Exploring the Relationship Between Social Choice and Machine Learning

Doctoral Consortium

Ben Armstrong
University of Waterloo
b8armstr@uwaterloo.ca

ABSTRACT

My thesis will study the intersection of social choice and machine learning, with a focus on recent or under-explored social choice paradigms, such as liquid democracy, and how social choice and ML can benefit each other. My initial results show the idea of using ML and social choice to understand the other holds promise. An early project of mine uses deep learning to enhance social choice by creating a neural network that acts as a voting rule, able to be trained to select a winner satisfying customizable sets of axioms. More recently, I have explored the idea of using liquid democracy as a framework for ensembles for classification problems. I am particularly interested in improving the real-world applicability of existing social choice methods and understanding how they can more beneficially impact the world. Going forward, my primary tools in these goals are simulation and provable axiomatic or performance guarantees.

KEYWORDS

Social Choice; Machine Learning; Liquid Democracy

ACM Reference Format:
Ben Armstrong. 2021. Exploring the Relationship Between Social Choice and Machine Learning: Doctoral Consortium. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 3 pages.

1 DEEP LEARNING TO SUPPORT SOCIAL CHOICE

A clear parallel can be drawn between voting rules in social choice and classification tasks in machine learning. Voting rules often take the form of a function taking in a set of preferences and returning a winning alternative (though may also return a ranking or set of winners) while classifiers take in data of a general form and predict a single class. An early project in my PhD builds on this similarity and focuses on the use of neural networks to act as novel voting rules with the goal of providing less manipulable social choice mechanisms with properties customizable to the domains in which they are used.

While relatively little work appears to exist on this subject, early positive results have shown that existing scoring rules can be efficiently learned [19], lending merit to the parallel between machine learning and voting. Other recent work has also focused on comparisons between existing voting rules and machine learning [16] and has outlined several benefits of novel social choice mechanisms created by machine learning [22].

1.1 My Contribution

My existing work in this area develops a system for training neural networks that use ranked preferences to find alternatives that satisfy a customizable set of axioms [1].

Each training example in this system takes the form of a set of ranked preferences, a single alternative, and an associated score corresponding to the utility of that alternative being “predicted” by the network as the winner under the given preferences. Before training, the system designer decides on a set of desirable axioms (e.g. Condorcet consistency, participation, non-dictatorship) and the relative importance of these axioms. During training, the scoring function gives a higher score to alternatives satisfying more of these axioms.

Many impossibility results, most famously Arrow’s theorem, have been found that identify sets of axioms which cannot be mutually satisfied for all possible sets of preferences by any voting rule [11]. My preliminary results have so far shown the ability to identify the Condorcet winner with the ultimate goal of showing the possibility of training networks that will satisfy mutually impossible axioms in common cases more likely to appear in practical situations (e.g. for voters drawn from common preference models such as Mallows [17] or Polya-Eggenberger Urn [3]).

There are two primary questions I plan to address in future work in this area. First, it is important to quantify, through one or both of theory and thorough experimentation, to what extent any mechanism can satisfy the axioms it is trained for. Second, as has been briefly explored in recent work [9, 18], novel voting rules (which may essentially be created on-demand for specific applications) can improve privacy and reduce manipulability. Both of these effects should be studied further in order to show the usefulness of such an approach and methodology.

2 LIQUID DEMOCRACY ENSEMBLES

My work utilizing neural networks to improve social choice outcomes uses machine learning to benefit social choice. This relationship can work in the other direction as well. In work currently under review I explore whether recent techniques from social choice can be used to improve results in machine learning. More specifically, I apply the liquid democracy framework to ensemble learning.

Liquid democracy is an emerging paradigm of delegative voting in which delegations are transitive. Voters may delegate to another voter or may cast a weighted vote, with the weight determined by the total number of direct and indirect delegations they received.

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

In basic models voters select a single other voter to act as their delegate or, instead, cast a vote directly. It is common for papers in this area to focus on models of "ground-truth" voting wherein voters aim to select a single correct alternative using a weighted majority vote [4, 14].

My work begins with theoretical results showing that in a basic weighted majority two-candidate liquid-democracy system delegation can always weakly improve group accuracy when it comes to determining the ground-truth. Similar results are also shown when this is extended to a model in which voters are classifiers, voting on (predicting the class of) multiple issues at once. The goal is to improve the accuracy of the ensemble through delegation which is seen to typically be possible, despite being NP-hard to optimize in the worst case [5].

However, I test a variety of delegation mechanisms aimed at increasing classifier diversity or retaining the more accurate classifiers and find that, when studied empirically, they provide no significant improvement to accuracy.

In fact, concerns about the benefit of using liquid democracy abound. While some results show situations in which liquid democracy can guarantee a benefit to group accuracy [14] it is common for other papers to focus on concerns largely unique to the system such as preventing delegation cycles [6] or mitigating dictatorships caused by delegations congregating at a small number of voters [13, 15]. In a critique using a model similar to my own, [5] show that finding optimal delegations may require very unintuitive actions, such as delegation from more accurate to less accurate voters.

Based on these issues and my experimental results my paper concludes that liquid democracy is likely unsuitable for ground truth voting and that ongoing research in liquid democracy should place more emphasis on empirical analysis in addition to theoretical results.

There is a great deal of further research to be done on liquid democracy. Christoff and Grossi study a theoretical model of liquid democracy in which agents vote on "logically interdependent propositions," which bears a certain resemblance to classifiers voting in our setting [6]. Results such as theirs may be applicable in my future exploration of this ML-social choice relationship. For this, I plan to begin with an axiomatic analysis providing further theoretical results on the guarantees that machine learning can provide in this or similar settings.

3 FUTURE DIRECTIONS

3.1 Theoretical Guarantees for Existing Results

While my existing research includes both theoretical and experimental results, to date it has had a strong empirical focus. In the near future, I plan to explore the possibility of theoretical analysis to provide more guarantees to existing projects.

Currently I am surveying the broad literature on the application of SAT problems to social choice. This subfield began with the use of SAT solvers to prove Arrow's Theorem and automatically generate many additional impossibility theorems [11, 12]. I believe this methodology can be applied far more widely than it has been, including to show the axiomatic potential (or lack thereof) in liquid democracy systems. I hope to also draw on this and connections with recent work in automated mechanism design [20] to strengthen my results combining social choice and deep learning.

3.2 Novel Social Choice Mechanisms

Recently a number of relatively modern social choice mechanisms have entered the research spotlight, such as liquid democracy and participatory budgeting [2]. An older mechanism now seeing renewed interest is sortition, the process of choosing a random subset of voters to make a decision on behalf of a larger population. The selection of voters for a sortition panel gives the opportunity for improved fairness guarantees and accurate representation of demographic diversity [10] while also increasing the difficulty of manipulations for common voting rules [21].

I see multiple potential connections between sortition and machine learning which I hope to explore in the future. First, both sortition and ML make heavy use of sampling techniques in order to more accurately understand a broader population (of voters, or of data). Second, perhaps overlapping with the first, is an emphasis on fairness. Both domains have results aimed at ensuring group fairness (e.g. demographic diversity in sortition, similar classification errors across subgroups in ML) which may be beneficial to the other domain.

3.3 Social Good

An intentional goal common to most of my work is a interest in results that can be applied, in a practical sense, to real world tasks. This culminates in a focus on providing some positive benefit through my research. This pursuit of social good is common in social choice (e.g. kidney exchanges, better representation through voting [8]) and is also a rapidly growing trend in machine learning [7]. There are many existing mechanisms from both of these areas that significantly affect the lives of most individuals. My ongoing and future research will bear in mind this responsibility and aim to improve the impact these systems have on us.

REFERENCES

[1] Ben Armstrong and Kate Larson. 2019. Machine Learning to Strengthen Democracy. *NeurIPS Joint Workshop on AI for Social Good* (2019).

[2] Haris Aziz and Nisarg Shah. 2021. Participatory budgeting: Models and approaches. In *Pathways Between Social Science and Computational Social Science*. Springer, 215–236.

[3] Sven Berg. 1985. Paradox of voting under an urn model: The effect of homogeneity. *Public Choice* 47, 2 (1985), 377–387.

[4] Daan Bloembergen, Davide Grossi, and Martin Lackner. 2019. On rational delegations in liquid democracy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1796–1803.

[5] Ioannis Caragiannis and Evi Micha. 2019. A Contribution to the Critique of Liquid Democracy. In *IJCAI*. 116–122.

[6] Zoé Christoff and Davide Grossi. 2017. Binary Voting with Delegable Proxy: An Analysis of Liquid Democracy. *Electronic Proceedings in Theoretical Computer Science* 251 (Jul 2017), 134–150. https://doi.org/10.4204/EPTCS.251.10

[7] Josh Cowls, Andreas Tsamados, Mariarosaria Taddeo, and Luciano Floridi. 2021. A definition, benchmark and database of AI for social good initiatives. *Nature Machine Intelligence* 3, 2 (2021), 111–115.

[8] Ulle Endriss. 2017. *Trends in Computational Social Choice*. Lulu. com.

[9] Daniel Firebanks-Quevedo. [n.d.]. Machine Learning? In MY Election? It's More Likely Than You Think: Voting Rules via Neural Networks. ([n. d.]), 19.

[10] Bailey Flanigan, Paul Gölz, Anupam Gupta, and Ariel Procaccia. 2020. Neutralizing Self-Selection Bias in Sampling for Sortition. *arXiv preprint arXiv:2006.10498* (2020).

[11] Christian Geist and Ulrich Endriss. 2011. Automated search for impossibility theorems in social choice theory: Ranking sets of objects. *Journal of Artificial*

Intelligence Research 40 (2011), 143–174.
[12] Christian Geist and Dominik Peters. 2017. Computer-aided methods for social choice theory. *Trends in Computational Social Choice* (2017), 249–267.
[13] Paul Gölz, Anson Kahng, Simon Mackenzie, and Ariel D Procaccia. 2018. The fluid mechanics of liquid democracy. In *International Conference on Web and Internet Economics*. Springer, 188–202.
[14] Anson Kahng, Simon Mackenzie, and Ariel Procaccia. 2018. Liquid democracy: An algorithmic perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
[15] Christoph Kling, Jérôme Kunegis, Heinrich Hartmann, Markus Strohmaier, and Steffen Staab. 2015. Voting behaviour and power in online democracy: A study of LiquidFeedback in Germany's Pirate Party. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 9.
[16] Hanna Kubacka, Marija Slavkovik, and Jan-Joachim Ruckman. [n.d.]. Predicting the winners of Borda, Kemeny and Dodgson elections with supervised machine learning. ([n. d.]), 18.
[17] Colin L Mallows. 1957. Non-null ranking models. I. *Biometrika* 44, 1/2 (1957), 114–130.
[18] Farhad Mohsin, Ao Liu, Francesca Rossi, and Lirong Xia. 2020. Learning to Design Fair and Private Voting Rules. *Cooperative AI workshop at NeurIPS 2020* (2020).
[19] Ariel D. Procaccia, Aviv Zohar, Yoni Peleg, and Jeffrey S. Rosenschein. 2009. The Learnability of Voting Rules. *Artificial Intelligence* 173, 12–13 (Aug 2009), 1133–1149. https://doi.org/10.1016/j.artint.2009.03.003
[20] Weiran Shen, Pingzhong Tang, and Song Zuo. 2019. Automated mechanism design via neural networks. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems*. 215–223.
[21] Toby Walsh and Lirong Xia. 2012. Lot-based voting rules.. In *AAMAS*, Vol. 12. Citeseer, 603–610.
[22] Lirong Xia. 2013. Designing social choice mechanisms using machine learning. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems*. 471–474.