Zeno: Distributed Stochastic Gradient Descent with Suspicion-based Fault-tolerance

Cong Xie, Oluwasanmi Koyejo, Indranil Gupta

June 12, 2019

Poster *158*

Security in Distributed ML

Zeno: distributed synchronous SGD that

- tolerates an arbitrary number of malicious workers
- provides convergence guarantees for non-convex problems

Goal: converge under attacks/failures, regardless of false negative

	Prev.	Ours
Tolerates a majority of malicious workers	No	**Yes**
Considers the progress of optimization	No	**Yes**
Tolerates stealth adversary (empirically)	No	**Yes**

Byzantine-tolerant SGD

m workers, distributed SGD:

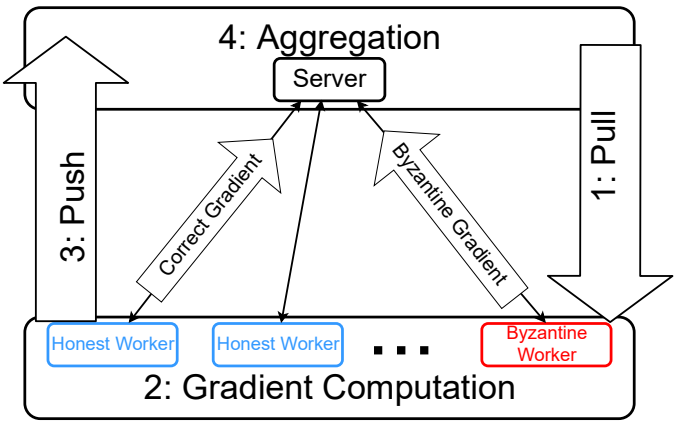

Poster *158*

Main Idea & Results

⋆ Sort $g_i(x), i \in [m]$ by the Stochastic descent score:

Definition

Stochastic descent score of any update u:

$$Score_{\gamma,\rho}(u, x) = f_r(x) - f_r(x - \gamma u) - \rho\|u\|^2,$$

$f_r(x)$: unbiased estimator of the loss $F(x)$, for validation.

⋆ Zeno: filter the b "worst" gradients $\frac{1}{m-b}\sum_{i=1}^{m-b} \tilde{v}_{(i)}$, $b > q$.

⋆ Convergence after T iterations:

$$\frac{\sum_{t=0}^{T-1} \mathbb{E}\|\nabla F(x^t)\|^2}{T} \leq \mathcal{O}\left(\frac{1}{\sqrt{T}}\right) + \mathcal{O}\left(\frac{(b-q+1)(m-q)}{(m-b)^2}\right).$$