Random Walk Distributed Dual Averaging Algorithm NEC

r

Cun Mu*, Asim Kadav, Erik Kruus, Donald Goldfarb* and Martin Rengiang Min NEC Labs, Princeton Columbia University*

Peer-to-peer parallel learning

Parallel learning over peer nodes without a central coordinating resource.

Performed over a network of sensors, mobile devices or geographically separated data centers.



Network is robust to failure due to lack of a central coordinating resource

Challenges in peer-to-peer learning

- Non-iid data: The peer replicas may process non independent and identically distributed datasets. For example, physical sensors or data centers may collect biased samples.
- Random link failures: There may be intermittent link failures that may affect convergence for large scale learning

Goal: Provide a distributed consensus algorithm that provably converges in presence of non-iid data and link failures

min
$$f(\boldsymbol{x}) = rac{1}{n} \sum_{i=1}^n f_i(\boldsymbol{x})$$
 s.t. $\boldsymbol{x} \in \mathcal{X} \subseteq \mathbb{R}^d$

