View Review

Paper ID
7255

Paper Title
PeerNomination: Relaxing Exactness for Increased Accuracy in Peer Selection

Track Name
IJCAI2020

REVIEW QUESTIONS

---

1. Overall score
Weak Accept (marginally above the acceptance threshold. Rejecting it would not be that bad)

2. Would you champion this paper for the IJCAI-PRICAI 2020 best paper award?
No

3. Confidence in your assessment
High

4. Comments to authors
The authors consider the peer selection problem, in which agents must choose a subset of themselves for an award or a prize. They present a novel algorithm for impartial peer selection, Peer Nomination. The authors touch a bit on two properties of this mechanism - strategy-proofness and monotonicity - and show empirically that it often achieves higher accuracy than some existing algorithms but not always, e.g. EDP gives better results for small values of the reviews per agent.

The problem is interesting. It is also relevant. I believe that this paper should be accepted. However, there is room for improvement.

Major comments:

- The main contribution is a new algorithm, which I really like because it is intuitive and much simpler to understand and easier to implement than EDP and some other existing algorithms.

- The empirical results are promising but nevertheless just incremental. As the authors point out, the testing framework already exists, see [Aziz et al. 2019]. However, they add results about PeerNomination.

- I like the novel connection with the classification measures. However, the authors could spend less space on it as these are common knowledge in the community. I am also not surprised by the ROC and PR curves as these are for a small nomination quota, i.e. 1.66(6).

- The authors sketch that PeerNomination satisfies strategy-proofness and monotonicity – both of which are not formally defined – but they do not show or discuss whether it satisfies committee monotonicity (CM). Such a discussion could improve the paper. For example, EDP satisfies CM (see Aziz et al. [2019]) and is outperformed by PN (see your experiments). Is it because PN violates CM?

Minor comments:

- "The unique profile which is consistent with the ground truth is called truthful." Why is this the case?

- "…the results by Bjelde et al. [2017] shows…"→"…the results by Bjelde et al. [2017] show…"

- The authors use (k/n)*m as the nomination quota. Is there any specific reason behind it?

- What if agents have ties in their rankings? I could extend the algorithm in order to accommodate ties in the agents' rankings. This could be perhaps an interesting future direction.

- ``…in the setting n=130, m=7 and k=30, Figure 2b…"→ This is perhaps a typo as Figure 2b is for n=130,m=9.

- "For most practical applications \epsilon\in [-0.05,0.15]…"→ Perhaps, the authors could provide more details about the applications they have in mind.

- "…Precision-Recall…"→"…Precision-Recall (PR)…"

- "We ran experiment has been repeated…"→Our experiment has been repeated… OR We ran an experiment that has been repeated…

- I would love to see how the new algorithm scales, say to the IJCAI 2020's dimensions, n~8000, m~5-10, k~2000. As m is small, it would be interesting if PeerNomination is again outperformed by EDP as it is for m=5 and smaller values of n and k in your experiment.

5. Why should this paper be presented at IJCAI-PRICAI 2020?
The work is relevant for this conference.

6. Confidential comments
I have no confidential remarks.