



HAL
open science

Agent-Mediated Social Choice

Umberto Grandi

► **To cite this version:**

Umberto Grandi. Agent-Mediated Social Choice. J.-F. Laslier; H. Moulin; M. R. Sanver; W. S. Zwicker. The Future of Economic Design. The Continuing Development of a Field as Envisioned by Its Researchers, Springer, pp.89-96, 2019, Studies in Economic Design book series (DESI), 978-3-030-18050-8. 10.1007/978-3-030-18050-8_14. hal-03066935

HAL Id: hal-03066935

<https://hal.science/hal-03066935>

Submitted on 15 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Agent-Mediated Social Choice

Umberto Grandi
Institut de Recherche en Informatique de Toulouse (IRIT)
University of Toulouse
umberto.grandi@ut-capitole.fr

Abstract

Direct democracy is often proposed as a possible solution to the 21st-century problems of democracy. However, this suggestion clashes with the size and complexity of 21st-century societies, entailing an excessive cognitive burden on voters, who would have to submit informed opinions on an excessive number of issues. In this paper I argue for the development of “voting avatars”, autonomous agents debating and voting on behalf of each citizen. Theoretical research from artificial intelligence, and in particular multiagent systems and computational social choice, proposes 21st-century techniques for this purpose, from the compact representation of a voter’s preferences and values, to the development of voting procedures for autonomous agents use only.

1 Introduction

Computational studies of voting are mostly motivated by two intended applications: the coordination of societies of artificial agents, and the study of human collective decisions whose complexity requires the use of computational techniques. Both research directions are too often confined to theoretical studies, with unrealistic assumptions constraining their significance for real-world situations. Most practical applications of these results are therefore confined to low-stakes decisions, which are of great importance in expanding the use of algorithms in society, but are far from high-stakes choices such as political elections, referenda, or parliamentary decisions, which societies still make using old-fashioned technologies like paper ballots.

In this paper I argue in favour of conceiving “voting avatars”, artificial agents that are able to act as proxies for voters in collective decisions at any level of society. Besides being an ideal test-bed for a large number of techniques developed in the field of multiagent systems and artificial intelligence in general, agent-mediated social choice may also suggest innovative solutions to low voter participation, a problem that is endemic in most practical implementations of electronic decision processes.

From low-stakes/high-frequency to high-stakes/high-frequency social choice. In their quest for practical applications, researchers in (computational) social choice

have argued in favour of moving away from political elections, and high-stakes/low-frequency collective decisions in general, where computational techniques and studies are less relevant. Low-stakes decisions such as answering to a personalised search query, or designing group recommender systems for retailers, were instead identified as more suitable applications for their research (Boutilier and Lu, 2011). A similar trend can be observed in existing platforms for electronic democracy such as LiquidFeedback¹, which were initially designed to be used for policy design by political parties, and evolved into decision systems for smaller committees and low-stakes decisions (Behrens et al., 2014). Many researchers, however, (including the author of this paper) became interested in social choice with the dream of having an impact on how large societies take high-stakes decisions. Technological advancements now make it possible to carry out polls and surveys among citizens in almost real-time. *Has the time come for high-stakes/high-frequency collective decision-making?*

Electronic, direct, participative, and interactive democracy. Electronic democracy (or e-democracy) is an umbrella term that groups several pieces of software and policies that aim at increasing citizens' participation in collective decisions by means of digital technologies. Among those, the most relevant applications for social choice theorists revolve around the development of interactive or direct forms of democracy through the design of electronic platforms. While on the one hand existing platforms tend to be rather simple, or, as mentioned in the previous paragraphs, they are restricted to low-stakes decisions, on the other hand more encompassing visions of direct democracy often seem excessively utopian in imagining an active participation of the entire electorate to collective debates and votes (Green-Armytage, 2015). Real-world experiments seem to suggest the opposite: when the citizens of Madrid were asked to vote directly (and electronically) on the renovation projects of one of the most important squares in the city, only 8% of eligible voters actually took part in the voting process.² The use of direct (or interactive) democracy platforms for high-stakes decisions seems to produce a very low citizen participation, a phenomenon that might be considered as an instance of the well-known paradox of voting (Downs, 1957), i.e., the cost of casting a single vote exceeding the expected benefit of affecting the result of the election. *Can we imagine a technological advancement that will solve the problem of low participation in direct democratic decisions?*

Artificial intelligence and democracy. Techniques from artificial intelligence (AI), such as the use of machine learning for user profiling and micro-targeting, have been widely used in recent political campaigns, and their effects have been widely debated by the press. However, most articles limit the future use of AI in elections to the design of centralised algorithms that take collective decisions from the collection of citizens' preferences and behaviours.³ While discussing the potential misuse of data analysis techniques by central governments, a recent article by Helbing et al. (February 2017) claims that "If data filters and recommendation and search algorithms would be se-

¹<http://liquidfeedback.org/>

²https://elpais.com/elpais/2017/02/28/inenglish/1488280371_057827.html

³A vision already proposed by Asimov (1955) in the short story "Franchise".

lectable and configurable by the user, we could look at problems from multiple perspectives, and we would be less prone to manipulation by distorted information”. *Can we develop personalised AI techniques that help people construct and motivate their views and voting behaviour in collective decisions?*

Agent-mediated e-commerce and computational mechanism design. Negotiation technology, trust-building, and a vast number of other techniques developed in multi-agent systems found application in e-commerce, where human and artificial agents participate with various roles in suitably designed markets. A research field that started more than fifteen years ago (Sierra, 2004; Feigenbaum et al., 2009; Dash et al., 2003) helped creating a new reality: “The anticipated agent-mediated economy is almost upon us” (Parkes, 2017). *Once people get used to delegate their consumer power to artificial agents, will they be ready to delegate their citizen power as well?*

Computational social choice and its applications. The field of computational social choice (COMSOC) started around ten years ago when a group of researchers with common interests in computer science, economics, and political science regrouped in an international workshop,⁴ and flourished up to the recent publication of a handbook (Brandt et al., 2016). Research in this field is rather theoretical, and its motivation often ambiguous between developing technologies for the coordination of artificial agents and the study of theoretical properties of voting mechanisms to be later used by humans. Unrealistic assumptions of complete knowledge or full rationality of individual agents limit the applicability of many of its results, a problem that would be less relevant in agent-mediated institutions. If the field is to succeed and prosper, it needs applications to feed its research agenda, experimenting with either real-world situations (as argued by Bouveret, 2017) or agent-based technologies. A recent positive example is the creation of the Spliddit webpage (Shah, 2017), which created a fruitful loop between practical applications and theoretical research. *Can we conceive of technologies enabling agent-mediated social choices, where theorems and algorithms from COMSOC research could be applied without changing their currently unrealistic assumptions?*

2 Voting Avatars - A Short Story

Sylvia (a human being) lives in a world where each citizen with the right to vote is paired with a *voting avatar*, an autonomous agent that is able to communicate with her and who is authenticated by a central voting authority to vote on her behalf.⁵ Sylvia is the only authorised person to communicate with her voting avatar (e.g., speech recognition, fingerprint authentication, ...). The central voting authority is a democratically elected government with executive power, supported by an elected bureaucracy who sets the agenda of debates, polls, and votes to be conducted among the entire electorate. Each morning Sylvia receives the daily political agenda, with issues classified

⁴www.illc.uva.nl/COMSOC/workshops.html

⁵This section may read like a science-fiction novel, but so do many academic papers advocating for direct democracy.

by themes and by interest: local, regional, national, and global (the world in which Sylvia lives is likely to be a world federation). Sylvia can simply ignore the message and have a good cup of coffee, as she does on most days: Her voting avatar has already been searching the internet for opinions, consulted influential avatars, and built a preliminary voting behaviour for her daily agenda. During the day, the voting avatar will follow all discussions and correlated votes, and update Sylvia's voting behaviour based on this information and on the level of strategic behaviour she set (currently she left the "strategic voting" button unchecked, like most of her friends claim to do). Her voting avatar has been training for years on a number of votes Sylvia takes directly every month, as well as by observing her conversations on social media, by reading her emails, and by having direct conversations with her every time that the avatar made a wrong or debatable decision on her behalf. Last night, for instance, Sylvia realised that the avatar suggested voting against issuing extra visas to refugees from the Mars colonies, based on a number of dubious sources that she had consulted a couple of days earlier out of curiosity. They discussed the issue for a good 5 minutes, clarifying her position on immigration, the job market, and charity (she actually found the discussion very helpful in constructing a solid view on these issues). Today Sylvia is quite interested in the debate on global freezing, and her voting avatar is proposing a vote in support of the current bill (decisions with long-term consequences involve a long series of iterated votes on improving proposals, in order to maximise consensus). Sylvia has access to a short summary of the reasons supporting the avatar's suggestions, with links to a number of articles by authors she finds reliable, extracts from email discussions she had with a friend on this topic, as well as a list of her past decisions on related issues. She notices that the coherence warning is yellow, suggesting that her vote clashes with some of the positions she defended on the energy market a couple of months ago, but not being a public figure she chooses to ignore the warning...

3 Multiagent Systems and Artificial Intelligence

In the past 20 years researchers in the field of multiagent systems, a research community in artificial intelligence, have been importing and adapting models from theoretical economics and political science to conceive and program societies of artificial agents (see, e.g., Shoham and Leyton-Brown, 2009). Social choice mechanisms have been proposed for collective decision-making in multiagent systems, stimulating research, e.g., on the computational properties of voting rules, the development of tractable procedures for strategic voting, or the development of approximation algorithms for computing the result of particularly hard rules.

Together with algorithmic game theory, computational social choice is currently one of the most well-represented research areas in conferences on multiagent systems.⁶ Many of the techniques introduced by COMSOC researchers would find a prime application in the development of voting avatars for agent-mediated collective decisions, such as those described in the short story above.

⁶See, e.g., the proceedings of the *International Conference on Autonomous Agents and Multiagent Systems*, AAMAS ([www.ifaamas.org](http://wwwifaamas.org)).

Strategic voting. Modelling human voters as perfectly rational agents is a useful simplification for obtaining intuitive theorems, but these assumptions limit significantly the applicability of these results (take the Gibbard-Satterthwaite Theorem as a classical example). Modelling bounded rationality is a big challenge in both AI and Economics, a problem that is absent, or less severe, in societies of artificial agents. Moreover, many of the techniques developed by COMSOC researchers for the analysis of strategic voting (see, e.g., Conitzer and Walsh, 2016; Faliszewski and Rothe, 2016) may find application test-beds in agent-mediated collective choices.

Machine learning. A voting avatar needs to be able to learn a voter's views and preferences from a set of voters' choices, that can moreover be evolving over time, and be more or less consistent with a set of existing views already present in society. This certainly is a challenging problem to frame, and one for which large personalised datasets are hard to construct. The use of machine learning techniques is still widely unexplored in social choice, and existing papers are mostly focused on creating novel voting methods (see, e.g., Xia, 2013). A fruitful starting point may be the conception of decision-support systems based on past voting behaviour.

Iterative voting. Voting methods that have been discarded as too complex or too unintuitive for human voters could be used successfully by artificial agents that act as proxies for voters. One example is iterative voting, in which repeated elections are staged in search for a more consensual voting outcome (see, e.g., Meir, 2017, for a survey). The use of voting avatars takes the burden of intensive communication away from the voter, and reinforcement learning techniques could be tested in this setting to obtain novel voting strategies and rules (Airiau et al., 2017).

Combinatorial voting and judgment aggregation. By making high-frequency collective decisions possible, the development of voting avatars will be confronted with a large and interconnected space of alternative choices. This would require compact representations for voters' views and preferences, such as those developed in the area of combinatorial voting (Lang and Xia, 2016), as well as a detailed understanding of the counterintuitive results and paradoxical situations that arise when aggregating them, such as those settings analysed by the theory of judgment aggregation (Endriss, 2016).

Argumentation. Political views and voters' behaviour are the result of discussions and debates, both in the social neighbourhood of a voter and in society in general. The field of (computational) argumentation is a vast enterprise ranging from the elicitation of arguments from natural language, to the identification of winning arguments in a debate between multiple agents. Theoretical research in this field (see, e.g., Rahwan and Simari, 2009) is recently being complemented by technologies that can form the basis of personal assistants for the construction of political views.

Social influence. A voting avatar will be scouting for opinions presented in influential newspapers or expressed by a designated set of influential voting avatars. Computational-friendly models of influence need to be developed, perhaps based on the

extensive research on trust and reputation in multiagent systems (Sabater and Sierra, 2005). The relation between social choices and social networks have recently been investigated by a number of papers (see, e.g., Grandi, 2017, for a survey)

4 Conclusions and Challenges

We may be quite far from a society in which collective choices of all sorts are taken every day in large numbers by means of artificial agents that vote on our behalf. However, in this paper I argue that the implementation of similar “voting avatars” is a viable solution to the lack of voter participation which affects e-democracy applications, and would moreover constitute a prime application for many of the techniques developed in computational social choice and multiagent systems in general. A first step in this direction could be the development of decision-aid systems, which help voters construct their political views and opinions on current issues, building on existing (simple) technologies such as Vote Match⁷ in UK, or Stemwijzer⁸ in the Netherlands.

There are, however, a number of scientific challenges that need to be tackled. First, representing a voter’s view may require combining techniques from computational knowledge representation with complex models of voting behaviour coming from political science and sociology. Second, data protection and security will be a key aspect of developing artificial agents acting on behalf of humans. However, this problem is shared by virtually all electronic applications of voting. A related problem is vote buying and vote influence in general, given the availability of data about voters’ views and votes that might be generated by moving to electronic platforms for collective decision-making. Solutions may be sought in cryptography algorithms, which may be easier to apply to societies of artificial agents rather than human ones (see, e.g., the recent work of Parkes et al. (2017)). Third, a functional voting avatar needs to be able to explain its behaviour to the voter in an understandable and convincing way. Explanation and interpretation are among the most important challenges for the deployment of artificial intelligence techniques in society, and a recent survey argues in favour of importing sociological theories of explanation to create transparent and trustworthy algorithms (Miller, 2017). Investigating how an artificial agent can explain simple voting behaviours in terms of the input preferences received may constitute a first step in this direction. Last, the design of voting avatars that can act as proxies for voters raises all sorts of ethical questions – from the level of autonomy of such an artificial proxy to the legal status of agent-mediated collective decisions – some of which are currently debated as general ethical issues related to the employment of artificial intelligence technologies in our societies. It turns out that voting itself might be a possible mean of developing collective views on similar ethical problems (Noothigattu et al., 2018).

⁷www.votematch.org.uk

⁸www.stemwijzer.nl

Acknowledgements

I am grateful to the participants and the speakers of the reading group on e-democracy at the University of Toulouse.⁹ Many thanks to Sylvie Doutre, Piotr Faliszewski, Davide Grossi, Arianna Novaro, Ashley Piggins, Ariel Procaccia, and Marija Slavkovic for their useful comments on previous versions of this paper.

References

- S. Airiau, U. Grandi, and F. S. Perotto. Learning agents for iterative voting. In *Proceedings of the 5th International Conference on Algorithmic Decision Theory (ADT)*, 2017.
- I. Asimov. Franchise. *If*, 1955.
- J. Behrens, A. Kistner, A. Nitsche, and B. Swierczek. *Principles of Liquid Feedback*. 2014.
- C. Boutilier and T. Lu. Probabilistic and utility-theoretic models in social choice: Challenges for learning, elicitation, and manipulation. In *Proceedings of the IJCAI-2011 Workshop on Social Choice and Artificial Intelligence*, 2011.
- S. Bouveret. Social choice and the web. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access, 2017.
- F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- V. Conitzer and T. Walsh. Barriers to manipulation in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- R. K. Dash, N. R. Jennings, and D. C. Parkes. Computational-mechanism design: A call to arms. *IEEE Intelligent Systems*, 18(6):40–47, 2003.
- A. Downs. *An Economic Theory of Democracy*. Harper, 1957.
- U. Endriss. Judgment aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- P. Faliszewski and J. Rothe. Control and bribery in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- J. Feigenbaum, D. C. Parkes, and D. M. Pennock. Computational challenges in e-commerce. *Communications of the ACM*, 52:70–74, 2009.

⁹www.irit.fr/~Umberto.Grandi/teaching/directdem

- R. Freedman, J. S. Borg, W. Sinnott-Armstrong, J. Dickerson, and V. Conitzer. Adapting a kidney exchange algorithm to align with human values. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- U. Grandi. Social choice on social networks. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access, 2017.
- J. Green-Armytage. Direct voting and proxy voting. *Constitutional Political Economy*, 26(2):190–220, 2015.
- D. Helbing, B. S. Frey, G. Gigerenzer, E. Hafen, M. Hagner, Y. Hofstetter, J. van den Hoven, R. V. Zicari, and A. Zwitter. Will democracy survive big data and artificial intelligence. *Scientific American*, February 2017.
- J. Lang and L. Xia. Voting in combinatorial domains. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- R. Meir. Iterative voting. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access, 2017.
- T. Miller. Explanation in artificial intelligence: Insights from the social sciences. *CoRR*, abs/1706.07269, 2017. URL <http://arxiv.org/abs/1706.07269>.
- R. Noothigattu, S. N. S. Gaikwad, E. Awad, S. Dsouza, I. Rahwan, P. Ravikumar, and A. D. Procaccia. A voting-based system for ethical decision making. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- D. C. Parkes. On AI, markets and machine learning. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2017.
- D. C. Parkes, P. Tylkin, and L. Xia. Thwarting vote buying through decoy ballots. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2017.
- I. Rahwan and G. R. Simari. *Argumentation in Artificial Intelligence*. Springer, 2009.
- J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artificial Intelligence Review*, 24(1):33–60, Sep 2005.
- N. Shah. Optimal social decision making. In *Proceedings of the 16th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2017.
- Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009.
- C. Sierra. Agent-mediated electronic commerce. *Autonomous Agents and Multi-Agent Systems*, 9(3):285–301, 2004.
- L. Xia. Designing social choice mechanisms using machine learning. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2013.