Assignment Algorithms to Prevent Quid-Pro-Quo in Conference Peer Review

Ryan Liu advisor: Nihar Shah collaborators: Fei Fang, Hanrui Wang, Steven Jecmen, Vincent Conitzer

Motivation

Modern day research results are organized and shared through large-scale conferences such as AAAI and ICML. These are conferences with thousands of papers and reviewers, where checking for colluding behavior is near impossible by the few program chairs. Prevention or reduction of dishonest behavior through automated assignment algorithms are therefore an important step in ensuring the longevity and reliability today's research publication system.

Problem Statement

At the start of the conference, program chairs construct similarity matrix $s \in [0,1]^{reviewers \cdot papers}$ to represent similarities between each (paper, reviewer) pair. They need to construct a matching $m \in \{0,1\}^{reviewers \cdot papers}$ where each reviewer has at most *k* papers, and each paper has at least *l* reviewers. We additionally have different requirements based on what kind of collusion we want to avoid.

Background Work Done and Usage

Previously, Charlin and Zemel created the Toronto Paper Matching System (TPMS), which maximizes the total sum similarity between assigned papers and reviewers:

$$\sum_{reviewer=i,paper=j} s_{i,j} m_{i,j}$$

My advisor Nihar Shah designed an alternative that focuses on maximizing fairness towards the worst-off paper:

$$\min_{papers} \sum_{\text{reviewer i assigned to j}} s_i j$$

Our assignment algorithms use the sum similarity as the objective of our linear program, & choose from a selection of constraints based on the type of collusion we want to prevent.

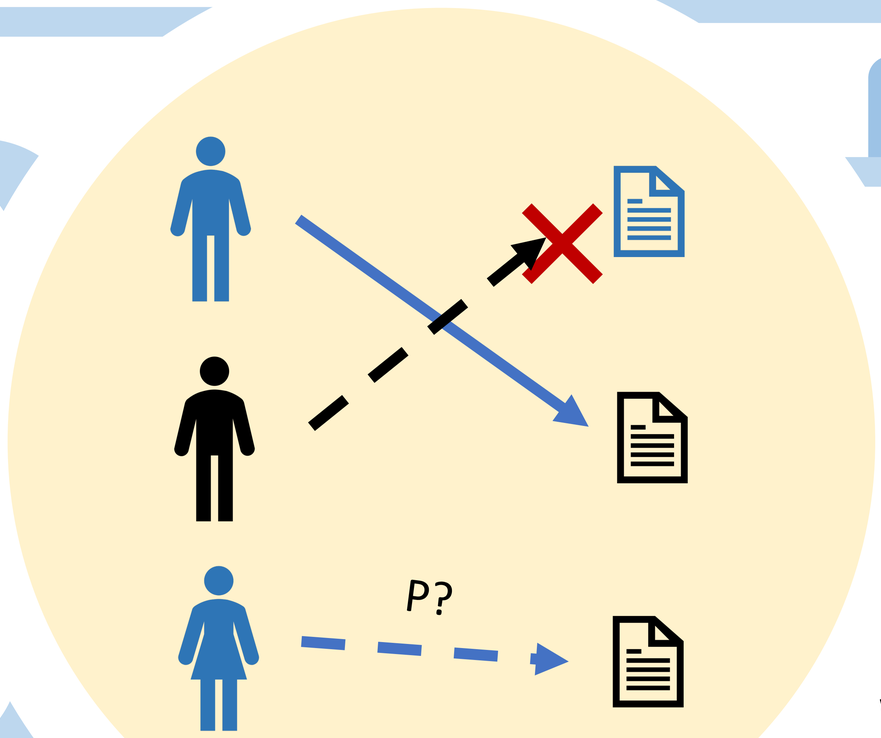

Contribution 1: Pair Collusion Not Totally Unimodular

Without any extra constraints, the constraint matrix of the LP of TPMS is totally unimodular. This is desirable because this means that the LP output will always be integral, which means we will never assign "half a paper".

We implemented a constraint matrix generator for constraint no-pair-collusion, which disallows the case where reviewer *i* reviews a paper by *j*, and reviewer *j* also review a paper by *i.*

We found that under this condition, the constraint matrix is no longer unimodular.

Contribution 2: Implementation & Real-World Testing of Algorithms

Through our collaborators' work, we found that we could upper bound the probability of any assignment by $0 < P < 1$ through the linear program, and then round it to integral values via a flow model. We implemented this and tested it on data from a previous ICLR conference under the TPMS objective. We also generated random similarity matrices to test under TPMS as well as for runtime. From what we can gather, the runtime of this algorithm for an 8k reviewer, 8k paper conference at P = 0.75 is 25 minutes, and sum similarities fall off half as fast as *P* as *P* goes from 1 to 0.

Future Plans & Goals

Continuing on with Contribution 1, we plan to look into ways we can transform a fractional LP solution to an integral one without breaking any of the constraints.

Continuing on with Contribution 2, we plan to check if the algorithm yields results that are equally good under Shah's maximizing fairness constraint.

Charlin and Zemel's TPMS paper: http://www.cs.toronto.edu/~zemel/documents/tpms.pdf Shah's maximizing fairness paper: https://arxiv.org/abs/1806.06237