

Adversary-resilient Distributed and Decentralized Statistical Inference and Machine Learning

[An Overview of Recent Advances Under the Byzantine Threat Model]

Zhixiong Yang, Arpita Gang, and Waheed U. Bajwa

Statistical inference and machine learning algorithms have traditionally been developed for data available at a single location. Unlike this *centralized setting*, modern datasets are increasingly being distributed across multiple physical entities (sensors, devices, machines, data centers, etc.) for a multitude of reasons that range from storage, memory, and computational constraints to privacy concerns and engineering needs. This has necessitated the development of inference and learning algorithms capable of operating on non-located data. We divide such algorithms into two broad categories in this article, namely, *distributed algorithms* and *decentralized algorithms*.

Distributed Algorithms: Distributed algorithms correspond to setups in which the data-bearing entities (henceforth referred to as “nodes”) require some form of coordination through a specially designated entity in the system in order to generate the final result. Depending on the application, this special entity is referred to as master node, central server, parameter server, fusion center, etc., in the literature. While distributed setups can take a number of forms, this exposition mostly revolves around the so-called *master-worker distributed architecture* in which data-bearing nodes only communicate with a single master node that is tasked with generating the final result. Among other applications, such distributed architectures arise in the context of parallel computing, where the focus is computational speedups and/or overcoming memory/storage bottlenecks, and federated systems, where “raw” data collected by individual nodes cannot be shared with the master node due to either communication constraints (e.g., sensor networks) or privacy concerns (e.g., smartphone data).

Decentralized Algorithms: Decentralized algorithms correspond to setups that lack central servers; instead, data-bearing nodes in a decentralized system are *collectively* tasked with generating the final result. In particular, individual nodes in a decentralized setup typically communicate among themselves over a network (often ad hoc) to reach a common solution (i.e., achieve *consensus*) at all nodes. Decentralized setups arise either out of the need to eliminate single points of failure in distributed setups or due to practical engineering constraints, as in the internet of things and autonomous systems.




Fig. 1. Inference and machine learning algorithms involving non-located data can be broadly divided into the categories of (i) distributed algorithms and (ii) decentralized algorithms. The former category includes master-worker distributed setups such as parallel computing and federated systems in which a single entity is tasked with generating the final result. The latter category deals with decentralized setups in which entities cooperate among themselves to produce the final result.

Is it distributed or is it decentralized? Inference and learning from non-located data have been studied for decades in computer science, control, signal processing, and statistics. Both among and within these disciplines, however, there is no consensus on use of the terms “distributed” and “decentralized.” Though many works share the definitions provided in here, there are numerous authors who use these two terms interchangeably, while there are some other authors who reverse these definitions.

The last few decades have witnessed the emergence of a plethora of applications that necessitate advances in inference and learning in distributed and decentralized setups. Indeed, distributed and decentralized statistical inference methods are increasingly being relied upon in urban traffic monitoring, environmental sensing, management of smart grids, distributed spectrum sensing, and homeland security, among other applications. Similarly, distributed and decentralized machine learning algorithms are increasingly being utilized in the context of networks of self-driving cars, control of robot swarms, pattern recognition in large-scale datasets, and federated learning systems for healthcare data. Collectively, these applications have resulted in the development of a huge body of work devoted to understanding the algorithmic and theoretical underpinnings of distributed and decentralized inference and learning. But much of this work assumes a non-adversarial setting in which individual nodes—apart from occasional statistical failures—operate as intended within the algorithmic framework.

In recent years, however, cybersecurity threats from malicious non-state actors and rogue entities—and the potentially disastrous consequences of these threats for the aforementioned applications—have forced

practitioners and researchers to rethink the robustness of distributed and decentralized algorithms against adversarial attacks. As a result, we now have an abundance of algorithmic approaches that guarantee robustness of distributed and/or decentralized inference and learning under different adversarial threat models; see, e.g., recent survey articles [1]–[3]. Driven in part by the world’s growing appetite for data-driven decision making, however, securing of distributed/decentralized frameworks for inference and learning against adversarial threats remains a rapidly evolving research area. In this article, we provide an overview of some of the most recent developments in this area under the threat model of *Byzantine attacks*. This threat model—which subsumes the case of malfunctioning nodes, referred to as *Byzantine faults/failures*—is one of the hardest to safeguard against since it allows for undetected takeover of nodes by the adversary. In particular, nodes affected under this threat model—termed *Byzantine nodes*—are assumed to have the potential to arbitrarily bias the outputs of the underlying algorithms by colluding among themselves, injecting false data and information into the distributed/decentralized system, etc. This is in stark contrast to the relatively simpler threat model of *crash faults*, in which individual nodes in the distributed/decentralized system continue to operate as intended till a crash fault occurs, at which point the faulty node ceases to interact with the system (rather than potentially injecting false information into the system). We refer the reader to [4], and the references within, for further discussion on both the generality and the hardness of the Byzantine threat model in relation to the crash-fault model.

Origination of the Byzantine threat model: The threat model of Byzantine attacks/faults/failures in its most general form was introduced and analyzed in [5] within the context of reliability of computer systems with potentially malfunctioning components. The overarching problem in [5] was abstracted as the *Byzantine Generals Problem*, in which several generals of the Byzantine army—some of whom are likely to be traitors—need to agree on attacking or retreating from an enemy city through exchange of messages via a messenger; see, e.g., Figure 2. The authors reported two key results in [5] for this problem. First, they established the impossibility of safeguarding against Byzantine nodes (traitor generals) when the number of uncompromised nodes (loyal generals) is not more than two-thirds of the total number of nodes. The threat from General 3 in Figure 2, therefore, cannot be neutralized by any algorithm since the number of loyal generals is only two in this scenario. Second, the authors proposed two




Fig. 2. A Byzantine army led by three generals, one of whom (General 3) is a traitor, surrounding an enemy city. The loyal generals are trying to reach a consensus on the plan of action against the enemy, while the traitor is trying to mislead them.

algorithms for *complete* and *regular* graphs that provably counter the actions of Byzantine nodes in the case of binary (e.g., attack/retreat) decisions as long as the number of uncompromised nodes exceeds two-thirds of the total number of nodes.

The Byzantine threat model, since its introduction in 1982, has been extensively studied in the context of statistical inference from non-located data [1], [2]. However, its potential impact on more general decentralized consensus and distributed/decentralized machine learning problems has only been recently studied. It is against this backdrop that this article summarizes recent developments in Byzantine-resilient processing of non-located data, with a majority of the discussion focused on machine learning problems in distributed and decentralized settings.

I. ADVERSARY-RESILIENT DISTRIBUTED PROCESSING OF DATA

In the classic master–worker setting of distributed systems tasked with processing of non-located data, there is one central server—referred to as master node, parameter server, fusion center, etc.—that coordinates with M data-bearing nodes (sensors, smartphones, worker nodes, etc.) for final decision making. While all nodes in this setting communicate with the server, they usually cannot communicate with each other (see “Distributed Setups” in Figure 1). The Byzantine threat model in this setting assumes at most b nodes in the system have been compromised. This parameter b , which typically corresponds to a crude upper bound on the exact number of Byzantine nodes, often plays an important role in both the analysis and the performance of Byzantine-resilient algorithms. We summarize some of these algorithms and their theoretical guarantees in the following for statistical inference and machine learning problems.

A. Distributed Statistical Inference

Statistical inference leverages data samples in order to draw conclusions about the underlying probability distribution(s) generating the data. While statistical inference can take many forms, we limit our discussion in this article to Byzantine-resilient distributed detection and distributed estimation.

Distributed Detection: Distributed detection under both the *Neyman–Pearson* (NP) and the *Bayesian* frameworks has a rich history. Given two hypotheses \mathcal{H}_0 and \mathcal{H}_1 , a typical distributed detection algorithm first involves each node taking a local decision in favor of either \mathcal{H}_0 or \mathcal{H}_1 based on its own data samples. The nodes then send their decisions to the central server, which applies a fusion rule to the local decisions in order to reach the final (global) decision. Unfortunately, distributed detection algorithms designed without consideration of potential Byzantine failures break down in the presence of Byzantine nodes.

Despite the brittleness of traditional distributed detection techniques, investigation of Byzantine-resilient distributed detection only took off in the last decade. The survey article [1] provides an overview of many

TABLE I
SUMMARY OF RECENT RESULTS CONCERNING BYZANTINE-RESILIENT DISTRIBUTED DETECTION

Reference	Hypothesis Test	Detection Setup	Critical Fraction α^*	Distributed Architecture
Nadendla et al. [6]	Q -ary	Neyman–Pearson	$\alpha^* = \frac{Q-1}{Q}$	Star topology
Kaillkhura et al. [7]	Binary	Bayesian	$\alpha^* = 0.5$	Star topology
Hashlamoun et al. [8]	Binary	Bayesian	Can have $\alpha^* = 1$	Quasi-star topology
Kaillkhura et al. [9]	Binary	Neyman–Pearson	$b_1 = \lceil \frac{M_1}{2} \rceil$	Tree topology

of the resulting methods and asymptotically characterizes the respective *critical fraction* α^* of Byzantine nodes, defined as $\alpha^* := \frac{b}{M}$, at (or beyond) which the central server cannot do better than random guessing for its final decision. An important insight from the earliest works on Byzantine-resilient distributed detection is that (asymptotically) α^* can be $\frac{1}{2}$ or higher [1]; in contrast, recall that $\alpha^* = \frac{1}{3}$ for the original Byzantine Generals Problem [5].

Since the appearance of [1] in 2013, a few other works on Byzantine-resilient distributed detection have appeared in the literature. It is shown in [6] that distributed detection can be more resilient to Byzantine failures in the case of a general Q -ary hypothesis testing problem, with the critical fraction given by $\alpha^* = \frac{Q-1}{Q}$. Byzantine-resilient distributed binary hypothesis testing is investigated for the first time under the Bayesian framework in [7], with the critical fraction also given by $\alpha^* = \frac{1}{2}$ in this case. Finally, it is established in [8] that—under certain conditions—it is possible to have $\alpha^* = 1$ in the case of Bayesian distributed binary detection as long as each node is allowed to replicate its message to the central server through one other node in the system. Strictly speaking, however, this coordination among pairs of nodes leads to a distributed architecture that differs from the distributed master–worker (star topology) architecture of prior works. We conclude by noting that a tree topology in which the central server sits at the root of the tree (Depth 0) and nodes in the distributed system route their messages to the server through their parent nodes is another distributed architecture that does not fall under the distributed master–worker setup. Byzantine resilience of such architectures in distributed detection tasks is investigated in [9], with the critical fraction of Byzantine nodes defined in terms of the number of Byzantine nodes b_1 and the total number of nodes M_1 at Depth 1 of the tree. We also refer the reader to Table I for a summary of all the results that have been discussed in this article in relation to Byzantine-resilient distributed detection.

Distributed Estimation: Byzantine-resilient distributed estimation has received significant attention lately in the context of state estimation in cyberphysical systems. Much of the developments in this regard have been limited to linear models, with a typical observation model at any given time at node

j expressed as $y_j = \mathbf{h}_j^T \mathbf{w} + \eta_j$, $j = 1, \dots, M$, where $\mathbf{h}_j \in \mathbb{R}^d$ denotes a known vector, $\mathbf{w} \in \mathbb{R}^d$ is the unknown state vector that needs to be estimated at the central server, and η_j denotes the observation noise at node j . Traditionally, distributed estimation has involved the nodes transmitting y_j 's (or their quantized versions) to the central server and the server estimating \mathbf{w} using some variant of the following least-squares formulation [10] (or a maximum likelihood one in the case of quantized transmissions):

$$\hat{\mathbf{w}} = \arg \min_{\tilde{\mathbf{w}}} \|\mathbf{y} - \mathbf{H}\tilde{\mathbf{w}}\|_2, \quad (1)$$

where the matrix $\mathbf{H} \in \mathbb{R}^{M \times d}$ has \mathbf{h}_j 's as its rows. The quality of such solutions are assessed in terms of gap of their *mean squared error* (MSE), $\mathbb{E}[\|\hat{\mathbf{w}} - \mathbf{w}\|_2^2]$, from either the *minimum mean squared error* (MMSE) or the *Cramer–Rao lower bound* (CRLB) for unbiased estimators.

In the presence of Byzantine nodes, in which case the y_j 's corresponding to the b Byzantine nodes can be arbitrarily different from $\mathbf{h}_j^T \mathbf{w}$, solutions of distributed estimation techniques based on (1) can be pushed far from the optimal (in terms of MMSE/CRLB); see, e.g., Figure 3. In this setting, one can again characterize the critical fraction α^* of Byzantine nodes at (or beyond) which the server can do no better than relying on prior knowledge about \mathbf{w} . It is, for instance, established in [6] that—under certain assumptions—it is possible to have $\alpha^* = \frac{Q-1}{Q}$ when nodes utilize Q -ary quantization in order to transmit their data to the server. The survey articles [1]–[3] provide additional discussion of works on Byzantine-resilient distributed estimation, many of which are based on the idea of either detection of Byzantine attacks and/or identification of individual Byzantine nodes as part of their mitigation strategy. Finally, while much of the focus in Byzantine-resilient distributed estimation has been on linear models, there are works like [11] that focus on nonlinear models. Since there is a significant overlap between statistical estimation and machine learning problems, we do not indulge in discussion of such works; instead, we move towards our review of recent results on Byzantine-resilient distributed machine learning.

Distributed estimation and Byzantine failures: We illustrate the impact of Byzantine nodes on distributed estimation through a simple example of two-dimensional parameter estimation, in which the linear observation model corresponds to a line in (x, y) plane, using $M = 8$ nodes. The observations at each node in this example correspond to $y_j = \mathbf{h}_j^T \mathbf{w} + \eta_j$ with $\mathbf{h}_j^T := [x_j \ 1]$ and the two-dimensional vector \mathbf{w} describing the slope and y -intercept of the line. It can be seen from Figure 3 that Byzantine failures of just two nodes returns a least-squares solution that results in a model $\mathbf{h}_j^T \hat{\mathbf{w}}$ that significantly differs from the true model $\mathbf{h}_j^T \mathbf{w}$.




Fig. 3. The impact of Byzantine failures on traditional least-squares estimation.

B. Distributed Machine Learning

A typical challenge in machine learning is to statistically minimize a function f , referred to as *loss function* or *risk function*, with respect to a candidate *model* $\mathbf{w} \in \mathbb{R}^d$ that describes the data, i.e.,

$$\min_{\mathbf{w}} \mathbb{E}_{\mathbf{z} \sim \mathcal{P}} [f(\mathbf{w}, \mathbf{z})], \quad (2)$$

where \mathbf{z} denotes data living in some Hilbert space that is drawn from an *unknown* probability distribution \mathcal{P} . To solve (2) without knowledge of \mathcal{P} , distributed machine learning focuses on minimizing an empirical variant of $\mathbb{E}_{\mathbf{z} \sim \mathcal{P}} [f(\mathbf{w}, \mathbf{z})]$ using samples of \mathbf{z} —termed *training data*—distributed across different nodes. The resulting objective, aptly termed distributed *empirical risk minimization* (ERM), can be expressed as $\min_{\mathbf{w}} \frac{1}{M} \sum_{j=1}^M f(\mathbf{w}, \mathbf{Z}_j)$, with \mathbf{Z}_j denoting the local training data at node j that comprises multiple samples drawn from \mathcal{P} and \mathbf{w} referred to as the *global* optimization variable. The model/variable \mathbf{w} in distributed machine learning is stored at the central server, which iteratively updates it based on messages received from individual nodes and subsequently sends the updated \mathbf{w} to all nodes. Each node, in turns, performs some computation according to its local data and the received \mathbf{w} , and sends a message back to the server.

Distributed Stochastic Gradient Descent: Similar to distributed statistical inference, Byzantine failures can lead to breakdowns of distributed machine learning methods. Motivated in part by the widespread adoption of deep neural networks and variance-reduction techniques like mini batching by the practitioners, we mainly focus here on robustification of synchronous distributed *stochastic gradient descent* (SGD) against Byzantine attacks. A typical distributed SGD algorithm proceeds iteratively. In each iteration, nodes compute the gradient of the loss function on their local data with respect to the current model and send their gradients to the server. The server, in turn, takes the average of all the gradients and updates the global variable according to the averaged gradient. In a faultless environment, distributed SGD—with proper choices of step size and batch size—has a linear convergence rate for strongly convex functions. However, a single Byzantine node can force the algorithm to converge to any model using a simple strategy. Suppose, for instance, the summation of gradients of faultless nodes is \mathbf{g} and the Byzantine node wishes the server to operate on an alternate gradient \mathbf{g}' . The Byzantine node can accomplish this by sending $M\mathbf{g}' - \mathbf{g}$ as its gradient to the server, thereby controlling the update step at the server. Of course, as simple as this strategy is, it is not an optimal one for the Byzantine node; indeed, a large M will result in a large gradient, which is likely to make the malicious message detectable. We refer the reader to [12]–[14] for more sophisticated strategies that can be employed by Byzantine nodes.

Several algorithms have been put forth recently to safeguard distributed SGD against Byzantine failures. The central idea in all these approaches that imparts Byzantine resilience to distributed SGD involves the use of a *screening* procedure at the server while it aggregates the local gradients. Using an appropriate

screening rule, the screened aggregated gradient can be shown to be close to the true average gradient, thereby enabling the server to approximately solve the distributed ERM problem. Algorithm 1 provides a general framework for Byzantine-resilient distributed SGD that is based on this idea of screening. While some algorithms alter this framework a little bit, the general idea remains the same. Note that if one were to remove the screening procedure in Step 10 of the algorithm, it becomes vanilla distributed SGD.

Algorithm 1 A General Framework for Byzantine-resilient Distributed SGD

```

1: for  $t = 1, 2, \dots$  do
2:   Server:
3:     Send the global optimization variable to all nodes
4:   Node:
5:     Receive the global optimization variable from the server
6:     Calculate gradient with respect to the local training set
7:     Send the local gradient to the server
8:   Server:
9:     Receive local gradients from all nodes
10:    Screen the received gradients for Byzantine resilience and aggregate them
11:    Update the global variable by taking a gradient step using the aggregated gradient
12: end for

```

The State of the Art: We now discuss the various screening (and aggregation) procedures adopted in different algorithms. The algorithm introduced in [15], termed *robust distributed gradient descent*, uses two screening methods: *coordinate-wise median* and *coordinate-wise trimmed mean*. In coordinate-wise median, the server aggregates the local gradients by taking the median in each dimension of the gradients. Coordinate-wise trimmed mean, on the other hand, involves the server eliminating the smallest and the largest b values in each dimension of the gradients and coordinate-wise averaging the remaining values for aggregation. The GeoMed algorithm [16], in contrast, uses the *geometric median* of local gradients as the screening and aggregation rule. The Krum [17] algorithm, on the other hand, finds the local gradient that has the smallest distance to its $M - b - 2$ closest gradients and uses this gradient for the update step. There is also a variant of the Krum algorithm, termed *Multi-Krum* [17], which finds $m \in \{1, \dots, M\}$ local gradients using the Krum principle and uses an average of these gradients for update.

The Bulyan algorithm [18], [19] is a two-stage algorithm. In the first stage, it recursively uses vector median methods such as Geometric median and Krum to select $M - 2b$ local gradients. In the second stage, it carries out a coordinate-wise operation on the $M - 2b$ selected gradients in which $M - 4b$

values in each coordinate are retained (and then averaged for aggregation) by eliminating $2b$ values that are farthest from the coordinate-wise median. Zeno/Zeno++ [20], [21] are related algorithms that require an oracle at the server, which can generate an estimate of the true gradient in each iteration, for screening/aggregation purposes. In particular, the screening procedure involves calculating a score for each local gradient based on its difference from the oracle gradient, with the surviving gradients being the ones that are the most similar to the oracle gradient.

The Byzantine-resilient distributed SGD framework of Algorithm 1 has also been investigated in [22] under a *generalized* Byzantine threat model. The basic assumption underlying this threat model is that the set of nodes under Byzantine attacks can change in each iteration of the algorithm, but the attackers can only inject malicious information in some dimensions of their respective gradients. Similar to [15], [16], the work in [22] puts forth the use of coordinate-wise median, a variant of coordinate-wise trimmed mean, and geometric median for screening purposes, and establishes resilience of the resulting methods under the assumption that less than half of the gradients in each dimension are attacked in each iteration.

We conclude by discussing a few methods that somewhat deviate from the distributed framework of Algorithm 1. In contrast to [15]–[22], the RSA algorithm of [23] robustifies distributed SGD by making individual nodes store and update local versions of the global optimization variable \mathbf{w} , which are then aggregated at the server in each iteration in a Byzantine-resilient manner. In order to reduce the communications overhead of Byzantine-resilient distributed SGD, [24] discusses the signSGD algorithm, which only uses the element-wise sign of the gradient—rather than the gradient itself—for the update step. In particular, it is established in [24] that signSGD with an element-wise majority vote on the signs of the local gradients as a screening/aggregation rule is a Byzantine-resilient learning method. Finally, [25] proposes and analyzes a variant of distributed SGD that requires only one pass over the entire training data. Unlike other works, however, resilience in this work is accomplished through explicit labeling of nodes as Byzantine, which are then excluded from all future computations at the server.

We next summarize key aspects of some of the Byzantine-resilient distributed learning algorithms in Table II in terms of the number of nodes M , the number of Byzantine nodes b , and the number of independent and identically distributed (i.i.d.) samples per node N . The convergence rates in the table correspond to the gap between the learned model after t iterations and the minimizer of the statistical risk (cf. (2)) in the limit of large N . The parameter c denotes a constant that may change from one algorithm to the other, the $\mathcal{O}(\cdot)$ scaling hides dependence on problem parameters (including dimension d in the top part of the table), N/A signifies an algorithm lacks a particular result, and — means guarantees for an algorithm are not directly comparable to other algorithms. The last column in the top part of the table

TABLE II

SUMMARY OF SOME RECENT RESULTS CONCERNING BYZANTINE-RESILIENT DISTRIBUTED MACHINE LEARNING

Algorithm	Convergence Rate	Statistical Learning Rate	Condition on (M, b)
Coordinate-wise Median (CM) [15]	$\mathcal{O}(c^t)$	$\mathcal{O}\left(\frac{b}{M\sqrt{N}} + \frac{1}{\sqrt{MN}} + \frac{1}{N}\right)$	$M \geq 2b + 1$
Coordinate-wise Trimmed Mean (CTM) [15]	$\mathcal{O}(c^t)$	$\mathcal{O}\left(\frac{b}{M\sqrt{N}} + \frac{1}{\sqrt{MN}}\right)$	$M \geq 2b + 1$
GeoMed [16]	$\mathcal{O}(c^t)$	$\mathcal{O}\left(\frac{\sqrt{b}}{\sqrt{MN}}\right)$	$M \geq 2b + 1$
Krum [17]	N/A	N/A	$M \geq 2b + 3$
Multi-Krum [17]	N/A	N/A	$M \geq 2b + m + 2$
Bulyan [18]	N/A	N/A	$M \geq 4b + 3$
Zeno/Zeno++ [20], [21]	$\mathcal{O}(c^t) + \mathcal{O}(1)$	N/A	$M \geq b + 1$
RSA [23]	$\mathcal{O}\left(\frac{1}{t}\right) + \mathcal{O}(1)$	N/A	$M \geq b + 1$
signSGD [24]	—	N/A	$M \geq 2b + 1$

Algorithm	CM, CTM, Zeno/Zeno++	GeoMed	Krum, Multi-Krum	Bulyan
Screening Complexity	$\mathcal{O}(Md)$	$\mathcal{O}\left(Md + bd \log^3\left(\frac{1}{\gamma}\right)\right)^*$	$\mathcal{O}(M^2d)$	$\mathcal{O}(M^2d + Md)$

* Screening computational complexity for GeoMed is for computing $(1 + \gamma)$ -approximate geometric median [16].

lists conditions on M and b necessary for well-posedness of different algorithms. The bottom part of the table lists per-iteration computational complexity of the screening procedure for algorithms that involve an explicit screening step at the server. While the convergence/learning rates in the table are for strongly convex and smooth loss functions, some of the works require further assumptions and/or also provide results under a relaxed set of assumptions on the loss function and training data. We refer the reader to the referenced works for further details.

Screening and aggregation in two dimensions: We illustrate the robustness of different screening/aggregation methods against Byzantine attacks through a simple two-dimensional example. Consider a total of seven two-dimensional points (g_1, g_2) , out of which six represent correct data (black discs in Figure 4) and one of which has been falsified by a Byzantine node (red disc in Figure 4).

It can be seen from Figure 4, which is displayed with non-uniform g_1 - and g_2 -axes to capture the effect of the Byzantine node, that the ordinary average operation is highly susceptible to the falsified data point. On the other hand, screening and aggregation rules such as coordinate-wise median, coordinate-wise trimmed mean, Krum and Bulyan (using Krum in Figure 4) all produce final results that stay close to the true average of the faultless data points.




Fig. 4. An illustration of the effects of data falsified by Byzantine nodes on screening methods.

Numerical Experiments: It can be seen from Table II that (coordinate-wise) median, (coordinate-wise) trimmed mean, and GeoMed have some of the best theoretical guarantees for strongly convex and smooth loss functions in terms of convergence and learning rates. We now numerically compare the performance of most of the algorithms listed in Table II on nonconvex loss functions. Our comparison excludes GeoMed, since it lacks an algorithm for computing the exact geometric median of a set of high-dimensional gradients, as well as signSGD and RSA, since they differ from the rest of the algorithms in terms of their approach to Byzantine resilience.

The nonconvex learning task in the experiments corresponds to classification of the ten classes in CIFAR-10 dataset using a convolutional neural network with three convolution layers followed by two fully connected layers. Each of the three convolution layers are followed by ReLU activation and max pooling, while the output layer uses softmax activation. The distributed setup corresponds to one server and $M = 20$ nodes, with the 50,000 sample training set uniformly at random distributed across the system (i.e., each node has $N = 2,500$ training samples). We run two rounds of experiments for each algorithm, where each round is repeated 10 times—with each trial corresponding to 20,000 algorithmic iterations—and the results on the CIFAR-10 test set are averaged over these 10 trials. In the first round, none of the nodes are taken to be Byzantine nodes. In the second round, four nodes are randomly selected as Byzantine nodes, with each one sending a random vector to the server in each iteration whose elements uniformly take values in the range $(0, 10^{-5})$ for odd-numbered iterations and $(0, 20)$ for even-numbered iterations. We limit ourselves to four Byzantine nodes in order to provide a fairer comparison between different algorithms, since five Byzantine nodes exceeds the theoretical limit of some algorithms (e.g., Bulyan). In both rounds of experiments, the algorithms operate under the assumption of $b = 4$. We conclude our discussion of the experimental setup by noting that the optimal Byzantine attack strategy that adversely affects all distributed learning algorithms in a uniform manner remains an open problem. Nonetheless, the attack strategy being employed in our experiments has been carefully designed in light of the discussions in prior works (see, e.g., [12]); in particular, it appears to be the uniformly most potent strategy for the distributed learning algorithms under consideration in this article.

The results of the experiments in terms of average classification accuracy are shown in Figure 5, which highlight some trade-offs that should help the practitioners select algorithms that best fit their needs. We first note from Table II that the computational complexity of screening steps in Krum and Bulyan scales quadratically with the number of nodes M , whereas it scales only linearly with M in median, trimmed mean, and Zeno/Zeno++. Next, it can be seen from Figure 5 that all screening-based methods fall short of distributed SGD’s performance in the faultless setting. This is in line with the conventional wisdom




Fig. 5. Comparison of different distributed learning methods (with median and trimmed mean being coordinate wise) based on stochastic gradient descent. The left panel corresponds to the faultless setting, while the right panel corresponds to the case of four Byzantine nodes in the system. In both cases, the parameter b is set equal to 4 in all Byzantine-resilient algorithms. We have also overlaid plots of two decentralized machine learning methods, namely, BRIDGE and DGD, on top of the ones for distributed learning methods; these two methods are discussed later in the article (see “Decentralized Machine Learning”).

that robustness often comes at the expense of correctness, with median and Krum paying the highest price in terms of correctness, while Bulyan, trimmed mean, and Zeno paying the least (and somewhat insignificant) price. In the presence of Byzantine nodes, on the other hand, the vanilla distributed SGD completely falls apart. In contrast, while varying degrees of reduction in the performance of screening-based methods are observed under Byzantine attacks, none of them breaks down to the level of distributed SGD. Nonetheless, median, trimmed mean, and Krum (in this particular order) are affected the most in terms of performance by the aggressive Byzantine attack strategy employed in our experiments. Bulyan and Zeno, on the other hand, seem to be the most stable under Byzantine attacks. In addition, both these methods offer a competitive tradeoff between correctness (in a faultless setting) and robustness (under attack). However, it is worth pointing out that Zeno’s resilience comes at the expense of an oracle that can provide it with some knowledge of the true gradient. Similarly, since Bulyan screens four times the number of Byzantine nodes in each iteration, its resilience comes at the expense of a limited number of Byzantine nodes that it can handle (cf. Table II).

II. ADVERSARY-RESILIENT DECENTRALIZED PROCESSING OF DATA

Distributed systems in general, and distributed master–worker architectures in particular, have become the workhorse framework for non-centralized processing of data. The reasons for this range from relative ease of implementation and subsequent scaling up or scaling down of the system through addition or

removal of nodes to relative simplicity of synchronization and communication protocols. Despite these and other related advantages, decentralized setups for inference and learning are increasingly being investigated in lieu of distributed setups for a multitude of reasons. Unlike the distributed setup, a decentralized system lacks a central server that is connected to every node (see “Decentralized Setup” in Figure 1). Instead, all nodes in a decentralized system maintain local copies of the decision variable/machine learning model and reach (approximate) consensus on a common solution through periodic exchange of messages over a network with a subset of other nodes, termed *neighbors*. Because of this reason, and unlike distributed systems, there is no single point of failure in decentralized systems. In the same vein, whereas the star communications topology in distributed master–worker architectures can create a communications bottleneck at the central server (see “Distributed Setups” in Figure 1), the network/communications topologies in decentralized systems can be carefully designed to avoid such bottlenecks. One of the biggest reasons for the study of Byzantine-resilient inference and learning in decentralized systems, however, is the emergence of inherently decentralized applications such as networks of self-driving cars and robot swarms. Both centralized and distributed setups typically cannot be engineered in a cost-effective manner in such applications, which are generally studied under the moniker of *multiagent systems*. This necessitates the use of decentralized frameworks, often with ad-hoc network topologies, for statistical inference and machine learning.

Since no single node in decentralized setups gets access to the entire set of local variables, ensuring robustness against Byzantine failures is generally harder in decentralized systems. In particular, unlike simple characterizations of the feasible and/or necessary relationships between the number of nodes M and the number of Byzantine nodes b in Table I and Table II, the necessary and/or sufficient conditions for Byzantine-resilient decentralized algorithms are stated in terms of topology of the underlying network graph. In addition, while practical benefits of asynchronous update rules for and impact of time-varying network topologies on decentralized processing of data have long been investigated by researchers, the literature on Byzantine-resilient decentralized processing is relatively sparse in this regard. Because of this reason, and due to space constraints, our discussion in this article is mostly limited to synchronous algorithms on decentralized networks whose topology does not change with time (i.e., *static graphs*). We begin by engaging in discussion of an algorithmic process, aptly termed *consensus*, that often forms the basis of algorithms for decentralized inference and machine learning.

A. Decentralized Consensus

Many decentralized inference and learning algorithms require a subprocess that ensures (approximate) agreement, i.e., consensus, among all nodes. In the context of Byzantine-resilient decentralized processing,

therefore, it is instructive to understand the fundamentals of Byzantine-resilient consensus. Our discussion in this regard will be limited to linear strategies for consensus, which form the basis for Byzantine-resilient decision consensus, averaging consensus, convex consensus, etc. [26], [27].

Suppose each node j in the decentralized system has a local variable \mathbf{w}_j^0 and it is desired to reach consensus on the average of these variables at each node through in-network message passing, i.e., $\mathbf{w}_j \approx \frac{1}{M} \sum_{i=1}^M \mathbf{w}_i^0$ for every node j . Consensus algorithms usually proceed by having each node iteratively take a weighted average of their neighbors' variables. Mathematically, therefore, a consensus algorithm is associated with an $M \times M$ *weight matrix* whose (j, i) -th entry corresponds to the weight assigned by node j to the variable it receives from node i . In the absence of Byzantine nodes, a well-known result says that a doubly stochastic weight matrix whose smallest nonzero entry is lower bounded guarantees convergence of all \mathbf{w}_j 's to the average $\frac{1}{M} \sum_{i=1}^M \mathbf{w}_i^0$. Suppose, however, that node k in the system is Byzantine and it transmits $\mathbf{w}_k = \mathbf{w}'$ to its neighbors in each iteration. This simple "lazy" Byzantine strategy is enough for all nodes to converge to \mathbf{w}' . And if there were another Byzantine node k' in the system that always transmitted the same $\mathbf{w}_{k'} \neq \mathbf{w}'$, nodes in the system will not reach consensus at all.

The main idea behind ensuring robustness of (averaging) consensus to Byzantine nodes is to screen potential outliers at each node before the weighting step in each iteration. Since averaging of finite-dimensional vectors is equivalent to averaging of their individual respective (scalar) coordinates, much of the discussion in Byzantine-resilient consensus has focused on scalar-valued problems. Some variant of the trimmed mean, in which a node removes the largest b and the smallest b values received from its neighbors (including itself), is in particular a common screening method used in scalar consensus. Clearly, each node has to have enough neighbors (e.g., more than $2b$ for trimmed mean) for screening-based consensus to be well defined. However, since the network must not become disconnected after screening, the exact constraint on network topology for different algorithms tends to be more involved; see [26], [27] for further discussion of different screening methods and the corresponding topology constraints. Note that decentralized inference and learning algorithms utilizing similar screening ideas for Byzantine resilience tend to have similar topology constraints.

In terms of the performance of Byzantine-resilient consensus, even the best screening/aggregation method cannot guarantee removal of all falsified data. However, equipped with a reasonable strategy, a Byzantine-resilient scalar consensus algorithm can guarantee two things in each iteration: (i) the retained values are between the smallest and the largest scalars at all nonfaulty nodes; and (ii) the difference between the largest and the smallest nonfaulty values decreases. Together, these two conditions can ensure that the nonfaulty nodes converge to a common value. However, since a node can also retain falsified data

and/or eliminate nonfaulty data, consensus to the true average is no longer possible. Instead, guarantees can only be given to an approximate average such as a convex combination of initial values. This in particular is what makes Byzantine-resilient decentralized inference and learning so challenging.

We conclude by noting that asynchronous variants of decentralized consensus in the presence of Byzantine attacks have been investigated in [28], [29]. The asynchronicity in [29] comes through the use of *randomized gossip* algorithms, but the focus in that work is primarily on detection of Byzantine nodes. In contrast, [28] extends the synchronous (screening) framework and analysis of [26] to asynchronous settings (and time-varying network topologies). A noteworthy implication of [28] is that asynchronicity does not impose additional topology constraints required for resilience of decentralized consensus algorithms. This is important since synchronous algorithms often incur additional latency; indeed, a synchronous algorithm can only be as fast as the slowest connection/node in the system.

B. Decentralized Statistical Inference

Similar to the case of distributed inference, we limit ourselves in here to decentralized detection and estimation. In both problems, decentralized consensus makes an integral part of the developed algorithms.

Decentralized Detection: Similar to distributed detection, a typical setup in decentralized detection involves nodes in the system making observations under one of two (or more) hypotheses. Unlike distributed detection, however, nodes must collaborate among themselves to reach consensus in favor of one of the hypotheses. Vulnerability of the vanilla consensus framework to Byzantine nodes, therefore, also makes decentralized detection highly susceptible to Byzantine attacks.

While several application scenarios call for Byzantine-resilient decentralized detection, it has received less attention compared to its distributed counterpart. Some of the most relevant works in this regard include [30]–[32], all of which focus on scalar-valued problems. In the spirit of Byzantine-resilient scalar consensus, an adaptive threshold-based screening method is utilized in [30] to mitigate the impact of Byzantine nodes on the final decision. Similarly, a robust variant of distributed *alternating direction method of multipliers* (ADMM) is proposed in [31] that uses trimmed mean to eliminate $2b$ scalars at each node in every iteration. But other than the trivial topology constraint imposed by the trimmed mean (i.e., in-degree of nodes being greater than $2b$), this work lacks guarantees. In contrast, [32] proposes and analyzes a robust detection scheme that involves identification of Byzantine nodes and use of a weighted average consensus algorithm. However, this work focuses on a particular variant of Byzantine attacks and does not characterize topology constraints as a function of the number of Byzantine nodes.

Decentralized Estimation: Decentralized estimation, in which nodes use local observations and messages from neighbors to reach consensus on the estimate of an unknown parameter $\mathbf{w} \in \mathbb{R}^d$, is often

also studied under the linear model. Similar to decentralized detection, Byzantine-resilient decentralized estimation is a relatively recent topic of interest. Among relevant works, [33] focuses on a specialized estimation problem in which node j need only estimate the j -th entry of \mathbf{w} and there are three types of nodes in the system: *reliable* nodes, which are special nodes having perfect knowledge of the entry of \mathbf{w} associated with them, ordinary (normal) nodes, and Byzantine nodes. Under this setup and the assumption of a time-varying directed graph, [33] proposes and analyzes a trimmed mean-type (scalar-valued) screening procedure for Byzantine resilience. The works [34], [35] study vector-valued dynamic estimation under the noiseless observation model of $\mathbf{y}_j[n] = \mathbf{w}$, $n = 1, \dots$, for nonfaulty nodes. While the threat model in [34], [35] is a specialized variant of the Byzantine model, it allows the indices of the nodes under attack to change from one time instance to the next one. Both [34], [35] use similar consensus+innovations algorithms coupled with time-varying gains to achieve resilience. The main difference in these works is that [35] can achieve linear convergence, as opposed to sublinear convergence for [34], but it can only tolerate less than 30% of the nodes being under attack, in contrast to 50% for [34]. Finally, [36] accomplishes Byzantine-resilient decentralized dynamic estimation under the general linear model $\mathbf{y}_j[n] = \mathbf{H}_j \mathbf{w} + \boldsymbol{\eta}_j[n]$ by focusing on explicit detection of adversaries. We conclude by noting that the review article [2] provides more detailed overview of [33], [34], [36].

C. Decentralized Machine Learning

Decentralized machine learning algorithms, which can be considered a combination of consensus and distributed learning frameworks, approximately solve (2) by minimizing a *global* loss function on the non-located data in a decentralized manner while reaching an agreement among all nodes, i.e.,

$$\min_{\{\mathbf{w}_1, \dots, \mathbf{w}_M\}} \frac{1}{M} \sum_{j=1}^M f(\mathbf{w}_j, \mathbf{Z}_j) \quad \text{subject to} \quad \mathbf{w}_i = \mathbf{w}_j \quad \forall i, j, \quad (3)$$

where $\mathbf{w}_j \in \mathbb{R}^d$ denotes the model learned at node j in the system. We refer to (3) as the *decentralized* ERM problem, which approximately solves the statistical risk minimization problem (2).

We focus here on synchronous gradient descent-based methods for solving (3). The classic *decentralized gradient descent* (DGD) method [37], for instance, involves each node j exchanging its current local iterate in every iteration t with all nodes in its neighborhood \mathcal{N}_j and then updating the local iterate using a consensus-type weighted averaging step and a local gradient descent step, i.e.,

$$\mathbf{w}_j^{t+1} = \underbrace{\alpha_{jj} \mathbf{w}_j^t + \sum_{i \in \mathcal{N}_j} \alpha_{ji} \mathbf{w}_i^t}_{\text{consensus}} \underbrace{- \rho(t) \mathbf{g}_j(\mathbf{w}_j^t)}_{\text{local gradient descent}}, \quad (4)$$

where $\{\alpha_{ji}\}$ is the collection of averaging weights, $\rho(t)$ denotes the step size, and $\mathbf{g}_j(\mathbf{w}) := \nabla_{\mathbf{w}} f(\mathbf{w}, \mathbf{Z}_j)$ denotes the local gradient. Similar to the case of ‘‘Decentralized Consensus,’’ however, a simple lazy Byzantine strategy at one or more nodes will lead to breakdown of DGD-based learning methods.

Scalar-valued Problems: Robustification of decentralized ERM-based learning in the presence of Byzantine attacks requires resilient variants of DGD. The works in [4], [38] focus on this for scalar-valued (i.e., $d = 1$) decentralized optimization. Similar to distributed learning and decentralized consensus, resilience in these methods is also achieved through the use of a screening-based aggregation procedure within the consensus step in (4). The works [4], [38] in particular focus on trimmed-mean screening: after node j receives $w_i \in \mathbb{R}$ from its neighbors, it removes the largest b and the smallest b w_i ’s and takes average of the remaining scalars for consensus. This leads to the following update rule:

$$w_j^{t+1} = \frac{1}{|\mathcal{N}_{j^*}^t|} \sum_{i \in \mathcal{N}_{j^*}^t} w_i^t - \rho(t) g_j(w_j^t), \quad (5)$$

where $\mathcal{N}_{j^*}^t$ is the set of nodes (possibly including node j itself) that survive the screening at node j . This update can be shown to be Byzantine resilient in the sense that the w_j^t ’s converge to the minimizer of *some* convex combination of local losses $f(w, \mathbf{Z}_j)$. We also have a result from [4] that the minimum of the decentralized ERM in (3), even when restricted to the set of nonfaulty nodes, cannot be achieved in the presence of Byzantine nodes. Despite this negative result, it can be shown that any convex combination of local losses will converge in probability to the global statistical risk in the case of i.i.d. training samples. This can then be leveraged to establish the robustness of (5) for scalar-valued decentralized learning [39].

Vector-valued Problems: The algorithms in [4], [38] cannot be directly utilized in vector-valued problems (i.e., $\mathbf{w} \in \mathbb{R}^d$ for $d > 1$). On the one hand, unless a problem decouples over different coordinates of the optimization variable \mathbf{w} , minimizing the objective function along one coordinate *independent* of the other coordinates does not yield the right solution. On the other hand, since the trimmed-mean procedure of [4], [38] requires sorting of values received from one’s neighbors, it cannot be directly applied to members of an *unordered space* like \mathbb{R}^d . This limitation of [4], [38] is overcome in [39], which proposes an algorithm termed *Byzantine-resilient decentralized coordinate descent* (ByRDIE) for vector-valued decentralized learning in the presence of Byzantine nodes. ByRDIE, fundamentally being a coordinate descent method, cyclically updates one coordinate at a time in a decentralized manner. And since each subproblem in coordinate descent becomes a scalar-valued problem, ByRDIE uses trimmed-mean screening in each inner (coordinate-wise) iteration for Byzantine resilience. The final update rule in each coordinate-wise iteration of ByRDIE takes a form similar to (5), except that the gradient term also depends on other coordinates of \mathbf{w} . For strictly convex and smooth loss functions, [39] guarantees

(algorithmic and statistical) convergence of the iterates of ByRDIE to the statistical minimizer, with the algorithmic convergence rate being sublinear and the statistical learning rate being $\mathcal{O}\left(a/\sqrt{MN}\right)$; here, N again denotes the number of i.i.d. training samples per node, while the parameter a differs from one problem setup (and Byzantine attack model) to another and satisfies $1 \leq a \leq \sqrt{M}$. Despite these guarantees, which establish that ByRDIE can result in robust and (sample-wise) fast statistical learning in decentralized setups that exceeds the local learning rate of $\mathcal{O}\left(1/\sqrt{N}\right)$ [39], the one-coordinate-at-a-time update of ByRDIE can be inefficient for high-dimensional (i.e., $d \gg 1$) problems. Indeed, since the coordinate-wise gradient update step depends on the updates of other coordinates, the iterates in ByRDIE cannot be updated in a coordinate-wise parallel fashion, leading to high network coordination and local computation overheads in decentralized learning.

To curtail the high overheads of ByRDIE, [40] presents another algorithm—termed *Byzantine-resilient decentralized gradient descent* (BRIDGE)—that is based on gradient descent and coordinate-wise trimmed mean. Similar to DGD, each node j in BRIDGE also exchanges its *entire* current iterate \mathbf{w}_j^t in every iteration t with all nodes in its neighborhood \mathcal{N}_j . The update step in BRIDGE, however, involves coordinate-wise screening/aggregation of \mathbf{w}_j^t 's using trimmed mean with parameter $2b$, which is followed by a local gradient descent step. Mathematically, the (parallel) update of the k -th coordinate is given by

$$\forall k \in \{1, \dots, d\} \text{ (in parallel), } w_j^{t+1}(k) = \frac{1}{|\mathcal{N}_{j^*}^{t,k}|} \sum_{i \in \mathcal{N}_{j^*}^{t,k}} w_i^t(k) - \rho(t)g_j^k(\mathbf{w}_j^t), \quad (6)$$

where $\mathcal{N}_{j^*}^{t,k}$ is the set of nodes whose k -th coordinates survive the screening at node j and $g_j^k(\mathbf{w})$ denotes the k -th coordinate of the local gradient $\mathbf{g}_j(\mathbf{w})$. In the case of strongly convex and smooth loss functions, theoretical guarantees for BRIDGE match those for ByRDIE.

Topology Constraints: We noted earlier (see ‘‘Decentralized Consensus’’) that Byzantine resilience in decentralized setups depends on network topology. We now describe two related topology constraints for trimmed mean-based decentralized learning. The constraint in [4] requires that, after removing all b Byzantine nodes and *any* combination of remaining b (incoming) edges from each node, there is always a group of nodes—termed *source component*—of cardinality at least $(b+1)$ that has a directed path to every other node. In words, this constraint means every nonfaulty node, even after trimmed-mean screening, can always receive information—directly or indirectly—from the source component. The topology constraints for ByRDIE and BRIDGE are also based on this condition. The constraint in [38], on the other hand, is based on the idea that any two arbitrary partitions of the network must result in one of the partitions having at least one node with $(2b+1)$ neighbors outside the partition. This, in turn, guarantees that b Byzantine nodes cannot isolate any subset of nonfaulty nodes during trimmed-mean screening.

Beyond Trimmed-mean Screening: Unlike distributed learning, where several screening methods and aggregation rules have been proposed and analyzed for Byzantine resilience (Table II), utilization of screening methods in decentralized learning has been limited to (coordinate-wise) trimmed mean. The main reason for this is the need for iterate consensus in decentralized learning, which makes analysis of other screening/aggregation methods challenging. (It is also worth noting that our focus here is on iterate screening, as opposed to gradient screening in distributed learning.) In practice, it is possible to merge the ideas behind BRIDGE and screening methods such as coordinate-wise median, Krum, and Bulyan to develop other Byzantine-resilient decentralized learning methods. This is indeed the approach we have taken in the numerical experiments discussed in the following. Consensus, convergence analysis, and statistical learning rates of such methods, however, remain an open problem.

Numerical Experiments: We first compare the performance of DGD, ByRDIE, BRIDGE and three variants of BRIDGE that come about from incorporation of the screening principles of (coordinate-wise) median, Krum, and Bulyan (with Krum) within the BRIDGE framework. The computational inefficiency of ByRDIE for high-dimensional datasets and machine learning models necessitates an experimental setup, both in terms of the dataset and the learning task, that differs from the one utilized as part of our discussion on Byzantine-resilient distributed machine learning. Specifically, note that each (colored) image in the CIFAR-10 dataset is 3,072-dimensional and the corresponding convolutional neural network used earlier in our discussion for this dataset gives rise to a model with $d = 122,410$. In contrast, we resort to a computationally tractable experimental setup in here that corresponds to multiclass classification of MNIST dataset (784-dimensional data samples) using a linear one-layer neural network (i.e., $d = 7,840$). In addition to being computationally manageable for ByRDIE, this setup also results in a loss function that satisfies the theoretical conditions for convergence of ByRDIE and BRIDGE.

The decentralized system in our experimental setup involves a total of $M = 20$ nodes in the network, with a communications link (edge) between two nodes decided by a random coin flip. Once a random topology is generated, we ensure each node has at least $4b + 1$ nodes in its neighborhood (a condition imposed due to Bulyan screening). The training data at each node corresponds to $N = 2,000$ samples randomly selected from the MNIST dataset. The performance of each method is reported in terms of classification accuracy, averaged over $(M - b)$ nodes and a total of 10 independent Monte Carlo trials, as a function of the number of scalars broadcast per node. The final results, shown in Figure 6, correspond to two sets of experiments: (i) the faultless setting in which none of the nodes actually behaves maliciously; and (ii) the setting in which two of the 20 nodes are Byzantine, with each Byzantine node broadcasting every coordinate of the iterate as a uniform random variable between -1 and 0 . Note that this Byzantine




Fig. 6. Performance comparison of different decentralized learning methods in both faultless (left panel) and Byzantine settings (right panel). Byzantine-resilient algorithms in both settings operate under the assumption of $b = 2$. The algorithms entitled *Median*, *Krum*, and *Bulyan* are effectively BRIDGE combined with the screening procedures advocated in distributed learning.

attack strategy is by no means the most potent in all decentralized settings. However, similar to the attack strategy utilized in our earlier discussion on distributed learning, this particular strategy has been selected after careful evaluation of the impact of different strategies proposed in works such as [12]–[14] on our particular experimental setup. Finally, with the exception of DGD, all methods are initialized with parameter $b = 2$ in both faultless and faulty scenarios.

It can be seen from Figure 6 that, other than ByRDIE and Krum-based screening, all methods perform almost as well as DGD in the faultless case. In the presence of Byzantine nodes, however, DGD completely falls apart, whereas the performances of all screening methods remain comparable to the faultless setting. It is also worth comparing these results to those for Byzantine-resilient distributed learning (see Figure 5). While (coordinate-wise) median and trimmed mean appear to be the worst performers in Figure 5, Krum-based screening is the least effective in Figure 6. In both cases, however, Bulyan is quite effective, except that it has stringent topology requirements.

We conclude by explicitly comparing the performance of DGD and BRIDGE in decentralized settings to that of (Byzantine-resilient) learning methods in distributed settings. Since DGD and BRIDGE are scalable to high-dimensional learning tasks, our experimental setup, dataset, data distribution, learning task, and Byzantine attack strategy for this comparison are identical to the ones described earlier for Byzantine-resilient distributed learning methods. In order to ensure the decentralized setup in the case of BRIDGE satisfies the topology constraints corresponding to four Byzantine nodes in the system, we use 0.7 as the probability of random connectivity between any two nodes. The final results for DGD and BRIDGE,

which are overlaid on top of the ones for distributed learning methods in Figure 5, show that BRIDGE offers competitive performance in both faultless and faulty systems. In fact, BRIDGE (and DGD) have faster convergence rates than the distributed learning methods for this particular nonconvex problem, decentralized setup, and Byzantine strategy. While this could be partly attributable to higher network connectivity of 0.7 in the decentralized setting, a careful head-to-head comparison and understanding of distributed and decentralized learning methods in both faultless and faulty settings is an open problem.

III. SOME OPEN RESEARCH PROBLEMS

Despite recent advances, Byzantine-resilient inference and learning remains an active area of research with several open problems. Much of the focus in distributed inference has been on the somewhat restrictive model in which Byzantine nodes do not collude. Collaborative Byzantine attacks, on the other hand, can be much more potent than independent ones. A fundamental understanding of mechanisms for safeguarding against such attacks remains a relatively open problem in distributed inference. Byzantine-resilient distributed estimation under nonlinear models is another problem that has been relatively unexplored. In the case of Byzantine-resilient distributed learning, existing works have only scratched the surface. Convergence and/or learning rates of many of the proposed methods remain unknown (Table II). In addition, while SGD is a workhorse of machine learning, approaches such as accelerated first-order methods (e.g., accelerated gradient descent), first-order dual methods (e.g., ADMM), and second-order methods (e.g., Newton’s method) do play important roles in machine learning. However, resilience of distributed variants of such methods to Byzantine attacks has not been investigated in the literature.

The lack of a central server, the need for consensus, and an ad-hoc topology make it even more challenging to develop and analyze Byzantine-resilient methods for decentralized inference and learning. Much of the work in this regard is based on screening methods such as trimmed mean and median that originated in the literature on Byzantine-resilient scalar-valued consensus. This has left open the question of how other screening methods, such as the ones explored within distributed learning, might handle Byzantine attacks—both in theory and in practice—in various decentralized problems. Unlike distributed learning, any such efforts will also have to characterize the interplay between network topology and effectiveness of the screening procedure. The fundamental tradeoffs between the robustness and the (faultless) performance of Byzantine-resilient methods also remain largely unknown for decentralized setups. Finally, existing works on decentralized learning only guarantee sublinear convergence for strictly/strongly convex and smooth functions. Whether this can be improved by taking advantage of faster distributed optimization frameworks or different screening methods also remains an open question.

IV. CONCLUSIONS

In this article, we have presented an overview of latest advances in Byzantine-resilient inference and learning. In the distributed master–worker setting, which is characterized by the presence of a central server that computes the final solution, we discussed recent results concerning resilience of distributed detection, estimation, and machine learning against Byzantine attacks. Within distributed machine learning, we focused on Byzantine-resilient variants of distributed stochastic gradient descent, whose performances were compared using numerical experiments. In the decentralized setting, which typically requires consensus due to lack of a central server, we first discussed the principles behind Byzantine-resilient consensus. This was followed by a discussion of latest results on decentralized detection, estimation, and learning in the presence of Byzantine nodes. We also compared and contrasted different Byzantine-resilient decentralized learning methods using numerical experiments, and discussed similarities and differences between them and distributed learning methods. Byzantine-resilient inference and learning has a number of research challenges that remain unaddressed, some of which are also briefly discussed in the article.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of the NSF (CCF-1453073, CCF-1907658), the ARO (W911NF-17-1-0546), and the DARPA Lagrange Program (ONR/SPAWAR contract N660011824020).

SHORT BIOGRAPHIES

Zhixiong Yang (zhixing.yang@rutgers.edu) received BS degree from Beijing Jiaotong University, China in 2011, MS degree from Northeastern University, MA in 2013, and PhD in electrical engineering from Rutgers University–New Brunswick, NJ in 2020. He is currently a system engineer at Blue Danube Systems. His research interests are distributed processing, machine learning, and massive MIMO systems.

Arpita Gang (arpita.gang@rutgers.edu) received BTech degree from the National Institute Of Technology, Silchar, India, and MTech degree from IIT-Delhi, India. Since 2017, she has been pursuing PhD in electrical engineering from Rutgers University–New Brunswick, NJ. Her research interests are distributed optimization and signal processing. She is a graduate student member of the IEEE.

Waheed U. Bajwa (waheed.bajwa@rutgers.edu) has been with Rutgers University–New Brunswick, NJ since 2011, where he is currently an associate professor in the Departments of Electrical & Computer Engineering and Statistics. His research interests include statistical signal processing, high-dimensional statistics, and machine learning. He is a senior member of the IEEE.

REFERENCES

- [1] A. Vempaty, L. Tong, and P. Varshney, “Distributed inference with Byzantine data: State-of-the-art review on data falsification attacks,” *IEEE Signal Process. Mag.*, vol. 30, no. 5, pp. 65–75, 2013.
- [2] Y. Chen, S. Kar, and J. M. F. Moura, “The Internet of Things: Secure distributed inference,” *IEEE Signal Process. Mag.*, vol. 35, no. 5, pp. 64–75, 2018.
- [3] J. Zhang, R. S. Blum, and H. V. Poor, “Approaches to secure inference in the Internet of Things: Performance bounds, algorithms, and effective attacks on IoT sensor networks,” *IEEE Signal Process. Mag.*, vol. 35, no. 5, pp. 50–63, 2018.
- [4] L. Su and N. Vaidya, “Fault-tolerant distributed optimization (Part IV): Constrained optimization with arbitrary directed networks,” *arXiv preprint arXiv:1511.01821*, 2015.
- [5] L. Lamport, R. Shostak, and M. Pease, “The Byzantine generals problem,” *ACM Trans. Programming Languages and Syst.*, vol. 4, no. 3, pp. 382–401, 1982.
- [6] V. S. S. Nadendla, Y. S. Han, and P. K. Varshney, “Distributed inference with M-ary quantized data in the presence of Byzantine attacks,” *IEEE Trans. Signal Process.*, vol. 62, no. 10, pp. 2681–2695, 2014.
- [7] B. Kailkhura, Y. S. Han, S. Brahma, and P. K. Varshney, “Distributed Bayesian detection in the presence of Byzantine data,” *IEEE Trans. Signal Process.*, vol. 63, no. 19, pp. 5250–5263, 2015.
- [8] W. Hashlamoun, S. Brahma, and P. K. Varshney, “Audit bit based distributed Bayesian detection in the presence of Byzantines,” *IEEE Trans. Signal Inform. Proc. over Netw.*, vol. 4, no. 4, pp. 643–655, 2018.
- [9] B. Kailkhura, S. Brahma, B. Dulek, Y. S. Han, and P. K. Varshney, “Distributed detection in tree networks: Byzantines and mitigation techniques,” *IEEE Trans. Inf. Forens. Security*, vol. 10, no. 7, pp. 1499–1512, 2015.
- [10] J. Haupt, W. U. Bajwa, M. Rabbat, and R. Nowak, “Compressed sensing for networked data,” *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 92–101, 2008.
- [11] A. Vempaty, O. Ozdemir, K. Agrawal, H. Chen, and P. K. Varshney, “Localization in wireless sensor networks: Byzantines and mitigation techniques,” *IEEE Trans. Signal Process.*, vol. 61, no. 6, pp. 1495–1508, 2013.
- [12] C. Xie, S. Koyejo, and I. Gupta, “Fall of empires: Breaking Byzantine-tolerant SGD by inner product manipulation,” *arXiv preprint arXiv:1903.03936*, 2019.
- [13] M. Fang, X. Cao, J. Jia, and N. Z. Gong, “Local model poisoning attacks to Byzantine-robust federated learning,” *arXiv preprint arXiv:1911.11815*, 2019.
- [14] G. Baruch, M. Baruch, and Y. Goldberg, “A little is enough: Circumventing defenses for distributed learning,” in *Proc. Advances in Neural Inf. Process. Syst. (NeurIPS’19)*, 2019, pp. 8632–8642.
- [15] D. Yin, Y. Chen, K. Ramchandran, and P. Bartlett, “Byzantine-robust distributed learning: Towards optimal statistical rates,” in *Proc. 35th Int. Conf. Machine Learning (ICML’18)*, 2018, vol. 80, pp. 5650–5659.
- [16] Y. Chen, L. Su, and J. Xu, “Distributed statistical machine learning in adversarial settings: Byzantine gradient descent,” in *Proc. ACM Measurement and Analysis of Computing Systems*, Dec. 2017, vol. 1, pp. 44:1–44:25.
- [17] P. Blanchard, R. Guerraoui, and J. Stainer, “Machine learning with adversaries: Byzantine tolerant gradient descent,” in *Proc. Advances in Neural Inf. Process. Syst. (NeurIPS’17)*, 2017, pp. 118–128.
- [18] E. El-Mhamdi, R. Guerraoui, and S. Rouault, “The hidden vulnerability of distributed learning in Byzantium,” in *Proc. 35th Int. Conf. Machine Learning (ICML’18)*, 2018, vol. 80, pp. 3521–3530.
- [19] E. El-Mhamdi and R. Guerraoui, “Fast and secure distributed learning in high dimension,” *arXiv preprint arXiv:1905.04374*, 2019.
- [20] C. Xie, O. Koyejo, and I. Gupta, “Zeno: Byzantine-suspicious stochastic gradient descent,” *arXiv preprint arXiv:1805.10032*, 2018.

- [21] C. Xie, O. Koyejo, and I. Gupta, “Zeno++: Robust asynchronous SGD with arbitrary number of Byzantine workers,” *arXiv preprint arXiv:1903.07020*, 2019.
- [22] C. Xie, O. Koyejo, and I. Gupta, “Generalized Byzantine-tolerant SGD,” *arXiv preprint arXiv:1802.10116*, 2018.
- [23] L. Li, W. Xu, T. Chen, G. Giannakis, and Q. Ling, “RSA: Byzantine-robust stochastic aggregation methods for distributed learning from heterogeneous datasets,” in *Proc. AAAI Conf. on Artificial Intelligence (AAAI’19)*, 2019, vol. 33, pp. 1544–1551.
- [24] J. Bernstein, J. Zhao, K. Azizzadenesheli, and A. Anandkumar, “signSGD with majority vote is communication efficient and Byzantine fault tolerant,” *arXiv preprint arXiv:1810.05291*, 2018.
- [25] D. Alistarh, Z. Allen-Zhu, and J. Li, “Byzantine stochastic gradient descent,” in *Proc. Advances in Neural Inf. Process. Syst. (NeurIPS’18)*, 2018, pp. 4618–4628.
- [26] H. J. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram, “Resilient asymptotic consensus in robust networks,” *IEEE J. Sel. Areas Commun.*, vol. 31, no. 4, pp. 766–781, 2013.
- [27] N. Vaidya, “Matrix representation of iterative approximate Byzantine consensus in directed graphs,” *arXiv preprint arXiv:1203.1888*, 2012.
- [28] L. Haseltalab and M. Akar, “Approximate Byzantine consensus in faulty asynchronous networks,” in *Proc. American Control Conf. (ACC’15)*. IEEE, 2015, pp. 1591–1596.
- [29] D. Silvestre, P. Rosa, J. P. Hespanha, and C. Silvestre, “Stochastic and deterministic fault detection for randomized gossip algorithms,” *Automatica*, vol. 78, pp. 46–60, 2017.
- [30] S. Liu, H. Zhu, S. Li, C. Chen, and X. Guan, “An adaptive deviation-tolerant secure scheme for distributed cooperative spectrum sensing,” in *Proc. IEEE Global Commun. Conf. (GLOBECOM’12)*, 2012, pp. 603–608.
- [31] B. Kailkhura, P. Ray, D. Rajan, A. Yen, P. Barnes, and R. Goldhahn, “Byzantine-resilient locally optimum detection using collaborative autonomous networks,” in *Proc. IEEE 7th Int. Workshop Computational Advances Multi-Sensor Adaptive Process. (CAMSAP’17)*, 2017, pp. 1–5.
- [32] B. Kailkhura, S. Brahma, and P. K. Varshney, “Data falsification attacks on consensus-based detection systems,” *IEEE Trans. Signal Inform. Proc. over Netw.*, vol. 3, no. 1, pp. 145–158, 2017.
- [33] H. J. LeBlanc and F. Hassan, “Resilient distributed parameter estimation in heterogeneous time-varying networks,” in *Proc. 3rd Int. Conf. High Confidence Networked Syst. (HiCoNS’14)*, 2014, pp. 19–28.
- [34] Y. Chen, S. Kar, and J. M. F. Moura, “Resilient distributed estimation: Sensor attacks,” *IEEE Trans. Automat. Control*, vol. 64, no. 9, pp. 3772–3779, 2018.
- [35] Y. Chen, S. Kar, and J. M. F. Moura, “Resilient distributed estimation: Exponential convergence under sensor attacks,” in *Proc. IEEE Conf. Decision Control (CDC’18)*, 2018, pp. 7275–7282.
- [36] Y. Chen, S. Kar, and J. M. F. Moura, “Resilient distributed estimation through adversary detection,” *IEEE Trans. Signal Process.*, vol. 66, no. 9, pp. 2455–2469, 2018.
- [37] A. Nedić and A. Ozdaglar, “Distributed subgradient methods for multi-agent optimization,” *IEEE Trans. Automat. Control*, vol. 54, no. 1, pp. 48–61, 2009.
- [38] S. Sundaram and B. Ghahserifard, “Distributed optimization under adversarial nodes,” *IEEE Trans. Automat. Control*, vol. 64, no. 3, pp. 1063–1076, 2019.
- [39] Z. Yang and W. U. Bajwa, “ByRDIE: Byzantine-resilient distributed coordinate descent for decentralized learning,” *IEEE Trans. Signal Inform. Proc. over Netw.*, vol. 5, no. 4, pp. 611–627, 2019.
- [40] Z. Yang and W. U. Bajwa, “BRIDGE: Byzantine-resilient decentralized gradient descent,” *arXiv preprint arXiv:1908.08098*, 2019.