On the Privacy-Utility Tradeoff in Peer-Review Data Analysis

Wenxin Ding
Advisors: Nihar B. Shah, Weina Wang

"The main reason behind the lack of empirical studies on peer-review is the difficulty in accessing data. In fact, peer-review data is considered very sensitive, and it is very seldom released for scrutiny, even in an anonymous form." Balietti et al. (PNAS, 2016)

"We are familiar with the literature around privacy preserving dissemination of data for statistical analysis and feel that releasing our data is not possible using current state-of-the-art techniques." Tomkins et al. (PNAS, 2017)

PROBLEM SETTING

Goal is to release some peer-review data while concealing reviewer identities for each paper.

Standard privacy techniques release perturbed versions of the data, and this perturbation can considerably hamper utility of released data.

How to (considerably) improve utility while ensuring the same level of privacy?

Private data:

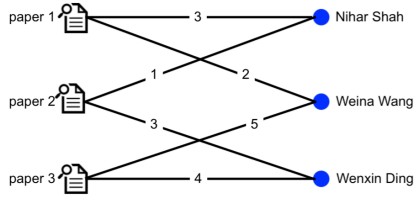

Public data:

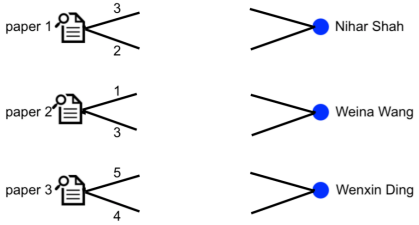

Data to be released:

- Sorted vector of mean scores given by the reviewers.
- This is a general framework which can also be used to release other properties such as:
  - Miscalibration of reviewers
  - Subjectivity of reviewers
- Accuracy (utility): mean squared error between output and true vectors
- Applicable to any privacy paradigm that perturbs data, including differential privacy

KEY OBSERVATIONS

1. Projecting noisy data onto any convex set containing all possible values of the true data **compromises neither privacy nor accuracy.**
2. Projection on "smaller" convex sets is desirable for a higher utility.
3. There is a non-trivial amount of peer-review data available publicly; use it to project on a small convex set to achieve high accuracy!

MAIN THEORETICAL RESULTS

One may be tempted to project on the set of all possible true values (which may not be a convex set).

Theorem 1: Projecting noisy data onto the set of all possible true values can increase the error.

The smallest convex set that contains all possible true values is the convex hull of the true data.

Theorem 2: Projecting on the convex hull of all possible true values is NP−hard.

Designing a polynomial-time algorithm.

Theorem 3: There exists an algorithm which can project any given vector onto a "small" convex set containing all possible true values, and runs in time polynomial in #reviewers.

What is a "small" set?

- Intuitively, if underlying review scores have nice structure, output should have very high accuracy.
- Axiomatic properties:
  1. When all review scores are identical, return a vector with all entries identical to that score.
  2. When every reviewer reviews 1 paper, return a sorted vector of the scores.
  3. When all but one papers receive all zero scores, output vector must have number of non-zero entries equal to number of reviews per paper.

ALGORITHM

Step 1: Compute lower and upper bounds for each entry of sorted mean scores:

- List all tuples consisting of scores from distinct papers, sorted by their means.
- Draw an edge between two tuples if their entries do not overlap.
- Lower bound of the *i*-th smallest mean score is given by the tuple:
  - with left chain of length ≥ *i*
  - leaves enough reviewers with higher means
- Upper bound computed in a similar way

Step 2: Project noisy data onto the convex set defined by:

- Bounds in Step 1
- Constraint 1: sum of values in the output = sum of all reviews scores divided by number of paper reviewed by a reviewer.
- Constraint 2: output is sorted.
- Objective is a simple L2 projection.

EMPIRICAL EVALUATIONS

Synthetic simulations

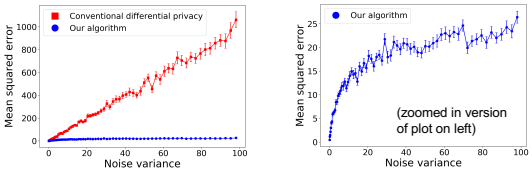

Real data: Grant proposal peer review