Robust Distributed Estimation: Extending Gossip Algorithms to Ranking and Trimmed Means

Can we design more robust gossip algorithms?

Anna van Elst, Igor Colin, Stephan Clémençon

LTCI, Télécom Paris, Institut Polytechnique de Paris

Towards More Robust Gossip Algorithms

- In [1], we introduce and analyze synchronous gossip algorithms for:
  - Rank estimation (GoRank)
  - Trimmed mean estimation (GoTrim)

- In [2], we extend:
  - Analysis to the asynchronous setting
  - GoTrim to rank-based statistics

[1] A. van Elst, I. Colin, S. Clémençon. *Robust Distributed Estimation: Extending Gossip Algorithms to Ranking and Trimmed Means*. NeurIPS, 2025.

[2] A. van Elst, I. Colin, S. Clémençon. *Asynchronous Gossip Algorithms for Rank-Based Statistical Methods*. FLTA, 2025.

Gossip Algorithms: Where and Why?

- When data is distributed over a network.
- If we have resource and communication constraints.
- If we don't have access to a central server.

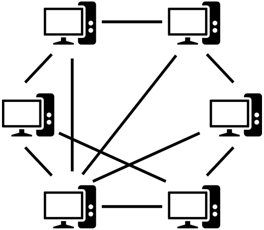

Peer-to-Peer (P2P) Networks

Wireless Sensor Networks

. . . and many more!

Simple Example: Gossip for Averaging

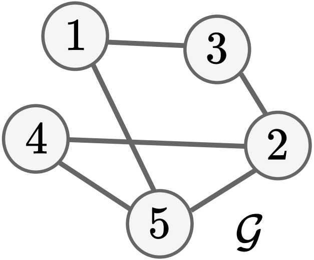

True mean:
$\bar{x} = 3$

Select random $(i, j) \in E$, average X_i and X_j

Simple Example: Gossip for Averaging

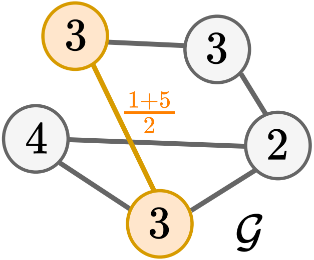

True mean:
$\bar{x} = 3$

Select random $(i, j) \in E$, average X_i and X_j

Simple Example: Gossip for Averaging

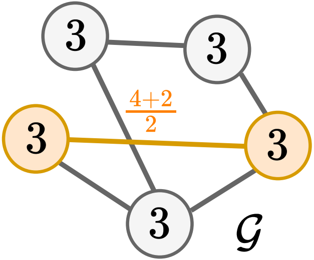

True mean:
$\bar{x} = 3$

Select random $(i, j) \in E$, average X_i and X_j

What Makes the Theoretical Analysis Work?

- Averaging matrix: $t > 0$, for a random edge $(i, j) \in E$,

$$\mathbf{W}_2(t) = \mathbf{I}_n - \frac{1}{2}\left(\mathbf{e}_i - \mathbf{e}_j\right)\left(\mathbf{e}_i - \mathbf{e}_j\right)^\top .$$

- Its expectation with respect to the edge sampling process:

$$\mathbf{W}_2 = \mathbb{E}\left[\mathbf{W}_2(t)\right] = \mathbf{I}_n - \boxed{\mathbf{L}}/(2|E|).$$

Graph Laplacian

- Properties: symmetric, doubly stochastic and

$$\left\|\mathbf{W}_2 - \mathbf{1}_n\mathbf{1}_n^\top/n\right\|_{\mathrm{op}} \leq 1 - \boxed{c}/2 < 1.$$

Graph's connectivity

- Convergence in expectation:

$$\|\mathbb{E}[\mathbf{Y}(t)] - \bar{x}_n\mathbf{1}_n\| \leq \left(1 - \frac{c}{2}\right)^t \|\mathbf{X} - \bar{x}_n\mathbf{1}_n\| .$$

Applications of Gossip for Decentralized Learning

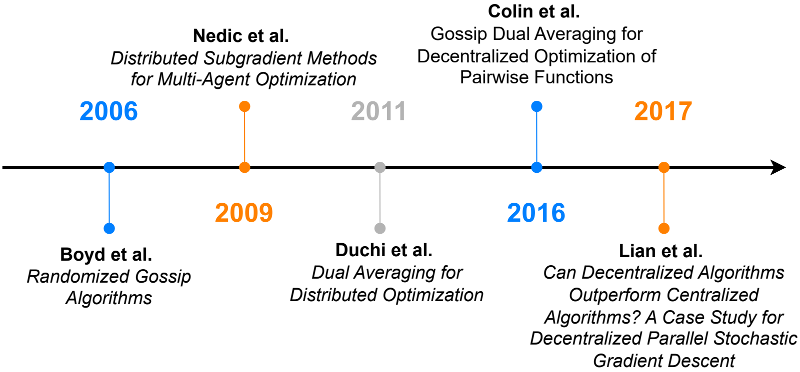

When Gossip Goes Wrong...

- What if a node's value is corrupted? All estimates will be corrupted!
- The average is not robust to outliers. We need more robust statistics, e.g., trimmed mean (see Huber's robustness).

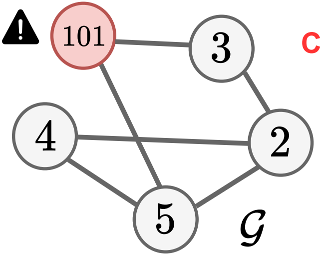

Corrupted mean:
$\bar{x} = 23$

True mean:
$\bar{x} = 3$

Robust Statistics: Trimmed Means

- Robust average: denoting $m = \lfloor n\alpha \rfloor$,

Order statistic

$$\underbrace{\bar{x}_\alpha}_{\alpha\text{-trimmed mean}} = \frac{1}{n-2m} \sum_{k=m+1}^{n-m} X_{n(k)}$$

- Breakdown point is equal to α.

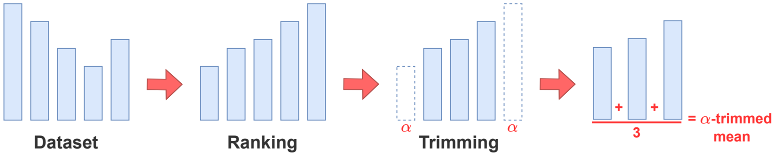

Problem Formulation and Framework

- Communication network: connected undirected graph $G = (V, E)$ with n nodes. Each node k holds X_k.
- Huber's contamination model: fraction ε of corrupted nodes.
- Synchronous/asynchronous setting: global/local iteration counter.

Using ranks, we can compute a robust average:

$$\underbrace{\bar{x}_\alpha}_{\alpha\text{-trimmed mean}} = \frac{1}{n-2m} \sum_{k=m+1}^{n-m} X_{n(k)}$$

$$= \frac{1}{n} \sum_k f\left(\underbrace{r_k}_{\text{Rank}} \right) \cdot \underbrace{X_k}_{\text{Observation}}$$

GoRank – A Gossip Algorithm for Ranking (Asynchronous)

$$\underbrace{r_k}_{\text{Rank of node } k} = 1 + n \left(\frac{1}{\underbrace{n}_{\#\text{ of nodes}}} \sum_{l=1}^{n} \mathbb{I}_{\{\underbrace{X_k > X_l}_{\text{Pairwise comparison}}\}} \right) = 1 + n r'_k$$

Estimates:

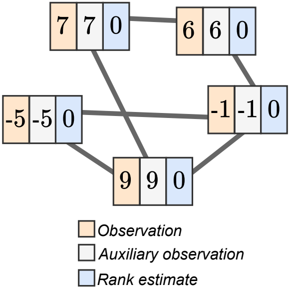

True ranks:

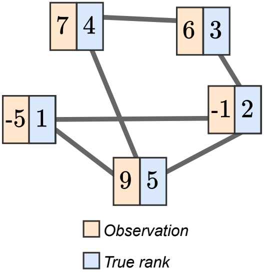

GoRank – A Gossip Algorithm for Ranking (Asynchronous)

$$\underbrace{\boxed{R_k(t)}}_{\text{Rank estimate}} = 1 + n\left(\frac{1}{\underbrace{\boxed{C_k(t)}}_{\text{Counter}}}\sum_{s=1}^{C_k(t)} \mathbb{I}_{\{X_k > \underbrace{\boxed{Y_k(s)}}_{\text{Auxiliary observation}}\}}\right) = 1 + nR'_k(t)$$

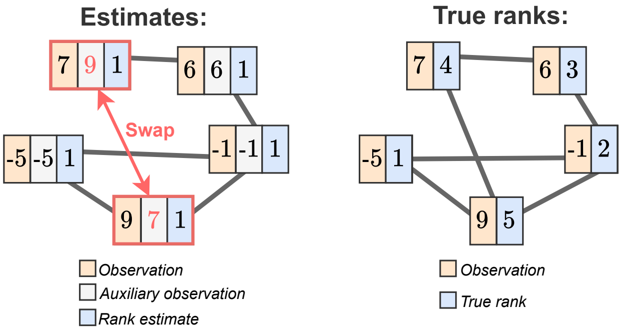

$$\boxed{R_k(t)} = 1 + n\left(\frac{1}{\boxed{C_k(t)}}\sum_{s=1}^{C_k(t)} \mathbb{I}_{\{X_k > \boxed{Y_k(s)}\}}\right) = 1 + nR'_k(t)$$

Rank estimate

Counter

Auxiliary observation

Estimates:

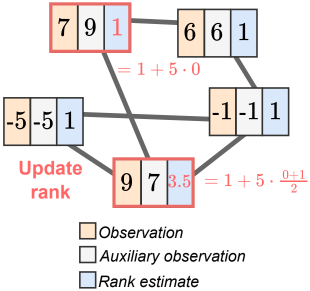

True ranks:

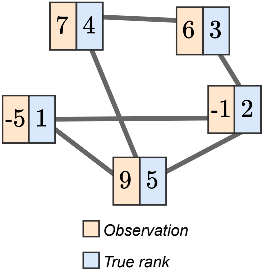

GoRank – A Gossip Algorithm for Ranking (Asynchronous)

$$\boxed{R_k(t)} = 1 + n\left(\frac{1}{\boxed{C_k(t)}}\sum_{s=1}^{C_k(t)} \mathbb{I}_{\{X_k > \boxed{Y_k(s)}\}}\right) = 1 + nR'_k(t)$$

Rank estimate

Counter

Auxiliary observation

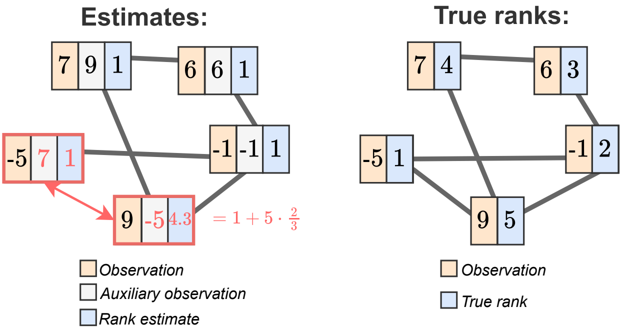

GoRank – A Gossip Algorithm for Ranking (Asynchronous)

$$\boxed{R_k(t)} = 1 + n\left(\frac{1}{\boxed{C_k(t)}}\sum_{s=1}^{C_k(t)} \mathbb{I}_{\{X_k > \boxed{Y_k(s)}\}}\right) = 1 + nR'_k(t)$$

Rank estimate

Counter

Auxiliary observation

Estimates:

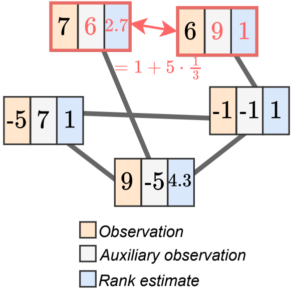

True ranks:

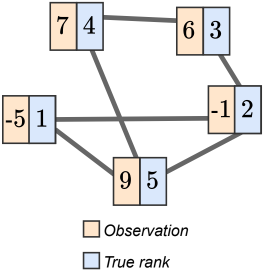

$$R_k(t) = 1 + n\left(\frac{1}{C_k(t)}\sum_{s=1}^{C_k(t)} \mathbb{I}_{\{X_k > Y_k(s)\}}\right) = 1 + nR'_k(t)$$

Rank estimate

Counter

Auxiliary observation

Estimates:

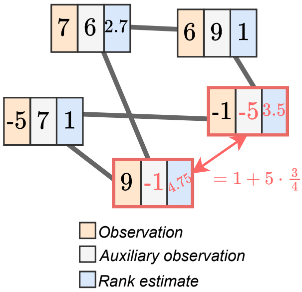

True ranks:

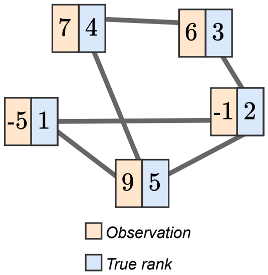

GoRank Shows Great Empirical Performance

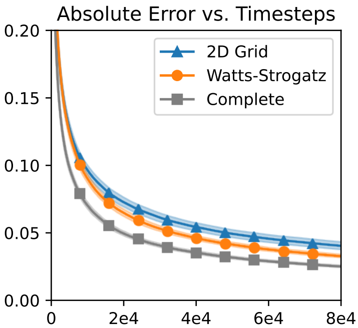

Figure 1: Graph topology has an impact on convergence.

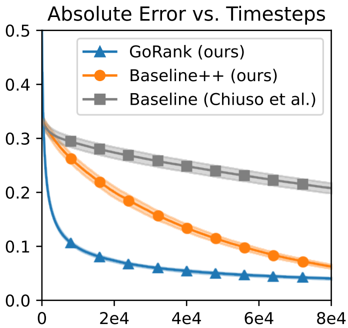

Figure 2: GoRank outperforms previous work in large graphs.

GoRank: Theoretical Analysis

	Synchronous	Asynchronous
Estimator	$\frac{1}{t}\sum_{s=1}^{t}\mathbb{I}_{\{X_k > Y_k(s)\}}$	$\frac{1}{C_k(t)}\sum_{s=1}^{t}\delta_k(s)\mathbb{I}_{\{X_k > Y_k(s)\}}$
Expected Abs. Error	$\mathcal{O}\left(\frac{\sigma}{\sqrt{ct}}\right)$	$\mathcal{O}\left(\frac{\sigma}{\sqrt{ct}}\right)$

Table 1: Convergence Rate Bounds.

σ = standard deviation

c = spectral gap of the Laplacian

"The absolute error converges in $\mathcal{O}\left(\frac{1}{\sqrt{t}}\right)$."

GoRank: What Makes the Theoretical Analysis Work?

- Swapping matrix: $t > 0$, for a random edge $(i, j) \in E$,

$$\mathbf{W}_1(t) = \mathbf{I}_n - (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top .$$

- Its expectation with respect to the edge sampling process:

$$\mathbf{W}_1 = \mathbb{E}\left[\mathbf{W}_1(t)\right] = \mathbf{I}_n - \boxed{\mathbf{L}}/|E|.$$

Graph Laplacian

- Properties: symmetric, doubly stochastic and

$$\left\|\mathbf{W}_1 - \mathbf{1}_n\mathbf{1}_n^\top/n\right\|_{\text{op}} \leq 1 - \boxed{c} < 1.$$

Graph's connectivity

- Computes a weighted average using current ranks:

$$\frac{1}{n}\sum_{k} f(r_k) \cdot X_k.$$

- **Idea 1**: Perform standard (weighted) gossip-based averaging using current ranks:

$$\text{weight} = f(r_k).$$

- **Idea 2**: Since ranks vary, correct past errors by injecting:

$$(\text{weight}^{\text{new}} - \text{weight}^{\text{prev}}) \cdot X_k \; .$$

GoTrim Shows Great Empirical Performance

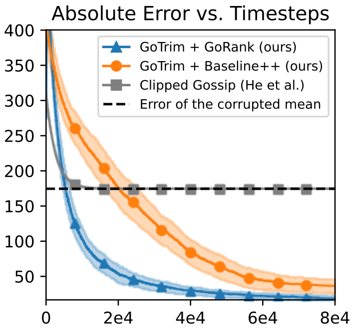

Figure 3: Simulated Dataset

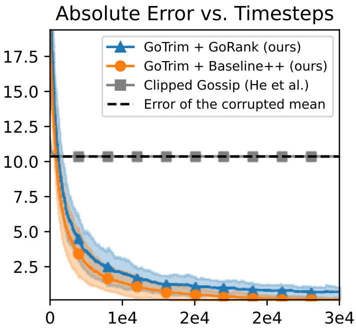

Figure 4: Basel Luftklima Dataset

"GoTrim quickly improves over the corrupted mean."

GoTrim Shows Great Empirical Performance

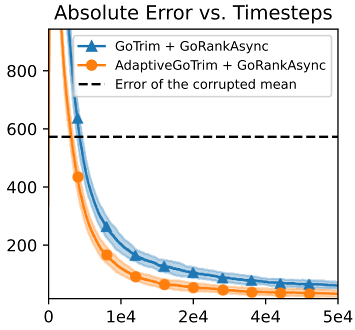

Figure 5: Trimmed Mean

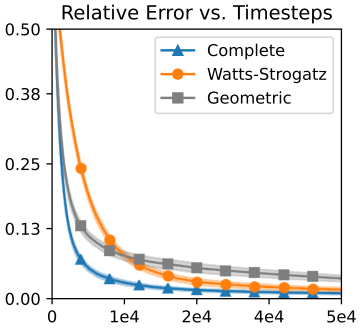

Figure 6: Wilcoxon Statistic

"GoTrim can be accelerated with a bias correction term and performs well when extended to Wilcoxon statistic estimation."

GoTrim: Theoretical Analysis

	GoTrim (synchronous)
Trimmed Mean Estimator of Node k	$\frac{1}{n}\sum_{k=1}^{n} f(R_k) \cdot X_k$
Convergence in Expectation (Theorem 3)	$\mathcal{O}\left(\frac{1}{c^2 t}\right)$
τ-Breakdown Point (Theorem 4)	$\approx \alpha - \frac{1}{\sqrt{t-T}}$

Table 2: Convergence Rate Bounds.

c = spectral gap of the Laplacian

"The bias converges in $\mathcal{O}\left(\frac{1}{t}\right)$."

Conclusion and Future Work

- We design novel gossip algorithms for ranking (GoRank) and trimmed means estimation (GoTrim).
- We extended our analysis to the asynchronous setting and our algorithm to rank-based statistics.
- We show strong empirical performance and have solid theoretical guarantees.
- **Future work**:
  - Extension to multivariate data
  - Decentralized robust optimization
  - Byzantine Robustness

References

[1] Peter J Huber and Elvezio M Ronchetti. **Robust statistics.** John Wiley & Sons, 2011.

[2] Stephen Boyd et al. **"Randomized gossip algorithms".** In: *IEEE transactions on information theory* 52.6 (2006), pp. 2508–2530.

[3] Igor Colin et al. **"Extending gossip algorithms to distributed estimation of U-statistics".** In: *NeurIPS* 28 (2015).

[4] Alessandro Chiuso et al. **"Gossip algorithms for distributed ranking".** In: *Proceedings of the 2011 American Control Conference*. IEEE. 2011, pp. 5468–5473.

[5] Lie He, Sai Praneeth Karimireddy, and Martin Jaggi. **"Byzantine-robust decentralized learning via clippedgossip".** In: *arXiv preprint arXiv:2202.01545* (2022).