

Computational Social Choice: Research & Development*

Blue Sky Ideas Track

Dorothea Baumeister
HS Bund, Federal University of
Applied Administrative Science
Brühl, Germany
dorothea.baumeister@hsbund.de

Sylvain Bouveret
Univ. Grenoble Alpes, CNRS, Inria,
Grenoble INP, LIG
Grenoble, France
sylvain.bouveret@imag.fr

Martin Lackner
Univ. of Applied Sciences St. Pölten
St. Pölten, Austria
martin.lackner@ustp.at

Ratip Emin Berker
Carnegie Mellon University
Pittsburgh, United States of America
rberker@cs.cmu.edu

Andreas Darmann
University of Graz
Graz, Austria
andreas.darmann@uni-graz.at

Jérôme Lang
CNRS, LAMSADE, Université
Paris-Dauphine, PSL
Paris, France
jerome.lang@dauphine.fr

Arianna Novaro
Université Paris 1
Panthéon-Sorbonne, CNRS, CES
Paris, France
arianna.novaro@univ-paris1.fr

Niclas Boehmer
Hasso Plattner Institute,
University of Potsdam
Potsdam, Germany
Niclas.Boehmer@hpi.de

Piotr Faliszewski
AGH University of Krakow
Kraków, Poland
faliszew@agh.edu.pl

Nicholas Mattei
Tulane University
New Orleans, United States of
America
nsmattei@tulane.edu

ABSTRACT

Computational social choice (COMSOC) studies principled ways to aggregate conflicting individual preferences into collective decisions. In this paper, we call for an increased effort towards *Computational Social Choice: Research & Development (COMSOC-R&D)*, a problem-driven research agenda that explicitly aims to design, implement, and test collective decision-making systems in the real world. We articulate the defining features of COMSOC-R&D, argue for its value, and discuss various roadblocks and possible solutions.

KEYWORDS

Computational social choice; collective decision-making; real-world deployment; research & development (R&D)

ACM Reference Format:

Dorothea Baumeister, Ratip Emin Berker, Niclas Boehmer, Sylvain Bouveret, Andreas Darmann, Piotr Faliszewski, Martin Lackner, Jérôme Lang, Nicholas Mattei, and Arianna Novaro. 2026. Computational Social Choice: Research & Development: Blue Sky Ideas Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*,

*The full version of our paper can be found at <https://arxiv.org/abs/2602.20074>.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/CHQP3419>

Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 6 pages. <https://doi.org/10.65109/CHQP3419>

1 A COMSOC-R&D AGENDA

The field of computational social choice (COMSOC) studies collective decision-making problems, where conflicting individual preferences must be aggregated into a collective outcome [11]. COMSOC offers a rich toolbox for achieving principled compromises, with applications ranging from participatory budgeting and fair resource allocation to group recommendations and deliberation.¹

Over the past 25 years, the COMSOC community has made remarkable progress in understanding the algorithmic and axiomatic foundations of collective decision-making problems, and COMSOC has become an integral part of the AAMAS community. However, partly due to its roots in mathematics, theoretical computer science, and economics, the field has remained predominantly focused on foundational, abstract-level questions, whereas studies that develop, deploy, and evaluate mechanisms in specific real-world application contexts remain the exception rather than the norm.

This strong theoretical basis positions COMSOC to confront pressing real-world challenges in collective choice. However, to

¹Our use of the term COMSOC does not include matching under preferences (which studies, e.g., school choice, course allocation, and kidney exchange). While these topics are considered part of COMSOC by some, applied efforts in these domains have their origins largely in independent communities and have been much more driven by concrete problems [49]; see also the monograph by Sönmez [48] describing best practices and challenges around policy-oriented (matching) market design. We therefore consider such works rather role models for a broader vision of COMSOC-R&D.

realize this potential, we need to complement COMSOC’s foundational work with more application-driven research that bridges theory and practice. Accordingly, this blue sky paper calls for a shift toward what we term *Computational Social Choice: Research & Development (COMSOC-R&D)*, a research agenda that views social choice not only as an analytical framework but also as an applied engineering discipline for designing, implementing, testing, and deploying collective decision-making systems in the real world.

1.1 What is COMSOC-R&D?

COMSOC-R&D aims to both apply existing social choice techniques and develop new ones in order to provide solutions for concrete, real-world problems. Rather than focusing primarily on developing new theoretical results, COMSOC-R&D emphasizes the engineering and translational research goals of applying this research to the real world to yield implementable solutions.² Accordingly, while a (standard) theoretical analysis might be part of it, a COMSOC-R&D project places its main emphasis on translating theory into practice by implementing proposed algorithms or protocols using an accessible interface, evaluating them empirically, and ideally deploying them in real-world contexts.

Running computational experiments or developing software implementations of a general methodology is not what we see as a proper COMSOC-R&D approach. The distinctive feature of COMSOC-R&D lies in its problem-driven orientation: every project should begin with a concrete problem in the real world and aim to progress through the entire cycle of design, implementation, testing, and application. This agenda draws inspiration from the distinction made between *project-based* and *product-based* engineering: while product-based initiatives are developed in anticipation of possible future demand, project-based efforts are initiated in response to an existing, concrete client’s need. Although the distinction between the two approaches is sometimes fuzzy, COMSOC-R&D particularly seeks to encourage more project-based efforts within COMSOC, which so far have remained comparatively rare.

1.2 Why do We Need More COMSOC-R&D?

Any situation involving a group of agents who need to choose from a set of alternatives is a collective decision-making problem that could be addressed using COMSOC tools. In general, and unlike other areas of computer science, people facing such collective choice problems (like where to go for lunch or who to elect as a leader) can think of an *ad hoc*, common sense solution. As a result, these decision-makers might not feel the need to turn to the scientific community. This problem is shared with other communities that look at multi-criteria decision making (however, in high-stakes settings, advisors do regularly get employed [52]). The resulting solutions for collective choice problems can often be sub-optimal, and developing or employing a better method requires scientific expertise that the COMSOC community can provide.

To illustrate this point, consider the example of participatory budgeting, where a municipality designates parts of its budget to be spent on projects proposed by citizens and selected according to a vote. In practice, the selection is often done by a simple greedy

rule, which generally funds the projects receiving the most votes yet is provably suboptimal in terms of both fairness and efficiency. Likewise, when the US National Science Foundation re-worked a mechanism to allocate telescope time, the resulting method was susceptible to strategic manipulation [37].

COMSOC-R&D work can thus bring numerous advantages to individuals, the greater AAMAS community, and society. First, it is an opportunity for actual societal impact by improving how collective decisions are taken; in fact, in many applications, it could introduce collective decision-making in settings that did not involve active social participation before. Moreover, each successful application can bring more attention to the full spectrum of problems that can be solved with COMSOC techniques. Given the current shortage of applied COMSOC work, additional COMSOC-R&D efforts can offer valuable inspiration and guidance to researchers and practitioners facing similar problems, helping to build a shared knowledge base and best practices. Consequently, successful COMSOC-R&D projects can bring substantial recognition within the community.

From a research perspective, working on specific real-world problems typically surfaces aspects of the problem that are not adequately captured by models studied in theory. COMSOC-R&D thus establishes a positive feedback loop between theory and practice: engineering abstract models to fit concrete settings reveals new constraints, objectives, and desiderata [34]—and also prevent misconceptions about the actual real-world relevance of some theoretical questions, or their importance in deployed systems.

1.3 Successful Examples of COMSOC-R&D

While COMSOC-R&D currently constitutes only a very small fraction of COMSOC research published at CS conferences, there are already several successful examples illustrating what we envision. We highlight three success stories that showcase different aspects of COMSOC-R&D. There are, of course, many more, including work on peer reviewing [36, 37, 47], designing conference timetables [46], allocating classrooms [29], managing resources for satellite missions [18, 32], and allocating food bundles [5].

Participatory Budgeting. The Method of Equal Shares (MES) [43, 44] was developed as a theoretical proof of concept demonstrating that in participatory budgeting, groups of voters can be represented proportionally to their size. A few cities recently adopted it, addressing concrete issues they were facing with the widely used greedy method, such as a well-coordinated group of citizens hijacking a large fraction of the budget. MES quickly turned into a COMSOC-R&D effort [38]. It is highly configurable, and different design choices can strongly affect its nature [20]. Tools have been developed to allow decision makers to experiment with the rule [45] and to present the results intuitively [8].

Sortition. Another recent success story and example of COMSOC-R&D is the development of methods and tools for composing citizens’ assemblies that are both representative and (individually) fair. The Panelot platform [42] has been developed to implement sortition-based methods that fulfill these two requirements. The underlying work, driven by concrete needs, has already been used on several occasions to create actual assemblies [23].

Nominated Proof-of-Stake Blockchains. An established application of multiwinner voting [21, 30]—and an example of successful

²We take inspiration from *translational research* in biomedicine, which aims to quickly move research into the real world with the help of a translational partner [15, 51].

COMSOC-R&D—is the nominated proof-of-stake system used by Polkadot [1, 13]. For this application, a tailored voting mechanism was developed, drawing on previous foundational work from the COMSOC community. The desired voting method had to comply with desiderata that had previously received little or no attention (e.g., over-representation, scalability, verifiability) and thus led to important general algorithmic developments [7, 12, 14, 22].

2 ROADBLOCKS AND POSSIBLE SOLUTIONS

Although there seems to be broad agreement within the community that COMSOC-R&D work is highly beneficial (see Section 1.2), such work remains comparatively rare (see Section 1.3). We discuss possible reasons, distinguishing challenges common to applied computer science research from issues that are more specific to COMSOC-R&D. The goal of this section is to raise awareness of potential roadblocks to COMSOC-R&D projects and to describe solution strategies, both at the individual and the community level.

2.1 Common Roadblocks

There is a rich literature on obstacles that academic computer science researchers face when trying to translate and apply their research [3, 4, 10, 17, 24, 33], including best practice guidelines for starting, conducting, and reviewing applied science and engineering work [16, 39], including work on social good [26]. We highlight three main roadblocks particularly salient for COMSOC.

Partnering. Building and sustaining links between researchers and real-world stakeholders requires an enabling network and substantial time. Misaligned timelines and values, limited mutual awareness, and the upfront work of establishing shared terminology and realistic expectations further complicate collaboration. At the community level, we can facilitate partnerships by creating venues and discussion forums where both academics and practitioners are regularly present; there are emerging efforts to build such systemic bridges, for example, the EDDY network [19]. Matchmaking platforms and forums like industry days have proven useful in other CS communities [25, 27], but they require continuous efforts. At the individual level, a common success strategy is to jointly search for grant opportunities, which creates buy-in from all parties.

Incentives. Researchers may doubt the ease of publishing academic articles of R&D-style work when theoretical novelty takes a back seat to engineering and deployment. Conventional evaluation rubrics often undervalue adaptation efforts, carefully reasoned design trade-offs, and negative or mixed results. Additionally, more engineering-focused work tends to attract fewer citations, as it often prioritizes context-specific design over generalizable methodology.

The incentive problem calls for community-wide action. One way to create space for such works is through designated tracks: some conferences, such as AAAI and IJCAI, already have both social impact tracks and deployed applications tracks; while there is not yet a strong tradition of COMSOC-R&D being published there, these tracks are well-aligned with our scope. At venues without dedicated tracks, a central challenge is that standard evaluation criteria are only partly applicable to COMSOC-R&D submissions, and reviewers may thus find it difficult to evaluate the quality of

COMSOC-R&D submissions; we propose in Section 3 a set of review criteria tailored to COMSOC-R&D research.

Operations. There is a huge gap between developing a proof-of-concept research software and a full-fledged, scalable tool [53, 54], which necessitates, e.g., developing an intuitive interface, managing security issues, and maintaining this infrastructure. Without special care for this invisible and undervalued work (within the academic ecosystem), a development project may fail. One particular hurdle is to ensure that stakeholders are able to continue their operations even after the research team moves on to other projects. A high-effort way of achieving this is by starting a spin-off company, but doing so is not always reasonable (see the Decision Analysis area of INFORMS for examples [28, 40]). Fortunately, there are often other solutions. For example, one city hired an external software company to integrate MES into its participatory budgeting system.

2.2 COMSOC-Specific Roadblocks

We now turn to more COMSOC-specific challenges; while some of these are shared by other areas, we find these challenges to be more salient for the COMSOC research community.

High-stakes Decisions. Many COMSOC applications are particularly high-stakes and closely scrutinized (e.g., democratic elections), which can erode researchers’ confidence. The lack of “perfect” solutions and the resulting need to balance competing normative criteria further increase the burden of justification and the risk of stakeholder push-back, a problem shared with many multi-criteria decisions [52]. A general mitigation strategy is to start with small problems in a local environment, building upon personal connections. For instance, the success of the Method of Equal Shares started with Wieliczka in Poland, a small city on the outskirts of Krakow. The mutual trust between the MES team and the city was based on previous personal contacts. Yet, even in this case, the city adopted MES only in an experimental ecology-oriented participatory budgeting exercise. However, after this initial success, other cities became more open to adopting MES. Even if no such obvious small-scale variant exists, at the level of broadly applied high-impact problems, there begins to be room to develop trial-run solutions, provided that expectations are managed accordingly. Further, at the community level, we can facilitate more “work-in-progress” forums for projects and ideas targeting real-world challenges, to provide low-threshold exchange mechanisms and to share responsibility.

Lack of Clear Best Practices. Visible role models, documented success stories, and codified best practices for COMSOC-R&D remain scarce. With every additional COMSOC-R&D paper, this problem will diminish. Still, there are more concrete efforts the community can undertake, such as highlighting success stories (as in Section 1.3), conducting meta-studies, publishing best practices, and, when writing about specific COMSOC-R&D projects, explicitly describing not only the results but also the process and lessons learned, so that others in the community can build on them (cf. Section 3). Taking these steps has led to success in other application-oriented parts of computer science [3], including software engineering [25].

Culture Gap. COMSOC researchers, accustomed to rigorous mathematical arguments and formalisms, are often less comfortable with

and less passionate about the rather complicated specifications, ambiguous requirements, and trade-off decisions inherent to real-world projects. Theoretical papers often address one specific issue at a time, whereas engineering work mixes several. Training settings would be highly valuable to address culture gaps and the associated lack of experience. One idea is a recurring challenge series (analogous to the PACE challenge in parameterized complexity [41]), in which a problem is posed to the community, and submitted mechanisms and ideas are evaluated by real-world stakeholders.

Status Quo Bias. Stakeholders may often be attached to existing, sometimes autocratic, procedures: as long as these are perceived as “working” and those in control are content with the outcomes, there is little incentive to replace them with more principled mechanisms. A natural response is outreach that clearly communicates the pitfalls and complexity of collective decision-making. Raising public awareness of specific mis-handlings of decision processes can also increase pressure for change, but typically comes with confrontation and must therefore be approached with care.

Lack of Trust. Non-academic stakeholders may view researchers as idealistic theoreticians who do not understand practical constraints and difficulties and who lack real-world experience. One approach to address their concerns is to show that the researchers can “speak the same language,” i.e., they are aware of various practical issues and possible problems. Building this competence requires effort and experience; this is another reason why starting with small-scale problems in the local environment is important. Here we can borrow from work in the human-computer interaction, design, and AI for social good communities, focusing on participatory design and methodologies of building long-term trust and engagement within communities [2, 6, 26].

3 MERITS OF COMSOC-R&D PAPERS AND PROPOSED REVIEWING GUIDELINES

As we have argued, COMSOC-R&D work faces several obstacles. One specific (but academically important) issue is the difficulty of publishing engineering-heavy work in a theory- and methods-inclined community. As a first step, we argue that COMSOC-R&D work provides *substantial benefits to the development of the field as a whole* (cf. Section 1.2). To highlight this value, we now outline aspects that are especially important when writing COMSOC-R&D papers, which should also be taken into account when evaluating and reviewing them. We do not expect every COMSOC-R&D paper to address every single one of these points.

Modeling and Motivation Any COMSOC-R&D work requires substantial effort in modeling a real-world scenario. The description of the model and the underlying rationale are a major part of the work. Since the application determines the model, it cannot be expected that these models are as ‘elegant’ as those typically studied in theory. Evaluation should therefore focus less on the model and its appeal in isolation and more on the application itself and the fidelity of the model *to that application*.

Literature An important feature of COMSOC-R&D work is its connection to the COMSOC literature. In practice, the proposed solution might only be loosely related to previous work. Still, discussing related work – where it was useful, where it failed – is

highly valuable: If the proposed solution is firmly grounded in previous (more theoretical) publications, it is a success for our community, and if not, this gap becomes a call for further research.

Empirical Assessment A central contribution of COMSOC-R&D work lies in evaluating proposed mechanisms in realistic settings. Comparisons with baseline methods (including the status quo) are valuable, in particular if they highlight the effect of approaching the problem from the perspective of collective decision-making.

Data Considerable effort has been put into collecting data sets suitable for COMSOC research [9, 20, 35]. Additional data sets from new applications are extremely valuable, especially as recent work [50] has shown that preference data can differ substantially depending on the domain.

Transparency Engineering work is inherently iterative. Papers should document the key design steps, discussions, and lessons learned during the project as valuable guidance for future COMSOC-R&D projects. Indeed, many small modeling choices and fine-tuning decisions must be made. While those cannot all be justified in depth, it is crucial to be transparent and explicit about them.

Reproducibility The full scope of COMSOC-R&D projects may not be reproducible. This can be due to physical impossibilities or contractual obligations. Nevertheless, authors should aim for reproducibility wherever possible. Providing an easily accessible interface or artifacts that others can build on is a strong plus.

Stakeholder Engagement COMSOC-R&D projects typically involve sustained collaboration with non-academic partners. Documenting this process – how requirements were elicited, how design decisions were made, and how feedback was incorporated – provides valuable guidance for future projects.

Implementation and Deployment The practical challenges of deploying a collective decision-making system (user interface design, integration with existing processes, computational optimization, etc.) are often substantial and rarely addressed in theoretical work. Detailed accounts of such deployment efforts are a distinctive contribution of COMSOC-R&D papers.

Incorporating these perspectives into the reviewing process does not necessarily require a change in reviewer guidelines. Rather, this is a call for consideration to the (wider) COMSOC community: COMSOC-R&D papers are valuable for the community, and hence should be judged by their merit rather than by how closely they resemble typical COMSOC papers. To give COMSOC-R&D papers an equal standing in the reviewing process, it has to be acknowledged that their merit differs from that of theory-focused papers.

4 A VISION FOR THE FUTURE

We envision a future where COMSOC-R&D papers will be equally represented at COMSOC venues, including AAMAS, alongside more classical, theory-focused papers. We believe that the COMSOC-R&D approach has great potential to bring concrete improvements in the way real-life decisions are made every day. Particularly promising application domains include *preferences-aware timetabling* (e.g., for university classes or work shifts; cf. [31]), *dynamic reallocation of objects under uncertainty* (e.g., for reassigning offices during a building renovation), and *fair public transport*. We invite researchers to pursue these challenges and call on the community to recognize and reward such efforts.

ACKNOWLEDGMENTS

This work was initiated in a work group during the Dagstuhl Seminar 25401 “Societal Impact of Computational Social Choice”.

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 101002854). Mattei was supported in part by NSF Awards IIS-RI-2134857 and IIS-RI-2339880 as well as the Harold L. and Heather E. Jurist Center of Excellence for Artificial Intelligence at Tulane University and the Tulane University Center of Excellence for Community-Engaged Artificial Intelligence (CEAI). Lang has received funding from ANR under grant PSAI-PRAIRIE, ANR-23-IACL-0008.



REFERENCES

- [1] Hanaa Abbas, Maurantonio Caprolu, and Roberto Di Pietro. 2022. Analysis of Polkadot: Architecture, Internals, and Contradictions. In *Proceedings of the IEEE International Conference on Blockchain (Blockchain)*.
- [2] Rediet Abebe, Solon Barocas, Jon M. Kleinberg, Karen Levy, Manish Raghavan, and David G. Robinson. 2020. Roles for computing in social change. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAT*)*.
- [3] Elena Agapie, Pamela Wisniewski, Kevin Siek, and Hazel Griffin. 2024. *Catalyzing Interdisciplinary Computing Research: Best Practices for Researchers*. Technical Report. Computing Research Association (CRA). https://cra.org/wp-content/uploads/2024/08/Catalyzing-Interdisciplinary-Computing-Research-Best-Practices-for-Researchers_August2024_FINAL.pdf
- [4] Nancy Amato, Ovidiu Daescu, Maria Gini, Georgia Hayes, Kevin Larson, Wei Li, Ming Lin, David Lopresti, Kevin Siek, Gene Spafford, and Benjamin Zorn. 2025. *Unique Considerations for Evaluating Computing Researchers*. Technical Report. Computing Research Association (CRA). <http://cra.org/cra/wp-content/uploads/2025/07/Unique-Considerations-for-Evaluating-Computing-Researchers.pdf>
- [5] Gerdus Benade and Aydin Alptekinoglu. 2023. Achieving Rawlsian Justice in Food Rescue. In *Proceedings of the ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*.
- [6] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Jason Gabriel, and Shakir Mohamed. 2022. Power to the people? Opportunities and challenges for participatory AI. In *Proceedings of the ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*, 1–8.
- [7] Niclas Boehmer, Markus Brill, Alfonso Cevallos, Jonas Gehrlein, Luis Sánchez-Fernández, and Ulrike Schmidt-Kraepelin. 2024. Approval-based committee voting in practice: A case study of (over-) representation in the Polkadot blockchain. In *Proceedings of the AAI Conference on Artificial Intelligence (AAI)*.
- [8] Niclas Boehmer, Piotr Faliszewski, Lukasz Janeczko, Dominik Peters, Grzegorz Pierczynski, Simon Schierreich, Piotr Skowron, and Stanislaw Szufa. 2024. Evaluation of Project Performance in Participatory Budgeting. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- [9] Niclas Boehmer and Nathan Schaar. 2023. Collecting, Classifying, Analyzing, and Using Real-World Ranking Data. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [10] Elizabeth Bondi, Lily Xu, Diana Acosta-Navas, and Jackson A. Killian. 2021. Envisioning Communities: A Participatory Approach Towards AI for Social Good. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES)*.
- [11] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of computational social choice*. Cambridge University Press.
- [12] Markus Brill, Rupert Freeman, Svante Janson, and Martin Lackner. 2024. Phragmén’s Voting Methods and Justified Representation. *Mathematical Programming* 203 (2024), 47–76.
- [13] Jeff Burdges, Alfonso Cevallos, Peter Czaban, Rob Habermeier, Syed Hosseini, Fabio Lama, Handan Kiling Alper, Ximin Luo, Fatemeh Shirazi, Alistair Stewart, and Gavin Wood. 2020. Overview of Polkadot and its Design Considerations. *arXiv preprint arXiv:2005.13456* (2020).
- [14] Alfonso Cevallos and Alistair Stewart. 2021. A verifiably secure and proportional committee election rule. In *Proceedings of the ACM Conference on Advances in Financial Technologies (AFT)*.
- [15] Francis S Collins. 2011. Reengineering translational science: the time is right. *Science translational medicine* 3, 90 (2011).
- [16] Computing Research Association. 2024. Best Practice Memos. <https://cra.org/resources/best-practice-memos/>.
- [17] Marian Daun, Jennifer Brings, Patricia Aluko Obe, and Bastian Tenbergen. 2023. An industry survey on approaches, success factors, and barriers for technology transfer in software engineering. *Software: Practice and Experience* 53, 7 (2023), 1496–1524.
- [18] James S Dyer and Ralph F Miles Jr. 1976. An actual application of collective choice theory to the selection of trajectories for the Mariner Jupiter/Saturn 1977 project. *Operations Research* 24, 2 (1976), 220–244.
- [19] EDDY Network. 2024. EDDY: European Digital Democracy Network. <https://www.eddy-network.eu/>.
- [20] Piotr Faliszewski, Jaroslaw Flis, Dominik Peters, Grzegorz Pierczynski, Piotr Skowron, Dariusz Stolicki, Stanislaw Szufa, and Nimrod Talmon. 2023. Participatory Budgeting: Data, Tools and Analysis. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- [21] Piotr Faliszewski, Piotr Skowron, Arkadij Slinko, and Nimrod Talmon. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access, Chapter 2, 27–47.
- [22] Luis Sánchez Fernández, Norberto Fernández García, Jesús A. Fisteus, and Markus Brill. 2024. The maximin support method: an extension of the D’Hondt method to approval-based multiwinner elections. *Mathematical Programming* 203, 1 (2024), 107–134.
- [23] Bailey Flanigan, Paul Gözl, Anupam Gupta, Brett Hennig, and Ariel D. Procaccia. 2021. Fair algorithms for selecting citizens’ assemblies. *Nature* 596, 7873 (2021), 548–552.
- [24] Luciano Floridi, Josh Cowsls, Thomas C. King, and Mariarosaria Taddeo. 2020. How to Design AI for Social Good: Seven Essential Factors. *Science and Engineering Ethics* 26, 3 (2020), 1771–1796.
- [25] Vahid Garousi, Kai Petersen, and Baris Ozkan. 2016. Challenges and best practices in industry-academia collaborations in software engineering: A systematic literature review. *Information and Software Technology* 79 (2016), 106–127.
- [26] Gregory D. Hager, Ann W. Drobnis, Fei Fang, Rayid Ghani, Amy Greenwald, Terah Lyons, David C. Parkes, Jason Schultz, Suchi Saria, Stephen F. Smith, and Milind Tambe. 2019. *Artificial Intelligence for Social Good*. Technical Report. Computing Community Consortium (CCC). <https://arxiv.org/pdf/1901.05406>
- [27] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé III, Miro Dudik, and Hanna Wallach. 2019. Improving fairness in machine learning systems: What do industry practitioners need?. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI)*.
- [28] INFORMS. 2020. 2020 Decision Analysis Software Survey. <https://pubsonline.informs.org/magazine/orms-today/2020-decision-analysis-software-survey>.
- [29] David Kurokawa, Ariel D Procaccia, and Nisarg Shah. 2018. Leximin allocations in the real world. *ACM Transactions on Economics and Computation (TEAC)* 6, 3-4 (2018), 1–24.
- [30] Martin Lackner and Piotr Skowron. 2023. *Multi-Winner Voting with Approval Preferences*. Springer International Publishing.
- [31] Marie-Louise Lackner. 2025. Schedules Need to be Fair Over Time. In *Proceedings of Digital Humanism: First Interdisciplinary Science and Research Conference, (DIGHUM)*.
- [32] Michel Lemaître, Gérard Verfaillie, Hélène Fargier, Jérôme Lang, Nicolas Bataille, and Jean-Michel Lachiver. 2003. Equitable allocation of earth observing satellites resources. In *Proceedings of the 5th ONERA-DLR Aerospace Symposium*.
- [33] Aditya Majumdar, Wenbo Zhang, Kashvi Prawal, and Amulya Yadav. 2025. The Hardness of Achieving Impact in AI for Social Impact Research: A Ground-Level View of Challenges & Opportunities. *arXiv preprint arXiv:2506.14829* (2025).
- [34] Nicholas Mattei. 2021. Closing the loop: Bringing humans into empirical computational social choice and preference reasoning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 5169–5173.
- [35] Nicholas Mattei and Toby Walsh. 2013. PrefLib: A Library for Preferences <http://www.preflib.org>. In *Proceedings of the International Conference on Algorithmic Decision Theory (ADT)*.
- [36] Reshef Meir, Jérôme Lang, Julien Lesca, Nicholas Mattei, and Natan Kaminsky. 2021. A market-inspired bidding scheme for peer review paper assignment. In *Proceedings of the AAI Conference on Artificial Intelligence (AAI)*.
- [37] Michael R Merrifield and Donald G Saari. 2009. Telescope time without tears: a distributed approach to peer review. *Astronomy & Geophysics* 50, 4 (2009), 4–16.
- [38] Method of Equal Shares. 2024. The Method of Equal Shares. <https://equalshares.net/>.
- [39] Nature Communications. 2024. Applied Science Research: Submission Guidelines. <https://www.nature.com/ncomms/submit/applied-science-research>.
- [40] David L Olson. 1997. Decision aids for selection problems. *Journal of the Operational Research Society* 48, 5 (1997), 541–542.
- [41] PACE Challenge. 2024. PACE: Parameterized Algorithms and Computational Experiments Challenge. <https://pacechallenge.org/>.
- [42] Panelot. 2024. Panelot: A Platform for Sortition. <https://panelot.org/>.
- [43] Dominik Peters, Grzegorz Pierczynski, and Piotr Skowron. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*.

- [44] Dominik Peters and Piotr Skowron. 2020. Proportionality and the Limits of Welfarism. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.
- [45] Pref.Tools. 2024. Pabutools: Tools for Participatory Budgeting. <https://pref.tools/pabutools/>.
- [46] Simon Rey and Ulle Endriss. 2024. *A Note on How to Build a Conference Schedule: The Case of ECAI-2024*. Technical Note X-2024-03. ILLC, University of Amsterdam. <https://doi.org/10.5281/zenodo.14019044>
- [47] Inbal Rozenzweig, Reshef Meir, Nicholas Mattei, and Ofra Amir. 2023. Mitigating skewed bidding for conference paper assignment. *arXiv preprint arXiv:2303.00435* (2023).
- [48] Tayfun Sönmez. 2025. Minimalist Market Design: A Framework for Economists with Policy Aspirations. *arXiv preprint arXiv:2401.00307* (2025).
- [49] Tayfun Sönmez and M. Utku Ünver. 2025. Matching under Non-transferable Utility: Applications. In *Handbook of the Economics of Matching*, Yeon-Koo Che, Pierre-André Chiappori, and Bernard Salanié (Eds.). Vol. 2. Elsevier, Chapter 3.
- [50] Stanislaw Szufa, Niclas Boehmer, Robert Bredereck, Piotr Faliszewski, Rolf Niedermeier, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2025. Drawing a map of elections. *Artificial intelligence* 343 (2025), 104332.
- [51] Elaine Wethington and Louis D Burgio. 2015. Translational research on caregiving: Missing links in the translation process. In *Family caregiving in the new normal*, Joseph E. Gaugler and Robert L. Kane (Eds.). Elsevier, 193–209.
- [52] Wikipedia contributors. 2025. Multiple-criteria decision analysis. https://en.wikipedia.org/wiki/Multiple-criteria_decision_analysis Revision date: 25 August 2025; Accessed: 8 December 2025.
- [53] Greg Wilson, Dhavide A Aruliah, C Titus Brown, Neil P Chue Hong, Matt Davis, Richard T Guy, Steven HD Haddock, Kathryn D Huff, Ian M Mitchell, Mark D Plumbley, et al. 2014. Best practices for scientific computing. *PLoS biology* 12, 1 (2014), e1001745.
- [54] Greg Wilson, Jennifer Bryan, Karen Cranston, Justin Kitzes, Lex Nederbragt, and Tracy K Teal. 2017. Good enough practices in scientific computing. *PLoS computational biology* 13, 6 (2017), e1005510.