DIFFERENTIAL PRIVACY FOR NETWORK CONNECTEDNESS INDICES

Tom A. Rutter, Yuxin Liu, M. Amin Rahimian

Stanford University, University of Pittsburgh

Social scientists interested in a range of social phenomena are increasingly using data on the *networks* in which individuals, firms, financial institutions and nations are embedded. The use of network data has provided valuable insights on the spread of contagious diseases, the role of peers and mentors in shaping people's life outcomes, the resilience of global supply chains, and the nature of systemic risk in the financial system. The use of such data, however, comes with concerns about the privacy protections provided to individuals and organizations in the dataset. The naive publication of aggregated statistics from these datasets entails risks that information may be leaked about the presence (or lack thereof) of particular connections in the dataset, or even about the individual characteristics of people or businesses represented in the data.

In this paper, we provide an approach to releasing a particular class of network statistics with a formal privacy guarantee that maintains high accuracy even for relatively small realistic networks. The statistics we consider are *network connectedness indices*, which describe the proportion of connections from nodes with a given characteristic that extend to nodes with either the same characteristic or a different characteristic. For example, what proportion of the friends of white individuals are also white? What proportion of the suppliers of US firms are based in China? What proportion of bank loans are directed to hedge funds?

The application we focus on, which motivated this paper, concerns the proportion of friendships that cut across the income distribution. In a widely cited pair of studies, Chetty et al. 2022a and Chetty et al. 2022b publicly released a range of social capital indices constructed from Facebook and Instagram data for counties, zip codes, universities, and schools in the US. In particular, Chetty et al. 2022a released a measure of the fraction of high-socioeconomic status (SES) friends among low-SES individuals in an area, which they single out as the strongest predictor of economic mobility in the US. In this abstract, we focus on *cross-type connectedness indices*: what proportion of connections from nodes in a set $\mathcal{A}$ extend to nodes in a set $\mathcal{B}$?

Privacy model and mechanism. We adopt **edge-adjacent differential privacy for labeled networks**: two labeled graphs are adjacent if their edge sets differ by at most one edge and at most one vertex's label differs (Blocki et al. 2013). This captures the disclosure risks most salient to public releases of connectedness—protecting both relationship information and a sensitive node attribute—while treating the vertex set as non-private. Our key design choice is to privatize node attributes first, and then compute connectedness from privatized labels, followed by an edge-private noise layer. Concretely:

- *Label privacy (ε_l).* Apply randomized response to each node label. Because this attenuates cross-group frequencies, we apply an analytic debiasing correction (a post-processing step that preserves privacy).
- *Edge privacy (ε_e).* Add Laplace noise to the final connectedness estimate scaled to its sensitivity to single-edge changes. For the key aggregate terms in the estimator, a single edge can affect at most two summands.

We show that the composition of an ε_l-DP mechanism for node characteristics and an ε_e-edge-private mechanism for the statistics of interest satisfies $(\varepsilon_l + \varepsilon_e)$-edge-adjacent-DP. This is crucial to allow us to bypass global sensitivity, which for many connectedness indices can span the entire range $[0, 1]$. A simple star-network example shows that two adjacent labeled graphs (differing in a single tie and a single label) can produce connectedness estimates at opposite extremes. As a result, a naive sensitivity calibration that adds Laplace noise based on global sensitivity can overwhelm the statistic, even for large graphs. Instead, by protecting the privacy of a database of node characteristics and then assuming that all downstream statistics are computed from these private statistics, we only have to consider the sensitivity of our statistics to changes in the edge set, for which we show the sensitivity is much lower.

We provide an outline of our approach for this particular statistic in Algorithm 1:

Algorithm 1 Differentially Private Cross-Type Connectedness Index

Require: Network $(\mathcal{V}, \mathcal{E})$ (with weights $e_{ij} \geq 0$), node labels $l_i \in \{a, b\}$, privacy parameter $\varepsilon_l, \varepsilon_e$.
Ensure: Differentially private connectedness index $\widehat{C}_{\mathrm{DP}}^{\mathcal{A}\to\mathcal{B}}$

$p \leftarrow \dfrac{1}{1+e^{\varepsilon_l}}$
for each node $i \in \mathcal{V}$ **do** ▷ Randomized response on labels
  $\hat{l}_i \leftarrow \begin{cases} \text{the opposite value in } \{a,b\}, & \text{with probability } p, \\ l_i, & \text{with probability } 1-p \end{cases}$
end for
for each node $i \in \mathcal{V}$ **do** ▷ Compute debiased weights and individual connectedness
  $d_i \leftarrow \sum_{j\in\mathcal{V}} e_{ij}$
  $\hat{\rho}_i \leftarrow \begin{cases} \dfrac{1}{d_i}\sum_{j\in\mathcal{V}} e_{ij}\,\mathbf{1}\{\hat{l}_j = b\}, & d_i > 0, \\ 0, & d_i = 0 \end{cases}$
  $w_i \leftarrow \dfrac{\mathbf{1}\{\hat{l}_i = a\} - p}{1-2p}$
  $\tilde{\rho}_i \leftarrow \dfrac{\hat{\rho}_i - p}{1-2p}$
end for
$S_0 \leftarrow \sum_{i\in\mathcal{V}} w_i$
$S_1 \leftarrow \sum_{i\in\mathcal{V}} w_i\,\tilde{\rho}_i$
$\widehat{C}_{\mathrm{DP}}^{\mathcal{A}\to\mathcal{B}} \leftarrow \dfrac{S_1}{S_0} + Z_n, \quad Z_n \sim \mathrm{Lap}\left(0, \dfrac{2(1-p)}{(1-2p)^2\,\varepsilon_e\,S_0}\right)$ ▷ Privatized Hájek estimator
return $\widehat{C}_{\mathrm{DP}}^{\mathcal{A}\to\mathcal{B}}$

The released index $\widehat{C}_{\mathrm{DP}}^{\mathcal{A}\to\mathcal{B}}$ has Laplace-induced variance $O\left(1/(\varepsilon_e\,|\mathcal{A}|)^2\right)$, plus additional variance from the randomized response and debiasing.

Connectedness indices are attractive precisely because they are interpretable by policymakers and practitioners. Our mechanism provides a modular, governance-friendly way to allocate privacy budget separately to sensitive attributes and relational information, supporting transparent privacy–utility tradeoffs and review. More broadly, the "labels-first" strategy illustrates a general template for privacy-preserving network statistics: privatize sensitive attributes at the record level, compute bounded aggregates, then add calibrated noise to privatize connections.

References

Blocki, Jeremiah et al. (2013). "Differentially private data analysis of social networks via restricted sensitivity". In: *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, pp. 87–96.

Chetty, Raj et al. (2022a). "Social capital I: measurement and associations with economic mobility". In: *Nature* 608.7921, pp. 108–121.

– (2022b). "Social capital II: determinants of economic connectedness". In: *Nature* 608.7921, pp. 122–134.