Reaching Consensus Under a Deadline

JAAMAS Track

Marina Bannikova
Universitat Autònoma de Barcelona
Spain
marina.bannikova@uab.cat

Lihi Dery
Ariel University
Israel
lihid@ariel.ac.il

Svetlana Obraztsova
Nanyang Technological University
Singapore
lana@ntu.edu.sg

Zinovi Rabinovich
Nanyang Technological University
Singapore
zinovi@ntu.edu.sg

Jeffrey S. Rosenschein
The Hebrew University of Jerusalem
Israel
jeff@cs.huji.ac.il

ABSTRACT

Group decisions are often complicated by a deadline. For example, in committee hiring decisions, the deadline might be a budget's start date or the beginning of a semester. It may be that if no candidate is supported by a strong majority, the default is to hire no one—an option that may cost dearly. Hence, committee members might prefer to agree on a reasonable, rather than the best, candidate, to avoid unfilled positions. Here, we propose a model for the above scenario—*Consensus Under a Deadline (CUD)*—based on a time-bounded iterative voting process. We provide theoretical convergence guarantees and an analysis of the resulting decision quality. An extensive experimental study demonstrates more subtle features of CUDs, e.g., the difference between two simple types of committee member behavior, lazy vs. proactive voters. Finally, a user study examines the differences between the behavior of rational voting bots and human voters, concluding that it may often be best to have bots play on the humans' behalf.

KEYWORDS

Social Choice; Consensus; Iterative Voting; Group Decisions

ACM Reference Format:
Marina Bannikova, Lihi Dery, Svetlana Obraztsova, Zinovi Rabinovich, and Jeffrey S. Rosenschein. 2022. Reaching Consensus Under a Deadline: JAAMAS Track. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022,* IFAAMAS, 3 pages.

1 INTRODUCTION

We study the problem of arriving at a joint decision (a consensus) under a deadline, based on the preferences of multiple voters. The voters' task is to find an alternative that the majority agrees upon, before some predefined deadline is reached. The majority is predefined, but can vary from 51% of the votes to unanimous agreement. As different individuals have different preferences, a decision does not occur immediately, and more than one round of voting may be required. Hence, an iterative voting process takes place, where voters potentially change their ballots as the deadline approaches.

Examples where consensus under a deadline is required include an academic hiring committee, selection of a company CEO, a scientific committee deciding where to hold next year's conference, and even the venue choice for a major holiday family dinner. These cases have several common features. First, they each have a strict deadline for reaching an agreement: the start of an academic year, national holidays, a budget approval deadline, etc. Second, assuming that individual voters at least somewhat differ in their preferences, a consensus is unlikely to appear immediately; rather, if reached at all, it is after several rounds of a sequential voting process.

Consequently, we define a strict, formal, time-bounded iterative voting process. The process begins with each voter revealing her most-preferred alternative. If no alternative reaches the required majority, a multi-stage voting process begins. At each stage, all voters that wish to change their ballot apply for a voting slot. For example, a voter may realize that her most-preferred alternative has no chance of being elected, and decide to vote for another prominent alternative. One voter is chosen randomly from all those who applied for a voting slot, the chosen voter casts her new ballot, and the voting result is updated and publicized. The process continues in these stages until either a consensus or the deadline is reached, the sooner of the two. The sequential vote modification process is motivated by real-life scenarios (see the full paper [1] for a lengthy discussion), and supports the great spectrum of the iterative voting literature (see, e.g., [4, 9–11, 13, 14]). We assume that each voter has a private, strict preference order over the alternatives and that the number of alternatives is fixed. Bargaining is not permitted. We further assume that the voters are rational and strategic. The number of strategic ballot changes is unlimited and subject only to the deadline constraint. Each stage (or round) is defined as one clock tick. Lastly, we assume that a failure to reach a consensus is the worst outcome for all voters.

A short version of this paper with preliminary results was presented at a workshop [2]. The CUD game was presented as a short conference paper [15]. The full paper is available at [1].

2 THEORETICAL PROPERTIES OF CUD

We provide, to the best of our knowledge, the first model for iterative voting with a deadline. Although the model can be generalized to other voting rules, we initiate this line of research with a specific model, namely, *Consensus Under a Deadline* (CUD), based on Plurality with a threshold (also known as Majority). We establish the

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online.

theoretical properties of CUDs, such as termination, guarantees of no-default outcomes, and additive Price of Anarchy bounds (PoA^+). Studying the theoretical features of CUD, we find that when there are candidates whose need for votes is smaller than the time until the deadline (that is, there is enough time so that voters could alter their ballots to make this candidate win), then CUD converges with one of these candidates as a winner. Moreover, if there is a candidate for whom at least half the voters voted, then CUD converges with one of these candidates as a winner.

Obviously not all individuals behave identically, which we capture by introducing two simple types of voters: proactive and lazy. Proactive voters are, in a sense, trigger-happy to change their vote, even if just to ensure that their preferred possible winner gets one more point. Lazy voters change their votes only when it is necessary to do so, i.e., when their vote is pivotal to keeping a particular alternative as a possible winner. As CUDs adhere to a simple protocol, we can effectively simulate them to investigate statistical properties For example, theoretical results show PoA^+ principal bounds, but yield no specific trade-offs. Hence, our experiments take a deeper look into additive PoA^+ as a measure of winner quality.

We experiment with a total of eight data sets: four simulated data sets and four real-world data sets, such as [6, 8]. In particular, we measure the effects of voter behavioral types (lazy vs. proactive) on the number of voting steps and PoA^+. First, we find that PoA^+ is identical for both lazy and proactive voters. Second, PoA^+ is low, i.e., in most cases, the plurality winner was the CUD winner. Second, though proactive voters are naturally perceived as actively seeking consensus, they reap no benefit from their activism. Specifically, convergence time of both proactive and lazy voter CUDs is the same. In contrast, as our experiments show, the number of vote changes until convergence is higher for proactive voters. In a way, their behavior is inefficient. Thus, our later voter bot designs were based on the lazy voter behavior.

3 CUD USER STUDY

Even with an experimentally supported theoretical model in place, we were left wondering how human voters would behave. To address this, we built a *CUD-Game*, a game that follows our CUD model structure. We then recruited university students to play it with one another and, at times, with rational lazy-voter bots.

The *CUD-Game* proceeds as follows: First, each voter logs in with a name and an identification number, and waits for the others to join. The game begins once a predefined number of voters is reached. If the number of logged voters exceeds the number of voters for a game, multiple game instances are run simultaneously and independently. At the start of the game, each voter receives a randomly chosen preference profile. The highest preference for each voter is selected automatically, and the current vote stats are shown (Figure 1). On each round, each voter decides whether to change her current selection. If so, they apply for a ballot alteration by changing their selected alternative. The voters are required to reply within a fixed time-span (usually set to 15 seconds). Next, the system randomly selects one of the voters who applied for a ballot change and recomputes the intermediate results. The system checks if either the deadline or a consensus has been reached. If neither occurs, a new round begins with the display of the current ballots and stats. The participants are aware at all times of the number of remaining rounds. If no consensus is reached by the deadline, all receive zero points. When a consensus is reached, each voter receives a score corresponding to the chosen preference.

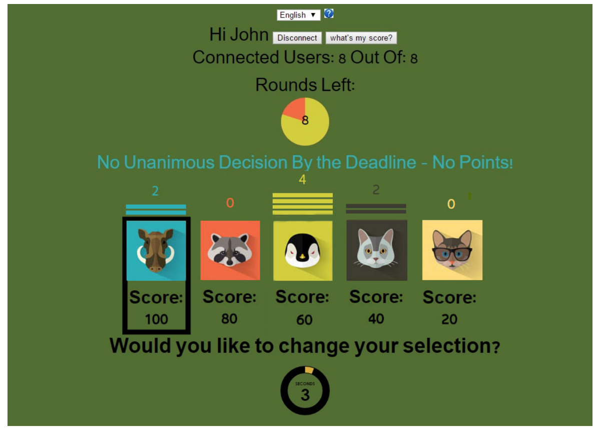

Figure 1: CUD game starts. Selector frame marks the voter's highest preference (here, the blue/boar card).

A total of 72 students played a total of 264 games: 144 mixed games with two bots and six students, and 120 games with no bots The students had no idea whether they were playing with other students or with bots. A total of 10, 000 bots played 1250 bot-only games, eight bots per game. So, do players behave in a rational manner? In our context, rational behavior is the tendency to conform to the group by voting for a candidate that was chosen by the majority. In contrast, irrational behaviors are actions that do not promote consensus-building. We followed the two types of irrational voting defined in a parallel study [5]: *Opposing Alignment (OA)*, where a candidate with a lower rank than the current winner receives a vote; and an *Inappropriate Alignment (IA)*, where a candidate with a lower score and rank than the current *selection/ballot* receives a vote. To elucidate this, consider the player from Figure 1. The candidate preferred by most players is the penguin, and it has a value of 60 for the observed player. An OA action would be to vote for one of the cards that have a lower rank (and value) than the penguin (e.g., the cats). An IA action would be to vote for the raccoon (valued at 80), which has a lower rank and a lower number of votes in the stats than the current selection/ballot of pig (valued at 100, and having 2 votes).

Decisions made purely by bots seem to have higher quality (game score) than decisions made by humans. This coincides with [3], who conclude that interacting via an agent leads to fairer results. Bots are playing ever-larger roles in our daily lives. We rely on them for navigation directions, information retrieval and perhaps one day, to reach a group decision. This research is a step in that direction. Future voters will set their preferences [7, 12] for decisions that they care about, and then use the system to find a consensus. To assist in this path, our framework can serve as one of the core testbeds for group decision making and psychological studies; e.g., CUD-Game has been used to study consensus in people with Alexithymia [5].

ACKNOWLEDGMENTS

This research has been partly funded by Israel Science Foundation grant #1340/18.

REFERENCES

[1] Marina Bannikova, Lihi Dery, Svetlana Obraztsova, Zinovi Rabinovich, and Jeffrey S Rosenschein. 2021. Reaching consensus under a deadline. *Autonomous Agents and Multi-Agent Systems* 35, 1 (2021), 1–42.

[2] Marina Bannikova, Lihi Naamani-Dery, Svetlana Obraztsova, Zinovi Rabinovich, and Jeffrey S Rosenschein. 2016. Between Fairness and a Mistrial: Consensus Under a Deadline. In *The 10th workshop on Advances in Preference Handling (M-PREF)*.

[3] Celso M de Melo, Stacy Marsella, and Jonathan Gratch. 2018. Social decisions and fairness change when people's interests are represented by autonomous agents. *Autonomous Agents and Multi-Agent Systems* 32, 1 (2018), 163–187.

[4] Lihi Dery, Svetlana Obraztsova, Zinovi Rabinovich, and Meir Kalech. 2019. Lie on the Fly: Strategic Voting in an Iterative Preference Elicitation Process. *Group Decision and Negotiation* 28, 6 (2019), 1077–1107.

[5] Hila Zahava Gvirts and Lihi Dery. 2020. Alexithymia and reaching group consensus. *Cognition and Emotion, to appear* (2020).

[6] T. Kamishima, H. Kazawa, and S. Akaho. 2005. Supervised ordering - an empirical survey. In *Proceedings of the 5th IEEE International Conference on Data Mining*. 673–676.

[7] MK Lee et al. 2018. *WeBuildAI: Participatory Frameworks for Fair and Efficient Algorithmic Governance*. Technical Report. CMU technical report.

[8] N. Mattei and T. Walsh. 2013. Preflib: A library for preferences. In *Proceedings of the 3rd International Conference on Algorithmic Theory*. 259–270.

[9] Reshef Meir. 2016. Strong and weak acyclicity in iterative voting. In *International Symposium on Algorithmic Game Theory*. Springer, 182–194.

[10] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. 2010. Convergence to Equilibria of Plurality Voting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*. 823–828.

[11] L. Naamani-Dery, S. Obraztsova, Z. Rabinovich, and M. Kalech. 2015. Lie on the Fly: Manipulative Voters under Iterative Voting Centre. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*. 2033–2039.

[12] Ritesh Noothigattu, Snehalkumar S Gaikwad, Edmond Awad, Sohan Dsouza, Iyad Rahwan, Pradeep Ravikumar, and Ariel D Procaccia. 2018. A voting-based system for ethical decision making. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

[13] S. Obraztsova, E. Markakis, M. Polukarov, Z. Rabinovich, and N. R. Jennings. 2015. On the Convergence of Iterative Voting: How Restrictive should Restricted dynamics be?. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*. 993–999.

[14] Annemieke Reijngoud and Ulle Endriss. 2012. Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, Vol. 2. 635–644.

[15] David Ben Yosef, Lihi Naamani-Dery, Svetlana Obraztsova, Zinovi Rabinovich, and Marina Bannikova. 2017. Haste Makes Waste: A Case to Favour Voting Bots. In *Proceedings of the International Conference on Web Intelligence* (Leipzig, Germany). 419–425.