A *One-Size-Fits-All* Approach to Improving *Randomness* in Paper Assignment

Yixuan Even Xu — Tsinghua
Steven Jecmen — Carnegie Mellon
Zimeng Song — Independent
Fei Fang — Carnegie Mellon

Paper Assignment as an Optimization

Setting: n_p **papers**, n_r **reviewers**, a **similarity matrix** $S \in [0,1]^{n_p \times n_r}$

Standard model: an **integer linear program**

Maximize Quality$(x) = \sum_{p,r} x_{p,r} \cdot S_{p,r}$ (maximize **Quality**)

Subject to $\sum_r x_{p,r} = \ell_p, \forall p$ (Each **paper** gets **exactly** ℓ_p **reviewers**)

$\sum_p x_{p,r} \le \ell_r, \forall r$ (Each **reviewer** gets **at most** ℓ_r **papers**)

$x \in \{0,1\}^{n_p \times n_r}$ (**Assignment** x)

Produces the Maximum-**Quality** Assignment

Randomness Is Crucial in Peer Review

Motivations of randomness:

1. Robustness to malicious behavior like author-reviewer collusion
2. Counterfactual evaluation of alternative paper assignments
3. Diversity of perspectives and expertise among reviewers
4. Reviewer anonymity after releasing paper assignment data

Randomized model: a **continuous linear program**

Maximize Quality$(x) = \sum_{p,r} x_{p,r} \cdot S_{p,r}$

Subject to $\sum_r x_{p,r} = \ell_p, \forall p$

$\sum_p x_{p,r} \le \ell_r, \forall r$

$x \in [0,1]^{n_p \times n_r}$ (**Randomized assignment** x)

Now $x_{p,r}$ denotes the **marginal probability** of assignment

It was shown by (Jecmen et al. 2020) that a **randomized assignment** can be converted into a **distribution of deterministic assignments**

Currently Deployed Algorithm: PLRA

Probability Limited Randomized Assignment (Jecmen et al. 2020):

Maximize Quality$(x) = \sum_{p,r} x_{p,r} \cdot S_{p,r}$

Subject to $\sum_r x_{p,r} = \ell_p, \forall p$

$\sum_p x_{p,r} \le \ell_r, \forall r$

$x \in [0, Q]^{n_p \times n_r}$

Hyperparameter Q:

Guarantees each **paper-reviewer** pair is matched w.p. $\le Q$

Mainly concerned with **robustness to malicious behavior**

PLRA has been deployed in multiple iterations of the **AAAI** conference and is implemented at popular conference management system **OpenReview.net**

Metrics for Randomness

Maximum Probability:

- Maxprob$(x) = \max_{p,r}\{x_{p,r}\}$
- Already used by PLRA as a constraint

Our proposed metrics:

1. **Average maximum probability:** AvgMax$(x) = \frac{1}{n_p}\sum_p \max_r\{x_{p,r}\}$
2. **Support size:** Support$(x) = \sum_{p,r} \mathbf{1}[x_{p,r} > 0]$
3. **Entropy:** Entropy$(x) = \sum_{p,r} x_{p,r} \cdot \log(1 / x_{p,r})$
4. **L2 norm:** L2Norm$(x) = \sqrt{\sum_{p,r} {x_{p,r}}^2}$

Our Proposal: Perturbed Maximization (PM)

A Problem with PLRA: The **randomness** of its assignment depends on **Q**

Not easy to set, and sometimes **suboptimal with any Q**

Example: 2 papers, 2 reviewers, equal similarities

Ideal assignment: uniform assignment (right)

For PLRA, multiple solutions maximize the objective function at the same time. Linear program solvers tend to choose a vertex solution (left).

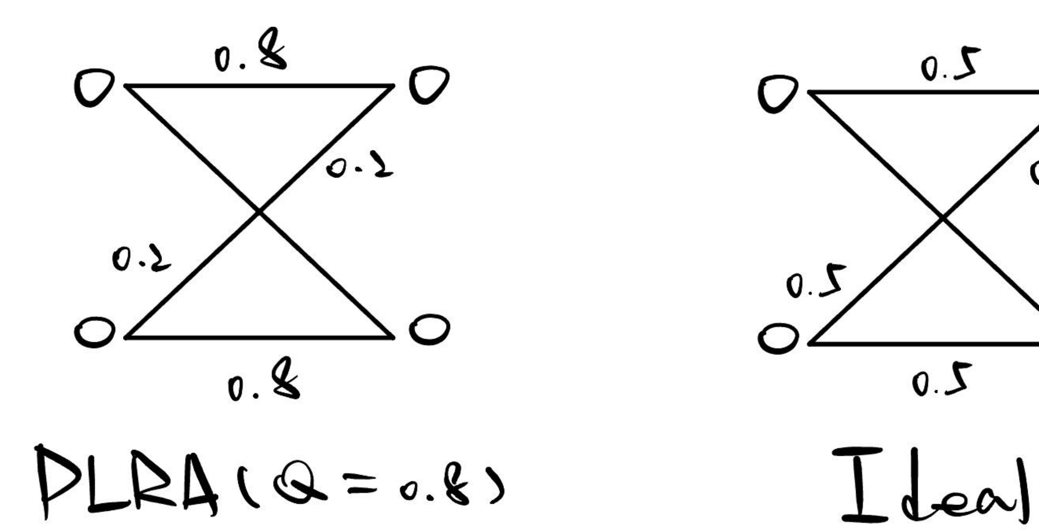

Perturbed Maximization (PM):

Maximize PQuality$(x) = \sum_{p,r} f(x_{p,r}) \cdot S_{p,r}$

Subject to $\sum_r x_{p,r} = \ell_p, \forall p$

$\sum_p x_{p,r} \le \ell_r, \forall r$

$x \in [0, Q]^{n_p \times n_r}$

Perturbation Function $f(\cdot)$**:**

A **non-decreasing**, **concave** function from [0,1] → [0,1]

Intuition: the higher $x_{p,r}$, the lower gain in **PQuality**

Breaks ties evenly, allows trading-off randomness and quality

In the example above:

As long as $f(\cdot)$ is **strictly concave**, by **Jensen's inequality** the only solution that maximizes **PQuality** is the uniform assignment.

PM Is a Provable Improvement

(Informal) On the same input instance, with the same probability limit Q and a **strictly concave** perturbation function $f(\cdot)$, PM outperforms PLRA under **any** of the proposed randomness measures **without loss** in solution **Quality** if the similarity matrix S is:

Blockwise Dominant or **Discrete & Random**

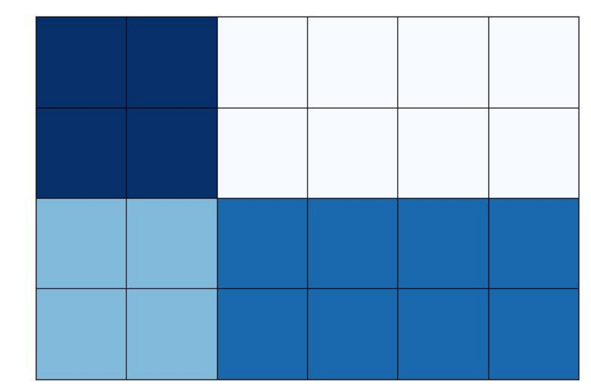

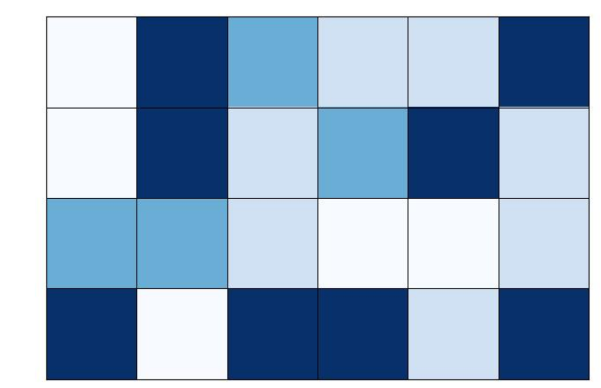

Experiments on Real-World Datasets

On the (discrete) bidding data of **AAMAS2015**, PM has exactly the same performance on **Maxprob** with PLRA (where PLRA is **optimal**), and outperforms PLRA on **all** other randomness measures

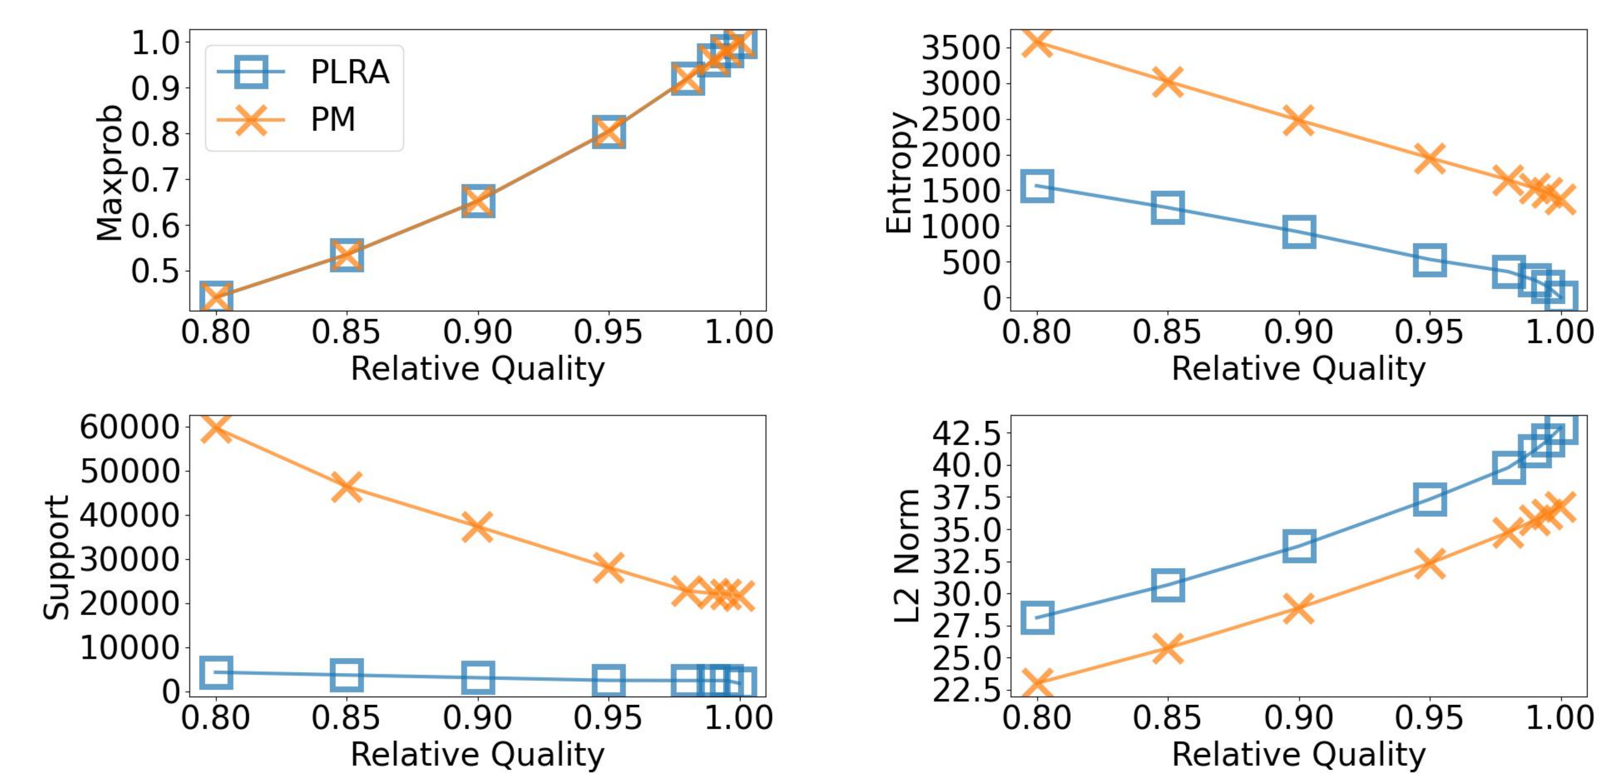

On the (continuous) similarities of **ICLR2018**, PM sacrifices **Maxprob** slightly compared with PLRA (the **optimal Maxprob**), but still improves over PLRA significantly on **all** other randomness measures