Collective Decision Making via Automated Reasoning

Doctoral Consortium

Ari Conati
University of Helsinki
Helsinki, Finland
ari.conati@helsinki.fi

ABSTRACT

Collective decision-making tasks, such as voting, matching, and resource allocation, are frequently encountered in multi-agent scenarios where a consensus is sought for based on the—often conflicting—preferences of individual agents. Deciding if a consensus can be reached and finding such consensus give rise to computationally hard decision and optimization problems, characterized by NP-completeness or even beyond-NP complexity. This complexity poses significant challenges for developing practical exact algorithms. At the same time, advances in automated logical reasoning techniques, such as Boolean satisfiability solvers, their extensions to higher-level constraints, and optimization have proven successful for capturing and solving a wide range of computationally hard real-world problems. My doctoral research harnesses automated logical reasoning for developing novel types of practical, exact algorithms for computational social choice scenarios.

KEYWORDS

computational social choice; judgment aggregation; fair allocation; Boolean satisfiability

ACM Reference Format:

Ari Conati. 2025. Collective Decision Making via Automated Reasoning: Doctoral Consortium. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

1 INTRODUCTION

Collective decision-making tasks, such as voting, matching, and resource allocation, are frequently encountered in multi-agent scenarios where a consensus is sought for based on the—often conflicting—preferences of individual agents [13]. Consensus is typically characterized by maximally satisfying criteria such as fairness or social welfare. Deciding if a consensus can be reached and finding such consensus give rise to computationally hard decision and optimization problems, characterized by NP-completeness or even beyond-NP complexity [11, 20, 21, 30].

This complexity poses significant challenges for developing practical algorithms in computational social choice (COMSOC) [13]. Interestingly, however, many of the problem settings considered in COMSOC—including, for instance, judgment aggregation [29] and fair allocation [2, 10, 38]— are naturally expressed in logic-oriented representations. At the same time, advances in automated logical reasoning techniques, such as Boolean satisfiability (SAT) solvers, their extensions to higher-level constraints (such as satisfiability modulo theories [5] and pseudo-Boolean reasoning [27]) and optimization (maximum satisfiability and extensions [39]) have proven a key technology for capturing and solving a wide range of computationally hard real-world problems. Critically, by enabling incremental usage under assumptions [26, 41], modern SAT solvers can even naturally capture decision and optimization procedures complete for the second level of the polynomial hierarchy with high efficiency.

In my PhD research, my goal is to develop efficient, practical algorithms for various relevant, computationally hard problems in COMSOC and to implement them open source. For beyond-NP problems, a polynomial-size SAT encoding presumably does not exist by complexity-theoretic argument, so this entails developing novel procedures based on strategic, iterative calls to SAT solvers. Thus far, I have published work on judgment aggregation [16, 18] and fair allocation [17], with plans to extend to other central problems within COMSOC.

In this extended abstract, I provide an overview of my doctoral research results thus far and outline my future plans toward completing my PhD.

2 JUDGMENT AGGREGATION

Judgment aggregation [29] is a general, logical framework which captures scenarios where agents reach a consensus by aggregating their preferences, judgments, or beliefs by social choice mechanisms. In judgment aggregation we consider a set of issues, each of which agents can either accept or reject, subject to an integrity constraint corresponding to logical constraints over the issues. A judgment aggregation rule then identifies a set of optimal collective judgment sets based on the opinions of the agents, similarly subject to an integrity constraint.

In my research so far I have considered both outcome determination [18] and strategic behavior [16] in judgment aggregation. The implementation of all of the algorithms described in this section is available open source at https://bitbucket.org/coreo-group/satcha.

2.1 Outcome Determination

Arguably the most central task in judgment aggregation is *outcome determination* [30]; that is, determining whether a given subset of the agenda is accepted under a given judgment aggregation rule. Outcome determination generalizes the winner determination [35] problem in voting, where the task is to determine if a given alternative is a winner of an election. It has been shown that outcome

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

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

determination is computationally NP-hard, and even complete for the second level of the polynomial hierarchy, under several proposed judgment aggregation rules [30].

In [18], I developed practical algorithms for outcome determination under a wide range of central judgment aggregation rules, namely the Kemeny [37, 43, 45], Slater [37, 43, 45], Young [37], Dodgson [43], MaxHamming [37], Reversal scoring [23], Condorcet [30, 37, 45], Ranked agenda [37], and Leximax [31, 44] rules. Under these rules, outcome determination is Θ_2^p-complete, Σ_2^p-complete, or Δ_2^p-complete, depending on the rule [30]: I therefore employed various algorithmic techniques adhering to the computational complexity of the problem, including direct approaches using MaxSAT encodings and iterative procedures based on the counterexample-guided abstraction refinement (CEGAR) paradigm [15]. I implemented and empirically evaluated the algorithms on preference data from the PrefLib [42] reference library containing real-world voting data arising from, e.g., elections and surveys, demonstrating scalability significantly beyond the reach of alternative approaches.

2.2 Manipulation and Bribery

An additional central area of study in judgment aggregation concerns strategic behavior [6, 19, 32, 36]. Strategic behavior is generally undesirable: for instance, it is nonideal if an agent is able to manipulate the outcome of a voting process. It is therefore important to understand the practical viability of identifying such strategies. In terms of forms of strategic behavior, I have looked at *manipulation* [7, 20], where the task is to determine if an agent can secure a desired outcome through indicating an insincere judgment set, and *bribery* [7, 20], where instead a third party enforces their preferred outcome by bribing several individual agents involved in the decision process.

In our research [16], we have extended previous complexity results showing manipulation and bribery to be Σ_2^p-complete under the Kemeny rule [20], demonstrating that Σ_2^p-completeness similarly holds for the Slater, MaxHamming, Young, and Dodgson rules. Towards capturing these problems, I developed CEGAR-based algorithmic approaches which make use of iterative calls to a MaxSAT solver. I implemented and evaluated these algorithms on voting data, demonstrating their practical viability despite the high theoretical complexity barriers.

3 FAIR ALLOCATION

In fair allocation of indivisible goods [2, 10, 38], the task is to distribute (bundles of) discrete items amongst a set of agents based on their individual preferences. Fair allocation is a widely-studied problem in COMSOC which arises from various real-world multi-agent settings, including, for instance, dividing computational resources in clusters [34], assigning courses to university students [14], and food distribution [1].

A central, desirable fairness property is *envy-freeness* [33, 46]: that is, an allocation is envy-free if each agent is at least as happy with their received bundle as they would be with that received by any other agent; i.e., they do not envy any other agent. An envy-free allocation is trivially achieved by withholding all items: however, this is typically not a satisfactory allocation. Therefore, envy-freeness is usually combined with an efficiency notion such as *completeness*, which requires that every item is allocated to an agent. Finding an allocation which is envy-free and complete, however, is already NP-hard [40]. For more refined notions of efficiency, the task is even harder. For instance, under *Pareto-efficiency*, we seek allocations where it is not possible to reallocate items in a way that makes some agent better off without making a different agent worse off. The task of determining the existence of a Pareto-efficient, envy-free allocation has been shown to be Σ_2^p-complete [11, 21].

Various types of agent preferences have been studied in the context of fair allocation. In settings where agents have *dichotomous* preferences [9, 12, 25], agents either approve or disapprove of any given bundle, but have no preference between bundles from either category. In these settings, preferences may be represented as propositional formulas. Thus, the application of SAT solvers is natural in this context.

In work published in the proceedings of AAMAS 2025 [17], I have designed and implemented iterative SAT-based algorithms for deciding the existence of envy-free, Pareto-efficient allocations under dichotomous preferences. I further extended these algorithms to minimizing envy in cases where an envy-free allocation does not exist, considering different notions of total envy. As far as we know, these are the first exact algorithmic approaches for problem variants in fair allocation which are complete for the second level of the polynomial hierarchy. I implemented and evaluated these approaches, demonstrating scalability up to hundreds of agents.

4 FUTURE PLANS

In fair allocation, I look to further consider other types of preferences such as additive preferences [40], where agents assign numerical values to each of the items. Encoding numerical weights in propositional logic is somewhat cumbersome, such that it may be necessary to consider other declarative paradigms such as integer programming or pseudo-boolean solving.

Additionally, our algorithmic approaches naturally extend to hedonic games [3], which model scenarios where agents are to form coalitions based on their individual preferences. I further aim to capture widely-studied solution concepts in hedonic games such as core stability and strict-core stability [4, 8, 24] through declarative methods. Different preference representations such as "friends and enemies" [24] and hedonic coalition nets [28] are also worth considering.

Finally, suitable benchmark datasets (i.e., diverse and ideally derived from real-world scenarios) are lacking for many problems within COMSOC. Simultaneously, we and others have observed that it can be non-trivial to estimate the difficulty of specific COMSOC instances, and that it does not necessarily correlate with the size of the instance (e.g., the number of agents) [17, 22]. As empirical evaluation is critical to algorithm development, we hope to further investigate methods for generating benchmarks that have similar structural and computational properties as typical real-world instances.

ACKNOWLEDGMENTS

Work financially supported by the Finnish Doctoral Program Network in Artificial Intelligence AI-DOC (VN/3137/2024-OKM-6) and the Academy of Finland (grants 347588 and 356046).

REFERENCES

[1] Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. 2015. Online Fair Division: Analysing a Food Bank Problem. In *IJCAI 2015*, Qiang Yang and Michael J. Wooldridge (Eds.). AAAI Press, 2540–2546.

[2] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. 2023. Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.* 322 (2023), 103965.

[3] Haris Aziz and Rahul Savani. 2016. Hedonic Games. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, 356–376.

[4] Suryapratim Banerjee, Hideo Konishi, and Tayfun Sönmez. 2001. Core in a simple coalition formation game. *Soc. Choice Welf.* 18, 1 (2001), 135–153.

[5] Clark W. Barrett, Roberto Sebastiani, Sanjit A. Seshia, and Cesare Tinelli. 2021. Satisfiability Modulo Theories. In *Handbook of Satisfiability - Second Edition*, Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh (Eds.). Frontiers in Artificial Intelligence and Applications, Vol. 336. IOS Press, 1267–1329.

[6] J. J. Bartholdi III, C. A. Tovey, and M. A. Trick. 1989. The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare* 6, 3 (1989), 227–241.

[7] Dorothea Baumeister, Gábor Erdélyi, Olivia Johanna Erdélyi, and Jörg Rothe. 2015. Complexity of manipulation and bribery in judgment aggregation for uniform premise-based quota rules. *Math. Soc. Sci.* 76 (2015), 19–30.

[8] Anna Bogomolnaia and Matthew O. Jackson. 2002. The Stability of Hedonic Coalition Structures. *Games Econ. Behav.* 38, 2 (2002), 201–230.

[9] Anna Bogomolnaia, Hervé Moulin, and Richard Stong. 2005. Collective choice under dichotomous preferences. *J. Econ. Theory* 122, 2 (2005), 165–184.

[10] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. 2016. Fair Allocation of Indivisible Goods. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, 284–310.

[11] Sylvain Bouveret and Jérôme Lang. 2008. Efficiency and Envy-freeness in Fair Division of Indivisible Goods: Logical Representation and Complexity. *J. Artif. Intell. Res.* 32 (2008), 525–564.

[12] Florian Brandl and Dominik Peters. 2022. Approval voting under dichotomous preferences: A catalogue of characterizations. *J. Econ. Theory* 205 (2022), 105532.

[13] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice.* Cambridge University Press.

[14] Eric Budish. 2010. The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. In *BQGT 2010*, Moshe Dror and Greys Sosic (Eds.). ACM, 74:1.

[15] Edmund M. Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. 2003. Counterexample-guided abstraction refinement for symbolic model checking. *J. ACM* 50, 5 (2003), 752–794.

[16] Ari Conati, Andreas Niskanen, Ronald de Haan, and Matti Järvisalo. 2024. Complexity Results and Algorithms for Manipulation and Bribery in Judgment Aggregation. In *ECAI 2024*, Ulle Endriss, Francisco S. Melo, Kerstin Bach, Alberto José Bugarín Diz, Jose Maria Alonso-Moral, Senén Barro, and Fredrik Heintz (Eds.). IOS Press, 3597–3604.

[17] Ari Conati, Andreas Niskanen, Ronald de Haan, and Matti Järvisalo. 2025. Computing Efficient and Envy-Free Allocations under Dichotomous Preferences using SAT. In *AAMAS 2025, to appear.*

[18] Ari Conati, Andreas Niskanen, and Matti Järvisalo. 2024. Declarative Approaches to Outcome Determination in Judgment Aggregation. *J. Artif. Intell. Res.* 81 (2024), 793–836.

[19] Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. 2007. When are elections with few candidates hard to manipulate? *J. ACM* 54, 3 (2007), 14.

[20] Ronald de Haan. 2017. Complexity Results for Manipulation, Bribery and Control of the Kemeny Judgment Aggregation Procedure. In *AAMAS 2017*, Kate Larson, Michael Winikoff, Sanmay Das, and Edmund H. Durfee (Eds.). ACM, 1151–1159.

[21] Bart de Keijzer, Sylvain Bouveret, Tomas Klos, and Yingqian Zhang. 2009. On the Complexity of Efficiency and Envy-Freeness in Fair Division of Indivisible Goods with Additive Preferences. In *ADT 2009 (LNCS, Vol. 5783)*, Francesca Rossi and Alexis Tsoukiàs (Eds.). Springer, 98–110.

[22] John P. Dickerson, Jonathan R. Goldman, Jeremy Karp, Ariel D. Procaccia, and Tuomas Sandholm. 2014. The Computational Rise and Fall of Fairness. In *AAAI 2014*, Carla E. Brodley and Peter Stone (Eds.). AAAI Press, 1405–1411.

[23] Franz Dietrich. 2014. Scoring rules for judgment aggregation. *Social Choice and Welfare* 42, 4 (2014), 873–911.

[24] Dinko Dimitrov, Peter Borm, Ruud Hendrickx, and Shao Chin Sung. 2006. Simple Priorities and Core Stability in Hedonic Games. *Soc. Choice Welf.* 26, 2 (2006), 421–433.

[25] Conal Duddy. 2015. Fair sharing under dichotomous preferences. *Math. Soc. Sci.* 73 (2015), 1–5.

[26] Niklas Eén and Niklas Sörensson. 2003. Temporal induction by incremental SAT solving. *Electronic Notes in Theoretical Computer Science* 89, 4 (2003), 543–560.

[27] Jan Elffers and Jakob Nordström. 2018. Divide and Conquer: Towards Faster Pseudo-Boolean Solving. In *IJCAI 2018*, Jérôme Lang (Ed.). ijcai.org, 1291–1299.

[28] Edith Elkind and Michael J. Wooldridge. 2009. Hedonic coalition nets. In *AAMAS 2009*, Carles Sierra, Cristiano Castelfranchi, Keith S. Decker, and Jaime Simão Sichman (Eds.). IFAAMAS, 417–424.

[29] Ulle Endriss. 2016. Judgment Aggregation. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, 399–426.

[30] Ulle Endriss, Ronald de Haan, Jérôme Lang, and Marija Slavkovik. 2020. The Complexity Landscape of Outcome Determination in Judgment Aggregation. *Journal of Artificial Intelligence Research* 69 (2020), 687–731.

[31] Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. 2014. Counting votes for aggregating judgments. In *AAMAS 2014*, Ana L. C. Bazzan, Michael N. Huhns, Alessio Lomuscio, and Paul Scerri (Eds.). IFAAMAS/ACM, 1177–1184.

[32] Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. 2009. How Hard Is Bribery in Elections? *J. Artif. Intell. Res.* 35 (2009), 485–532.

[33] Duncan K. Foley. 1967. Resource allocation and the public sector. *Yale Economic Essays* 7 (1967), 45–98.

[34] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. 2011. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types. In *NSDI 2011*, David G. Andersen and Sylvia Ratnasamy (Eds.). USENIX Association, 323–336.

[35] Edith Hemaspaandra, Lane A. Hemaspaandra, and Jörg Rothe. 1997. Exact Analysis of Dodgson Elections: Lewis Carroll's 1876 Voting System is Complete for Parallel Access to NP. In *ICALP (Lecture Notes in Computer Science, Vol. 1256)*, Pierpaolo Degano, Roberto Gorrieri, and Alberto Marchetti-Spaccamela (Eds.). Springer, 214–224.

[36] Edith Hemaspaandra, Lane A. Hemaspaandra, and Jörg Rothe. 2007. Anyone but him: The complexity of precluding an alternative. *Artif. Intell.* 171, 5-6 (2007), 255–285.

[37] Jérôme Lang, Gabriella Pigozzi, Marija Slavkovik, and Leendert W. N. van der Torre. 2011. Judgment aggregation rules based on minimization. In *TARK 2011*, Krzysztof R. Apt (Ed.). ACM, 238–246.

[38] Jérôme Lang and Jörg Rothe. 2016. Fair Division of Indivisible Goods. In *Economics and Computation, An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*, Jörg Rothe (Ed.). Springer, 493–550.

[39] Chu Min Li and Felip Manyà. 2021. MaxSAT, Hard and Soft Constraints. In *Handbook of Satisfiability - Second Edition*, Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh (Eds.). Frontiers in Artificial Intelligence and Applications, Vol. 336. IOS Press, 903–927.

[40] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *ÉC 2004*, Jack S. Breese, Joan Feigenbaum, and Margo I. Seltzer (Eds.). ACM, 125–131.

[41] João Marques-Silva, Inês Lynce, and Sharad Malik. 2021. Conflict-Driven Clause Learning SAT Solvers. In *Handbook of Satisfiability - Second Edition*, Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh (Eds.). Frontiers in Artificial Intelligence and Applications, Vol. 336. IOS Press, 133–182.

[42] Nicholas Mattei and Toby Walsh. 2013. PrefLib: A Library for Preferences http://www.preflib.org. In *ADT 2013 (Lecture Notes in Computer Science, Vol. 8176)*, Patrice Perny, Marc Pirlot, and Alexis Tsoukiàs (Eds.). Springer, 259–270.

[43] Michael K. Miller and Daniel N. Osherson. 2009. Methods for distance-based judgment aggregation. *Social Choice and Welfare* 32, 4 (2009), 575–601.

[44] Klaus Nehring and Marcus Pivato. 2019. Majority rule in the absence of a majority. *Journal of Economic Theory* 183 (2019), 213–257.

[45] Klaus Nehring, Marcus Pivato, and Clemens Puppe. 2014. The Condorcet set: Majority voting over interconnected propositions. *Journal of Economic Theory* 151 (2014), 268–303.

[46] Hal R. Varian. 1974. Equity, envy, and efficiency. *Journal of Economic Theory* 9, 1 (1974), 63–91.