Preference learning for e-democracy

EDMITT PhD project
Umberto Grandi and Jérôme Mengin

Starting date: September 2023

We look for applicants with a master degree in computer science, mathematics, or related disciplines, with a motivation to pursue theoretical research.

1 Context: e-democracy, social choice, and preference learning

There is growing interest on e-democracy applications both from the general public, who is in demand of an increase in participation to collective and institutional choices, and from researchers in computer science and economics. While many e-democracy applications are still based on simple techniques, some are implementing innovative non-trivial algorithms. One example is participatory budgeting, where the voters of a community need to choose which projects will be funded subject to an overall budget. A large research effort proposed rules and elicitation protocols for such collective choices, such as the recent method of equal share that is now under empirical test in a few cities around Europe [Peters et al., 2021]. Another example are deliberation platforms, where citizens can input proposals and vote on them dynamically in search of an agreement. A prime example is Pol.is which was extensively used in Taiwan in 2014, and which uses a matrix factorisation algorithm to decide which of proposals are seen and considered for voting by each user.

The research community on social choice and computational social choice has in the past decades developed a wide number of theoretical results and algorithms that can be put to test in novel e-democracy applications.[1] Traditionally, social choice has considered situations where a set of voters submit full information on their preferences over a set of alternatives, for instance in the form of a ranking. However, in e-democracy applications such preferences need to be collected under heavy time and space constraints, as the participation of users is typically very limited in time, and the choice of the few questions that can be queried is thus of prime importance. Hence the need of exploring preference learning techniques, in particular recent algorithms that take ordinal

[1] To get an idea about this research community consult this page: `http://research.illc.uva.nl/COMSOC` or consult the handbook by Brandt et al. [2016].

information as input, such as learning via pairwise comparison data (see, e.g., Bengs et al. [2022]) or active learning (see, e.g., Mikhailiuk et al. [2021]).

2 Research questions

This research subject wants to tackle the problem of partial participation in collective decisions, studying e-democracy applications through the lens of computational social choice using techniques from preference learning.

The first problem that will be considered is the elicitation protocol of preferences for one voter only. While a solution using optimisation techniques can be considered, we will be interested in the use of active learning. The novelty will be the use of non-standard preference elicitation languages (such as mixing approvals and preferences).

The second problem that we will consider is the definition of a query strategy on preference items and a set of users to learn most efficiently the outcome of a collective decision. Approaches such as PAC learning will be relevant here, understanding whether it will be possible to learn efficiently the result of a given collective choice rule.

The main techniques that will be used are mathematical modelling and analysis, simulations, and if possible data analysis on real data obtained via e-democracy applications.

Examples of related recent work are the following papers:

- Lee et al. [2014], which is based on earlier work by Xia and Conitzer [2011], studies elicitation of preferences to compute the result of voting rules
- Halpern et al. [2022] is a very recent paper inspired by Pol.is, where the authors consider how to rank proposals to users with only very sparse data

3 Job description

The thesis is going to be supervised by Umberto Grandi and Jérôme Mengin from the LILaC (`https://www.irit.fr/LILaC/site/`) and the ADRIA (`https://www.irit.fr/ADRIA/site/`) groups, which are part of the AI Department at IRIT. The PhD student will be given an office and a workstation at University of Toulouse 1 Capitole, at the Manufacture des Tabacs campus located in allée de Brienne.

The PhD student will be expected to conduct research activities leading to publications in scientific conferences and journals, including presenting the work at international conferences. Teaching is possible and encouraged, up to 64 frontal hours per year, and is remunerated around 40 euros per hour on top of the PhD scholarship. Working hours are flexible.

The PhD student will be part of the EDMITT doctoral school, who provide further supervision and training courses (some of them mandatory). For more information: `https://ed-mitt.univ-toulouse.fr/`

As of 2022, the current (gross) salary for PhD students is 1 975 EUR.

4 Selection procedure

A two-step selection procedure will take place:

- applicants will be evaluated based on their CV, application letter, and eventual zoom interview conducted by the two supervisors, selecting one candidate to be presented for funding. Timeline: end of April
- the selected candidate and the supervisors will then apply to the selection committee of the doctoral school who is in charge of allocating the scholarships. The selection will be based on the adequacy between the project and the candidate profiles, as well as the candidate's track record, together with a 20 minutes presentation in front of the committee (if possible in Toulouse, or on zoom). Timeline: interviews usually take place end of May/beginning of June, and the ranking is published in the days following the interviews

5 Contacts

For any question please contact by email Umberto Grandi (`umberto.grandi@irit.fr`) or Jérôme Mengin (`jerome.mengin@irit.fr`).

References

V. Bengs, R. Busa-Fekete, A. El Mesaoudi-Paul, and E. Hüllermeier. Preference-based online learning with dueling bandits: A survey. *J. Mach. Learn. Res.*, 22(1), jul 2022.

F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.

D. Halpern, G. Kehne, A. D. Procaccia, J. Tucker-Foltz, and M. Wüthrich. Representation with incomplete votes. *arXiv preprint arXiv:2211.15608*, 2022.

D. Lee, A. Goel, H. Landemore, and T. Aitamurto. Crowdsourcing for participatory democracies: Efficient elicitation of social choice functions. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 2, 06 2014.

A. Mikhailiuk, C. Wilmot, M. Perez-Ortiz, D. Yue, and R. K. Mantiuk. Active sampling for pairwise comparisons via approximate message passing and information gain maximization. In *25th International Conference on Pattern Recognition (ICPR)*, 2021.

D. Peters, G. Pierczyński, and P. Skowron. Proportional participatory budgeting with additive utilities. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, 2021.

L. Xia and V. Conitzer. Determining possible and necessary winners given partial orders. *J. Artif. Intell. Res.*, 41:25–67, 2011.