View Review

Paper ID
11249

Paper Title
Differentially Private Condorcet Voting

Track Name
Main Track

REVIEW QUESTIONS

1. {Summary} Please briefly summarize the main claims/contributions of the paper in your own words. (Please do not include your evaluation of the paper here).
The authors propose three randomized voting rules based on the Condorcet method under the framework of differential privacy: Laplace Condorcet method (CMLAP), exponential Condorcet method (CMEXP), and Condorcet method with randomized response (CMRR). They estimate the errors these three mechanisms introduce to the voting process and show that CMEXP is the most accurate mechanism in most cases. Next, they propose probabilistic variants of Condorcet criterion, Pareto criterion, and monotonicity criterion and prove that CMRR satisfies all of them, while CMEXP and CMLAP only satisfy probabilistic Pareto criterion and probabilistic monotonicity criterion. The authors also discuss that their rules do not satisfy the classical (i.e. non-probabilistic) versions of these criteria. Finally, the authors prove that CMEXP and CMLAP asymptotically satisfy the probabilistic Condorcet criterion.

2. {Novelty} How novel are the concepts, problems addressed, or methods introduced in the paper?
Fair: The paper contributes some new ideas.

3. {Soundness} Is the paper technically sound?
Good: The paper appears to be technically sound, but I have not carefully checked the details.

4. {Impact} How do you rate the likely impact of the paper on the AI research community?
Fair: The paper is likely to have moderate impact within a subfield of AI.

5. {Clarity} Is the paper well-organized and clearly written?
Excellent: The paper is well-organized and clearly written.

6. {Evaluation} If applicable, are the main claims well supported by experiments?
Good: The experimental evaluation is adequate, and the results convincingly support the main claims.

7. {Resources} If applicable, how would you rate the new resources (code, data sets) the paper contributes? (It might help to consult the paper's reproducibility checklist)
Not applicable: For instance, the primary contributions of the paper are theoretical.

8. {Reproducibility} Are the results (e.g., theorems, experimental results) in the paper easily reproducible? (It may help to consult the paper's reproducibility checklist.)
Fair: key resources (e.g., proofs, code, data) are unavailable but key details (e.g., proof sketches, experimental setup) are sufficiently well-described for an expert to confidently reproduce the main results.

9. {Ethical Considerations} Does the paper adequately address the applicable ethical considerations, e.g., responsible data collection and use (e.g., informed consent, privacy), possible societal harm (e.g., exacerbating injustice or discrimination due to algorithmic bias), etc.?
Not Applicable: The paper does not have any ethical considerations to address.

10. {Reasons to Accept} Please list the key strengths of the paper (explain and summarize your rationale for your evaluations with respect to questions 1-9 above).
1. The paper presents novel voting rules for the case when there might be a privacy leakage in the preferences of

voters.
2. The error evaluation is theoretical and empirical.
3. The contribution is feasible due to the provable tight bounds of the proposed algorithms.

11. {Reasons to Reject} Please list the key weaknesses of the paper (explain and summarize your rationale for your evaluations with respect to questions 1-9 above).
There are some typos and I am not sure whether some of them ruin some of the proofs. For this reason, please read my detailed questions and comments. If these cannot be fixed by the authors during the rebuttal, then this paper cannot be accepted with wrong proofs. The data and code are missing even in the supplementary material

12. {Questions for the Authors} Please provide questions that you would like the authors to answer during the author feedback period. Please number them.
1. How many voters and alternatives are present in the simulations in Figure 1 (i.e. n=? and m=?) and Figure 2 (i.e. n=?)? Please add these details.

2. Theorem 3: "Finally, when w>=2, we have" --> Shouldn't this be w<=2? Does the result change, or is this a typo? Also, what is the difference between w and |w|?

3. Theorem 6: The goal is to show that if ai $\succ_l$ aj holds for any l $\in$ N then P[CMRR(P) = ai]>=P[CMRR(P) = aj] holds. This is Definition 10 (p-Pareto). However, the authors show that P[CMRR(P) = ai]<=P[CMRR(P) = aj] holds. So, there must be a mistake in the proof, which might be a typo. Indeed, if P[CMRR(P) = ai]<P[CMRR(P) = aj] holds whenever ai $\succ_l$ aj holds for any l $\in$ N, then CMRR does not satisfy p-Pareto and this would contradict the statement of the theorem. Could the authors elaborate about it?

13. {Detailed Feedback for the Authors} Please provide other detailed, constructive, feedback to the authors.
MAJORS:

The paper is interesting.

1. The reviewer is not fully convinced why the authors study three different rules. By Table 1, we can see that CMRR satisfies all the proposed probabilistic criteria. For this reason, the analysis of the other two rules requires a further and deeper justification.

2. The authors could have also proposed probabilistic version of strategy-proofness (p-SP), that relaxes strategy-proofness (SP) as it may not be possible to guarantee SP. Then, it would be interesting to see whether any of these new rules satisfy p-SP and, if not, whether we can combine p-SP with either p-Pareto, p-monotonicity, or p-Condorcet. This would motivate the design of new rules that are p-SP and p-Pareto, p-SP and p-monotonic, as well as p-SP and p-Condorcet.

3. The data and code are missing even in the supplementary material. This partly prevents the reproducibility of the paper as one cannot fully verify the simulations.

4. The empirical results in Figure 1 are inline with the statement of Theorem 3. However, to confirm that the empirical results in Figure 2 are also inline with the statement of Theorem 5, the authors could calculate and include the values of \alpha for LAP and EXP in the caption of Figure 2, as well as include in Figure 2 trends that correspond to the equality P[M(P) = CW(P)]=$\alpha \cdot$ P[M(P) = ai]. Thus, the readers can clearly verify that Definition 9 holds in the simulation presented in Figure 2.

MINORS:

1. "transitive, antirefleiive, antisymmetric, and complete" --> "transitive, antireflexive, antisymmetric, and complete"

2. "by lifting" --> Specify what this means. I know it, but the general reader may not.

3. "According to Equation 9, the error rate of CMLAP is" --> Equation 9? Perhaps, "According to Equation 3, the error rate of CMLAP is"

4. ", the “erorr” requires Xij –Xji >w[aj,ai],then the error rate is" --> ", then the “error” requires Xij –Xji >w[aj,ai] and the error rate is"

5. The end of the proof of Proposition 4: "P[CMLAP(P) = a1]<=P[CMLAP(P) = a2]". In fact, this is P[CMLAP(P) = a1]<P[CMLAP(P) = a2] because P[CMLAP(P) = a1]=0.2749 and P[CMLAP(P) = a2]=0.2759. Please write "P[CMLAP(P) = a1]<P[CMLAP(P) = a2]" so that the reader see clearly the violation of the p-Condorcet axiom.

6. Theorems 5 and 6: "... Hence ..." --> "... Hence, ..."

7. "For PMR-EXP, we have" --> What is PMR?

Finally, despite some issues, I recommend this paper for acceptance. I believe that the authors can clarify my concerns during the rebuttal process. However, if some of the proofs cannot be easily fixed, then the paper cannot be accepted, in which case I will decrease my score.

AFTER THE RESPONSE

Thank you for addressing my questions in the response! I encourage you to fix the typos. In particular, I believe that |w| in Theorem 3 should be |w[ai,aj]|. At least, the errors that you define in Theorem 2 use |w[ai,aj]| and not |w|. You should therefore change this notation everywhere because it may lead to confusion with the anti-symmetric matrix w in your "Preliminaries" section. Thus, you let w denote a matrix, but replace w (|w|) in Theorem 3 and in other places where you use w (|w|) as a number and not a matrix with w[ai,aj] (|w[ai,aj]|). Congratulations!

14. (OVERALL EVALUATION) Please provide your overall evaluation of the paper, carefully weighing the reasons to accept and the reasons to reject the paper. Ideally, we should have: - No more than 25% of the submitted papers in (Accept + Strong Accept + Very Strong Accept + Award Quality) categories; - No more than 20% of the submitted papers in (Strong Accept + Very Strong Accept + Award Quality) categories; - No more than 10% of the submitted papers in (Very Strong Accept + Award Quality) categories - No more than 1% of the submitted papers in the Award Quality category

Accept: Technically solid paper, with high impact on at least one sub-area of AI or moderate to high impact on more than one area of AI, with good to excellent evaluation, resources, reproducibility, and no unaddressed ethical considerations.

15. (CONFIDENCE) How confident are you in your evaluation?

Quite confident. I tried to check the important points carefully. It is unlikely, though conceivable, that I missed some aspects that could otherwise have impacted my evaluation.

16. {Confidence-Justification} Please provide a justification for your confidence (only visible to SPC, AC, and Program Chairs).

I am working in COMSOC and, even though in fair division, I often read papers in voting and am familiar with the main issues there. I also teach voting and differential privacy to my students at the university.

17. (EXPERTISE) How well does this paper align with your expertise?

Very Knowledgeable: This paper significantly overlaps with my current work and I am very knowledgeable about most of the topics covered by the paper.

18. {Expertise-Justification} Please provide a justification for your expertise (only visible to SPC, AC, and Program Chairs).

I am working in COMSOC and, even though in fair division, I often read papers in voting and am familiar with the main

issues there. I also teach voting and differential privacy to my students at the university.

20. I acknowledge that I have read the author's rebuttal and made whatever changes to my review where necessary.

Agreement accepted