the inequality, we lower bound $\mathbb{E} [\|x_t - x^*(x_t)\|]$ as:

$$\mathbb{E} [\|x_t - x^*(x_t)\|] = \mathbb{E} \left[\frac{\|x_t - x^*(x_t)\|^2}{\|x_t - x^*(x_t)\|} \right] \quad (25)$$

$$\geq \frac{\mathbb{E} [\|x_t - x^*(x_t)\|^2]}{R}$$

Therefore (24) becomes

$$\mathbb{E} [\|x_{t+1} - x^*(x_{t+1})\|^2] \leq \gamma_t^2 \quad (26)$$

$$+ \mathbb{E} [x_t - x^*(x_t)]^2 \left(1 - \gamma_t \frac{2(\cos \phi(1 - p) - p)}{R} \right).$$

Plugging γ_t , we obtain:

$$\mathbb{E} [\|x_{t+1} - x^*(x_{t+1})\|^2] \leq \frac{\gamma^2}{t^2} \quad (27)$$

$$+ \mathbb{E} [x_t - x^*(x_t)]^2 \left(1 - \frac{(\cos \phi(1 - p) - p) 1}{(\cos \phi(1 - q) - q) t} \right)$$

$$\leq \frac{\gamma^2}{t^2} + \mathbb{E} [x_t - x^*(x_t)]^2 \left(1 - \frac{1}{t} \right), \quad (28)$$

where the last inequality uses $\frac{\cos \phi(1-p)-p}{\cos \phi(1-q)-q} \geq 1$ for $q \in [p, \frac{\cos \phi}{1+\cos \phi})$. Finally, a telescopic sum gives:

$$\mathbb{E} [\|x_{T+1} - x^*(x_{T+1})\|^2] \leq \gamma^2 \sum_{t=1}^T \frac{1}{t^2} \prod_{i=t+1}^T \left(1 - \frac{1}{i} \right) \quad (29)$$

$$+ \|x_1 - x^*(x_1)\|^2 \prod_{t=1}^T \left(1 - \frac{1}{t} \right)$$

$$= \gamma^2 \sum_{t=1}^T \frac{1}{t^2} \frac{t}{T} = \frac{\gamma^2}{T} \sum_{t=1}^T \frac{1}{t} \leq \frac{\gamma^2(1 + \log T)}{T}, \quad (30)$$

which completes the first part of the proof. To prove the almost sure convergence of $\|x_t - x^*(x_t)\|^2$, we go back to (23) and take the expectation conditioned on x_t to obtain:

$$\mathbb{E} [\|x_{t+1} - x^*(x_{t+1})\|^2 | x_t] \leq \|x_t - x^*(x_t)\|^2 + \gamma_t^2 \quad (31)$$

$$- 2\gamma_t \frac{\cos \phi(1 - p) - p}{R} \|x_t - x^*(x_t)\|^2$$

$$= \|x_t - x^*(x_t)\|^2 \left(1 - 2\gamma_t \frac{\cos \phi(1 - p) - p}{R} \right) + \gamma_t^2 \quad (32)$$

$$\leq \|x_t - x^*(x_t)\|^2 \left(1 - \frac{1}{t} \right) + \frac{\gamma^2}{t^2} \quad (33)$$

We now apply Theorem 1 with $V_t = \|x_t - x^*(x_t)\|^2$, $\alpha_t = 0$, $\eta_t = \|x_t - x^*(x_t)\|^2/t$, and $\chi_t = \gamma^2/t^2$ to conclude that $\|x_t - x^*(x_t)\|^2$ converges almost surely and $\sum_t \|x_t - x^*(x_t)\|^2/t$ is finite almost surely. In order to determine where $\|x_t - x^*(x_t)\|^2$ converges as well, we use (30) to obtain:

$$\lim_{t \rightarrow \infty} \mathbb{E} [\|x_t - x^*(x_t)\|^2] = 0. \quad (34)$$

Finally, since $\|x_t - x^*(x_t)\|^2 \geq 0$, $\lim_{t \rightarrow \infty} \|x_t - x^*(x_t)\|^2 = 0$ almost surely and therefore x_t converges to a minimum point x^* almost surely. ■

Theorem 2 is a general result for convex functions satisfying the acute angle condition (i.e., convex functions that meet Definition 3). The next corollary states a similar result for strongly convex and smooth functions (i.e., functions that meet Definitions 1 and 2).

Corollary 1. *Suppose that $f(\cdot)$ is μ -strongly convex and β -smooth. Assume that the parameter set \mathcal{X} contains a non-empty ball centered at x^* . Define condition number of $f(\cdot)$ as $\kappa := \beta/\mu$. If the corruption probability satisfies $p < \frac{1}{1+\kappa}$, then the iterates generated by Algorithm 1 with $\gamma = \frac{\kappa R}{2(1-q)-q\kappa}$, $\gamma_t = \gamma/t$ for some $q \in [p, \frac{1}{1+\kappa})$ have the following properties:*

$$\mathbb{E} [\|x_{T+1} - x^*\|^2] \leq \frac{\gamma^2(1 + \log T)}{T} \quad (35)$$

and

$$\lim_{t \rightarrow \infty} x_t = x^*, \quad a.s. \quad (36)$$

where x^* is the unique minimizer of $f(\cdot)$.

Proof: If the parameter set contains a non-empty ball centered at x^* , then $g(x^*) = 0$. Using strong convexity and

smoothness properties, we write:

$$\begin{aligned} \langle g(x), x - x^* \rangle &\geq \mu \|x - x^*\|^2 = \frac{\mu \|x - x^*\|^2 \|g(x)\|}{\|g(x)\|} \\ &\geq \frac{\mu \|x - x^*\|^2 \|g(x)\|}{\beta \|x - x^*\|} = \frac{1}{\kappa} \|x - x^*\| \|g(x)\|. \end{aligned} \quad (37)$$

By setting $\cos \phi = \frac{1}{\kappa}$, we get the required condition in (10) and hence the results from Theorem 2 follow. ■

The class of functions for which Algorithm 1 provides convergence in a randomly corrupted subgradient setting go beyond convex and strongly convex functions. The next corollary is for weakly pseudo-convex functions satisfying the acute angle condition (i.e., functions that meet Definitions 3 and 4).

Corollary 2. *Suppose that $f(\cdot)$ is weakly pseudo-convex and satisfies the acute angle condition for some $\phi \in [0, \pi/2)$. If the corruption probability p satisfies $p < \frac{\cos \phi}{1 + \cos \phi}$, then the iterates generated by Algorithm 1 with $\gamma = \frac{R}{2((1-q)\cos \phi - q)}$, $\gamma_t = \gamma/t$ for some $q \in [p, \frac{\cos \phi}{1 + \cos \phi})$ satisfy:*

$$\mathbb{E} [\|x_{T+1} - x^*(x_{T+1})\|^2] \leq \frac{\gamma^2(1 + \log T)}{T} \quad (38)$$

and

$$\lim_{t \rightarrow \infty} \|x_t - x^*(x_t)\|^2 = 0, \quad a.s. \quad (39)$$

Proof: By Definition 4, if $f(\cdot)$ is weakly pseudo-convex, then it does not have local extrema nor saddle points. Therefore, $g(x) = 0$ only if $x \in \arg \min_{x \in \mathcal{X}} f(x)$. The proof is then identical to that of Theorem 2. ■

Remark 1. *The class of non-convex smooth functions satisfying the acute angle condition also meet Corollary 2. If f is smooth and the diameter of \mathcal{X} is bounded, then $\|g(\cdot)\|$ is bounded. By Proposition 2 in [15], if f has bounded gradient and satisfies the acute angle condition, then f is weakly pseudo-convex. Therefore, Corollary 2 holds for f .*

We would like to highlight the implication of the sufficient condition on the corruption probability, which sets an upper bound on p for Theorem 2 and Corollaries 1 & 2. Observe that the upper bound is a decreasing function of the angle ϕ of the acute angle condition (or κ , the condition number for strongly convex functions). The acute angle condition determines an upper bound on the angle between $g(x)$ and $x - x^*(x)$, $\forall x \in \mathcal{X}$. Therefore, the smaller the angle between $g(x)$ and $x - x^*(x)$, the more accurate the update direction, since the goal is to find $x^*(x)$. This reflects to the upper bound on p : The functions that meet the acute angle condition with small ϕ (or small condition number κ , i.e., well-conditioned problems) can tolerate a higher corruption probability and therefore NSM is more robust to adversarial corruption. The intuition behind is that because the worst adversary may know the optimal solution, they can set the feedback along the direction of $x^*(x) - x$ (which meets the Cauchy-Schwarz inequality used for Equation (20) to (21)). Therefore, the more aligned the direction of the trustworthy subgradient with $x - x^*(x)$, the more effective




Fig. 1. Performances of first-order algorithms for Linear Regression using mean squared error as metric.

the trustworthy updates and the NSM can tolerate higher corruption probabilities.

In the next section, we present the experiments on the robustness of the normalized subgradient method.

IV. NUMERICAL EXPERIMENTS

In this section, we demonstrate the robustness of the normalized subgradient method in an adversarial setting by applying it to linear regression and logistic classification problems. Furthermore, we include performances of several other first-order algorithms for comparison. Lastly, we demonstrate the effect of the threshold corruption probability using the toy example discussed in Section II.

A. Linear Regression

Consider the following problem:

$$\min_{x \in \mathcal{X}} \|y - Ax\|^2, \quad (40)$$

where $A \in \mathbb{R}^{N \times d}$ is a matrix containing the N data vectors in its rows and $y \in \mathbb{R}^N$ is the vector containing the N associated label values or outputs. The task is to minimize the empirical quadratic error between a function $y_i = g(A_i)$ of d variables and its linear regression $\hat{g}(A_i) = A_i x$ on the original dataset (for $i = 1, \dots, N$), where A_i denotes the i 'th row of A . We picked $d = 100$, $N = 1000$ and randomly generated the entries of A from $\mathcal{N}(0, 1)$. To generate the output y_i for data point A_i , we first sampled a vector w from the interior of the d -ball with radius $R = 10$ and added noise: $y_i = A_i w + \xi_i$, where $\xi \sim \mathcal{N}(0, R^2/16)$ is the noise. We determined the optimal solution x^* using the closed-form solution $x^* = (A^T A)^{-1} A^T y$.

The problem is strongly convex with condition number $\kappa = \sigma_{\max}(A)/\sigma_{\min}(A)$, where $\sigma_{\max}(A)$ and $\sigma_{\min}(A)$ are the maximal and the minimal singular values of A , respectively. We set $p = (1/2)(1/(1 + \kappa))$ and $q = (3/4)(1/(1 + \kappa))$. At each iteration, we picked the corrupt gradient as $b_t = (R/\gamma_t)((x^* - x_t)/\|x^* - x_t\|)$, which can be shown to be worst possible adversarial value at iteration t .

In Figure 1 we compare the performances of several first-order optimization algorithms using mean squared error $\|x_t - x^*\|^2$ as the performance metric. The first-order




Fig. 2. Performances of first-order algorithms for Logistic Classification using training cost as metric.

algorithms such as Gradient Descent (GD), ADAM [6], Nesterov’s Accelerated Gradient Method (NAG) [2], RMSprop [5], and AMSGrad [43] fail to converge; whereas the normalized subgradient method (NSM) outlined in Algorithm 1 succeeds. Algorithms that do not normalize the gradient fail to converge due to the presence of the adversarial perturbations in the gradient.

B. Logistic Classification

We study the robustness of Algorithm 1 on L_2 -regularized multi-class logistic classification using the MNIST dataset [44]. The task is to determine $m = 10$ linear classifiers in order to separate $N = 6000$ $d = 784$ -dimensional image vectors. The problem can be stated as:

$$\min_{x \in \mathbb{R}^{m \times d}} -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{x_{y_i} A_i^T}}{\sum_{j=1}^m e^{x_{y_j} A_j^T}} + \lambda \|x\|^2, \quad (41)$$

where $A \in \mathbb{R}^{N \times d}$ is the matrix containing N data vectors in its rows (A_i denotes the i ’th row of A) and $y \in \{0, 1, \dots, 9\}^N$ is the vector containing the N associated classes. The decision parameter consists of m vectors with dimension d , where each vector corresponds to a class (hence, x_{y_i} corresponds to the y_i ’th row of x , where y_i is the class of i ’th data vector). We note that the problem is unconstrained since $\mathcal{X} = \mathbb{R}^{m \times d}$ and therefore the parameter can go unbounded. Furthermore, although the problem is convex, it is not trivial to determine whether the cost function satisfies the acute angle condition. Nevertheless, we numerically experiment whether normalized subgradient method can provide convergence in an adversarial setting.

We set $p = 0.25$, $\gamma_t = 0.1/t$, and $\lambda = 100$. At each iteration, we picked the corrupt gradient as $b_t = -15g(x_t)$, i.e., 15 times the opposite of the trustworthy gradient. In Figure 2 we compare the performances of several first order optimization algorithms using the value of the cost function $f(x_t)$ during training as metric. Similar to Linear Regression experiment, NSM converges while other algorithms diverge.

C. Threshold Probability

In this experiment, we demonstrate the significance of the threshold probability, for which our theoretical convergence




Fig. 3. Iterates generated by NSM for the simple example outlined in Section II for several corruption probabilities.

results hold, using the toy example outlined in Section II. Note that for the cost function $f(x) = x_{d-1}^4$, the acute angle condition holds with equality for $\cos \phi = 1/\sqrt{d}$. Therefore, the probability $1/(1 + \sqrt{d})$ is the exact threshold below which NSM converges and above which the adversary can cause NSM to diverge. We test this numerically for $d = 10$, $R = 10$, and $p_{\text{threshold}} = 1/(1 + \sqrt{d}) = 0.2403$. We initialized $x_{0,i} = 5$, $\forall i \in [1, \dots, d]$, and ran NSM for $p = \{0.1, 0.2, p_{\text{threshold}} - 0.01, p_{\text{threshold}}, p_{\text{threshold}} + 0.01, 0.4\}$ with $\gamma_t = 200/t$. At each iteration, we set $b_t = -x_t$ if $x_{t,0} \neq 0$; and $b_t = [1, \dots, 1]^T$ otherwise. In Figure 3, we plot $x_{t,0}$ versus the iteration index t for each corruption probability. We observe that for $p < p_{\text{threshold}}$, the iterate converges to the optimal solution and for $p > p_{\text{threshold}}$, the iterate diverges (saturates at the bound). For $p = p_{\text{threshold}}$, the iterate does not diverge, however, it converges to a suboptimal solution.

V. CONCLUSIONS

In this paper, we investigated the robustness of the normalized subgradient method in a randomly corrupted subgradient setting, where the subgradient at each iteration is arbitrarily corrupted with probability p . We have shown that if the cost function satisfies the acute angle condition for some angle ϕ and if p is less than $\cos \phi / (1 + \cos \phi)$, then the normalized subgradient method, with stepsizes that diminish as $\mathcal{O}(1/t)$, converges almost surely to a minimizer of the objective function. We provide sufficient conditions for convergence of the normalized (sub)gradient method for convex, strongly convex, and weakly pseudo-convex cost functions. Numerical examples show the robustness of the normalized subgradient method for linear regression and logistic classification problems.

REFERENCES

- [1] N. Qian, “On the momentum term in gradient descent learning algorithms,” *Neural networks*, vol. 12, no. 1, pp. 145–151, 1999.
- [2] Y. Nesterov, “A method for unconstrained convex minimization problem with the rate of convergence $\mathcal{O}(1/k^2)$,” in *Doklady an ussr*, vol. 269, 1983, pp. 543–547.
- [3] J. Duchi, E. Hazan, and Y. Singer, “Adaptive subgradient methods for online learning and stochastic optimization,” *Journal of machine learning research*, vol. 12, no. 7, 2011.
- [4] M. D. Zeiler, “Adadelta: an adaptive learning rate method,” *arXiv preprint arXiv:1212.5701*, 2012.

- [5] T. Tieleman and G. Hinton, "Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude," *COURSERA: Neural networks for machine learning*, vol. 4, no. 2, pp. 26–31, 2012.
- [6] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [7] T. Dozat, "Incorporating nesterov momentum into adam," in *Proceedings of 4th ICLR*, 2016.
- [8] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 10, no. 2, pp. 1–19, 2019.
- [9] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," in *Concurrency: the Works of Leslie Lamport*, 2019, pp. 203–226.
- [10] Y. Nesterov, *Introductory Lectures on Convex Optimization*. Springer Science & Business Media, 2004, vol. 87.
- [11] N. Z. Shor, *Minimization methods for non-differentiable functions*. Springer Science & Business Media, 2012, vol. 3.
- [12] K. Y. Levy, "The power of normalization: Faster evasion of saddle points," *arXiv preprint arXiv:1611.04831*, 2016.
- [13] R. Murray, B. Swenson, and S. Kar, "Revisiting normalized gradient descent: Fast evasion of saddle points," *IEEE Transactions on Automatic Control*, vol. 64, no. 11, pp. 4818–4824, 2019.
- [14] E. Hazan, K. Levy, and S. Shalev-Shwartz, "Beyond convexity: Stochastic quasi-convex optimization," in *Advances in Neural Information Processing Systems*, 2015, pp. 1594–1602.
- [15] X. Gao, X. Li, and S. Zhang, "Online learning with non-convex losses and non-stationary regret," in *International Conference on Artificial Intelligence and Statistics*, 2018, pp. 235–243.
- [16] C. Fang, C. J. Li, Z. Lin, and T. Zhang, "Spider: Near-optimal non-convex optimization via stochastic path-integrated differential estimator," in *Adv Neural Inf Process Syst*, 2018, pp. 689–699.
- [17] F. E. Curtis, K. Scheinberg, and R. Shi, "A stochastic trust region algorithm based on careful step normalization," *Inform Journal on Optimization*, vol. 1, no. 3, pp. 200–220, 2019.
- [18] L. Huang, A. D. Joseph, B. Nelson, B. I. Rubinstein, and J. D. Tygar, "Adversarial machine learning," in *Proceedings of the 4th ACM workshop on Security and artificial intelligence*, 2011, pp. 43–58.
- [19] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial machine learning at scale," *arXiv preprint arXiv:1611.01236*, 2016.
- [20] D. Lowd and C. Meek, "Adversarial learning," in *Proc. of the 11th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2005, pp. 641–647.
- [21] U. Jang, X. Wu, and S. Jha, "Objective metrics and gradient descent algorithms for adversarial examples in machine learning," in *Proceedings of the 33rd Annual Computer Security Applications Conference*, 2017, pp. 262–277.
- [22] N. Dalvi, P. Domingos, S. Sanghai, and D. Verma, "Adversarial classification," in *Proc. of the 10th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, 2004, pp. 99–108.
- [23] Y. Chen, C. Caramanis, and S. Mannor, "Robust sparse regression under adversarial corruption," in *International Conference on Machine Learning*, 2013, pp. 774–782.
- [24] K. Bhatia, P. Jain, and P. Kar, "Robust regression via hard thresholding," in *Adv Neural Inf Process Syst*, 2015, pp. 721–729.
- [25] B. McWilliams, G. Krummenacher, M. Lucic, and J. M. Buhmann, "Fast and robust least squares estimation in corrupted linear models," in *Adv Neural Inf Process Syst*, 2014, pp. 415–423.
- [26] I. Diakonikolas, G. Kamath, D. Kane, J. Li, J. Steinhardt, and A. Stewart, "Sever: A robust meta-algorithm for stochastic optimization," in *International Conference on Machine Learning*, 2019, pp. 1596–1606.
- [27] S. Mahloujifar and M. Mahmoody, "Blockwise p-tampering attacks on cryptographic primitives, extractors, and learners," in *Theory of Cryptography Conference*. Springer, 2017, pp. 245–279.
- [28] S. Mahloujifar, D. I. Diochnos, and M. Mahmoody, "Learning under p-tampering attacks," in *ALT*, 2018, pp. 572–596.
- [29] B. Turan, C. Uribe, H.-T. Wai, and M. Alizadeh, "Resilient primal-dual optimization algorithms for distributed resource allocation," *IEEE Transactions on Control of Network Systems*, 2020.
- [30] A. Ghosh, R. K. Maity, S. Kadhe, A. Mazumdar, and K. Ramchandran, "Communication-efficient and byzantine-robust distributed learning," *arXiv preprint arXiv:1911.09721*, 2019.
- [31] K. Pillutla, S. M. Kakade, and Z. Harchaoui, "Robust aggregation for federated learning," *arXiv preprint arXiv:1912.13445*, 2019.
- [32] Z. Wu, Q. Ling, T. Chen, and G. B. Giannakis, "Federated variance-reduced stochastic gradient descent with robustness to byzantine attacks," *IEEE Trans. Signal Process.*, vol. 68, pp. 4583–4596, 2020.
- [33] A. Ghosh, J. Hong, D. Yin, and K. Ramchandran, "Robust federated learning in a heterogeneous environment," *arXiv preprint arXiv:1906.06629*, 2019.
- [34] D. Yin, Y. Chen, K. Ramchandran, and P. Bartlett, "Byzantine-robust distributed learning: Towards optimal statistical rates," *arXiv preprint arXiv:1803.01498*, 2018.
- [35] D. Yin, Y. Chen, R. Kannan, and P. Bartlett, "Defending against saddle point attack in byzantine-robust distributed learning," in *International Conference on Machine Learning*. PMLR, 2019, pp. 7074–7084.
- [36] Y. Chen, L. Su, and J. Xu, "Distributed statistical machine learning in adversarial settings: Byzantine gradient descent," *POMACS*, vol. 1, no. 2, pp. 1–25, 2017.
- [37] J. Bernstein, J. Zhao, K. Azizzadenesheli, and A. Anandkumar, "signsgd with majority vote is communication efficient and fault tolerant," *arXiv preprint arXiv:1810.05291*, 2018.
- [38] S. Boyd and A. Mutapcic, "Stochastic subgradient methods," *Lecture Notes for EE364b, Stanford University*, 2008.
- [39] V. S. Borkar, *Stochastic approximation: a dynamical systems viewpoint*. Springer, 2009, vol. 48.
- [40] D. P. Bertsekas and J. N. Tsitsiklis, "Gradient convergence in gradient methods with errors," *SIAM Journal on Optimization*, vol. 10, no. 3, pp. 627–642, 2000.
- [41] Y. Nesterov and B. T. Polyak, "Cubic regularization of newton method and its global performance," *Mathematical Programming*, vol. 108, no. 1, pp. 177–205, 2006.
- [42] H. Robbins and D. Siegmund, "A convergence theorem for non negative almost supermartingales and some applications," in *Optimizing methods in statistics*. Elsevier, 1971, pp. 233–257.
- [43] S. J. Reddi, S. Kale, and S. Kumar, "On the convergence of adam and beyond," *arXiv preprint arXiv:1904.09237*, 2019.
- [44] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.