

Multi-Agent Meeting Scheduling: A Negotiation Perspective

Bram M. Renting
Leiden University
Delft University of Technology
The Netherlands
b.m.renting@liacs.leidenuniv.nl

Holger H. Hoos
RWTH Aachen University
Leiden University
Germany & The Netherlands
hh@aim.rwth-aachen.de

Catholijn M. Jonker
Delft University of Technology
Leiden University
The Netherlands
c.m.jonker@tudelft.nl

ABSTRACT

In multi-agent systems (MAS), applications that directly interface with daily human activities represent a rich avenue for exploration. This paper dives into a potentially impactful application of MAS, targeting a well-known real-world challenge: meeting scheduling. While there have been previous efforts to address this challenge, we believe that the time is right to revisit this task as a blue-sky challenge for the MAS community.

Traditional scheduling methodologies rely on static, sub-optimal support tools that are susceptible to inefficiencies that include repeated rescheduling, and the overhead for the humans affected per scheduling attempt remains substantial. This opens an intriguing challenge for the MAS community: What if a collection of autonomous agents could extend human capabilities, designed to adapt and negotiate, making scheduling more dynamic and less time-consuming? The potential of collective time saved is substantial, not only in a reduction of human effort due to fewer rescheduling attempts, but also in better alignment of schedules. Furthermore, the privacy of participants can be better preserved.

We argue that the richness of this domain is of interest to the MAS community and that recent advances in AI open up new ways for tackling this challenge. In this paper, we set the stage for this research direction, focussed on the use of MAS to support an age-old, yet fundamental and pervasive task.

KEYWORDS

Multi-Agent Systems, Negotiation, Meeting Scheduling

1 INTRODUCTION

Meetings with others, for social and work-related interactions, form a crucial part of our daily lives. A 2014 survey by Ovum¹ found that employees meet eight times per week on average and that this number has been rising over the years. More specifically, executive management and higher meet on average 12 times a week and VPs, directors and C-level roles in highly collaborative industries reach an even higher average of 17 meetings per week.

These business meetings need to be scheduled, which takes an average of 26–30 minutes per meeting per participant according to a blog by Doodle². This makes scheduling meetings a major time investment for the average employee, who likely has to schedule their own meetings. Higher-ranking roles often have assistants who perform this scheduling task on their behalf and only occasionally ask their bosses for confirmation. Furthermore, manual scheduling

can lead to sub-optimal schedules, due to the complexity of the problem. Tools that support solving this problem are popular³, but are often suboptimal as they only solve part of the problem. There is much to be gained from improving the process of scheduling, which is also recognised by industry⁴.

We informally define the meeting scheduling problem (MSP) as the problem of finding a time slot of a desired duration in which the intended set of participants (or an acceptable subset thereof) commit to attending the meeting at an agreed location. We note that the location can be on-site, online, or a mixture thereof. Furthermore, the notion of an “acceptable subset” makes this *de facto* a family of problems, as it leaves unspecified who determines what defines the acceptability of that subset. In terms of complexity, the problem becomes easy, if this is determined by the one that initiates the scheduling (authority), and most complex if acceptability is determined by a group process amongst the intended participants.

The meeting scheduling problem, being such a major part of daily human life, has seen decades of attention from the computer science community, first appearing in the 80s [25, 33], often modelled as a (form of a) constraint satisfaction problem (CSP) [46]. Researchers have attempted solving the multi-agent MSP using market-based approaches [18] and negotiation approaches, where agents, representing users, negotiate over meeting time slots [26, 40]. In the years after, the problem consistently continued to receive attention among researchers (e.g., [9, 20, 23, 27, 30, 49, 51]) across several communities.

Despite the MSP being a common and relatable problem that has seen considerable effort from the research community, we are still not close to a system that alleviates most of the burden. Difficulties in learning human preferences, communication with humans, and the complexity of decentralized mixed-motive multi-agent problems render the MSP challenging. Many of these challenges are recognised as open problems in cooperative AI [12]. With the recent successes in (multi-agent) (deep) reinforcement learning [45], large language models (LLM), and reinforcement learning from human feedback (RLHF) [44], we believe that now the time is right to revisit the MSP as a rich and rewarding real-world challenge for the MAS community. We also believe that the various communities within computer science that have studied this problem can come together to meet this challenge and jointly achieve far better solutions than currently available.

In this paper, we lay out the necessary groundwork for tackling this problem. We discuss the characteristics of the problem and try to isolate its distinct components. We believe that a decentralised negotiation-based solution is the best-fitting approach to solving

¹Ovum 2014 - Collaboration 2.0: Death of the Web Conference (As We Know It)

²<https://doodle.com/en/resources/blog/study-reveals-time-spent-with-scheduling/>

Proc. of the Adaptive and Learning Agents Workshop (ALA 2024), Avalos, Müller, Wang, Yates (eds.), May 6-7, 2024, Online, <https://ala2024.github.io/>. 2024.

³<https://www.gartner.com/reviews/market/scheduling-automation-software>

⁴<https://www.mckinsey.com/capabilities/operations/our-insights/smart-scheduling-how-to-solve-workforce-planning-challenges-with-ai>

the MSP for scalability and practical reasons; this, therefore, forms the basis of our effort.

2 THE MEETING SCHEDULING PROBLEM

The following anecdotal meeting scheduling process illustrates the richness of the MSP: Alice must schedule a meeting with 4 colleagues, of which 2 are notoriously busy. In his role as organiser, Alice first asks the 2 busy participants about their constraints and options. The first slot found is too far in the future, so the organiser reduces the meeting duration, allowing for two earlier slots. Alice proposes these slots to the other participants. One of them, Bob, already has other obligations conflicting with both slots, but might be able to reschedule a meeting and requests the others to agree on one of the two slots. After the agreement is made, Bob commits to that slot after rescheduling his previous commitment.

In the introduction, we provided a simple and informal definition of MSP that nonetheless already introduces complexity by referring to “an acceptable subset”. Furthermore, under the hood of this definition lurk additional complexities, as mentioned, *e.g.*, by Berger et al. [4]: Each participant must be able to reach the meeting location, attend for the entire duration and reach the next meeting location on time. This refers to travel time between meetings and to means of transportation. Even in online meetings, one must be in a place where one is allowed to speak, and that is quiet enough to hear what is being discussed. Aside from such practicalities that complicate MSPs, there are also numerous human aspects to consider, *e.g.*, participants having ulterior motives and/or hidden agendas, strategic voting, powerplay, and incomplete revelations of potential meeting slots. Any and all of these have an impact on what information they are willing to share and when, how much importance they attach to the meeting and some of its intended participants, and how many attempts have to be made to arrive at a feasible solution.

2.1 Earlier formalisations

The MSP naturally lends itself to be formalised as a constraint satisfaction problem (CSP) [46]. Formalised in this manner, all the techniques for solving CSPs are applicable. This includes centralised and distributed approaches. Centralised approaches boil down to efficient search strategies in the solution space defined by means of hard constraints or strategies for optimising the utility when using soft constraints. In the distributed approaches, solving the CSP is distributed to the agents as a local problem [50]. This often scales better than centrally solving a given CSP instance and is more sensitive to information sharing on a need-to-know basis. Examples of models of specific MSPs as CSP include the assignment problem [13], private incremental multi-agent agreement problem (piMAP) [19], group activity selection problem (GASP) [14, 15], valued constraint satisfaction optimisation problem (VCSOP) [39], group scheduling problem (GSP) [29], and stable group scheduling problem (SGSP) [30].

All of the above-mentioned formalisations of the MSP simplify part of the problem and fail to capture the full richness of the MSP. Such simplifications include full visibility of other agents’ preferences and assuming that decisions are made centrally. In the following, we attempt to describe the full richness of the MSP as a

basis for future research on the topic. In doing so, we refrain from fully formalising the problem, as multiple viable approaches exist. Instead, we focus on the characteristics of the MSP that must be considered when solving this problem.

2.2 Characteristics

We identified a set of characteristics of the problem of scheduling meetings as well as of the participants involved in the meetings. An overview of these characteristics can be found in Table 1. We deliberately make a distinction between these two sets of characteristics. The problem characteristics map the different aspects of the problem that we can either consider or exclude when scheduling meetings. The participant characteristics, describe the potential differences that participants have in relation to others within the considered characteristics of the problem. As soon as at least one participant has a certain characteristic, the problem must accommodate this.

To give a few examples of participant characteristics; within an MSP, a participant might have full visibility of the calendar of only part of the group of participants. Participants might also have different views on the importance of other participants’ attendance and substitutability. For example, Alice might feel that Bob can be substituted by Carol, whereas Bob feels he cannot be substituted at all. As another example, Alice might feel that Dan’s attendance is crucial, whereas Dan does not think the meeting is that important and will only attend if Erin will.

The characteristics we listed can be used to approach the problem in a systematic manner. Due to the richness of the problem, the complexity is also high, and we might want to approach it in incremental steps of complexity. Our characteristics form a structure of challenges that can be attempted in isolation. We will describe the characteristics one by one.

- **Substitutability:** If included, allows for the substitution of (some or all) participants by others. This information is to be observable by the participants. As a participant characteristic, we consider two cases: *Simple*; the set of participants that this participant can be substituted with. This models only the view of a given participant on who can substitute them. *Full*; for this participant, their full view on who (including themselves) can be substituted by whom. Differences in opinion need to be evaluated by all participants.
- **Importance of attendance:** If included, considers that some participants are more important for the meeting than others. Per participant, this is a value in $[0, 1]$ representing the importance of attendance of this or another participant, according to this participant.
- **Calendar Observability:** If included, allows sharing of calendars between participants. Per participant, there are three possibilities: *None* if no part of this participant’s calendar can be directly observed. *Availability* if only the availability of this participant can be observed. *Full*: the participant’s calendar is fully observable by others.

Table 1: Characteristics of the Meeting Scheduling Problem

Description	Problem		Participant		
	Type	Type	Values		
Substitutability	Boolean	Set	Other participants		
Importance of attendance	Boolean	Continuous	0 - Irrelevant	1 - Crucial	
Calendar Observability	Boolean	Categorical	None	Availability	Full
Rescheduling	Boolean	Boolean			
Role	Boolean	Categorical	Organiser	Participant	Observer
Repeated encounters	Boolean	Categorical	Yes	Maybe	No
Multiple rounds	Boolean	N/A			
Preferences	Boolean	Categorical	Invisible	On options	Visible
Arguments	Boolean	Categorical	None	Restricted format	Free text

- **Rescheduling:** If included, meetings can be rescheduled to clear slots for other more important meetings. Per participant, whether this participant has the authority and capability to reschedule existing meeting commitments to free a slot.
- **Roles:** If included, participants can have different roles in the MSP. Per participant, these are: Organiser if this participant is the organiser. Participant if this participant is intended to attend the meeting as a participant. Observer if this participant is intended to attend the meeting as an observer.
- **Repeated encounters:** If included, allows for multiple encounters between agents over the course of scheduling different meetings. Per participant, whether the participant is encountered repeatedly.
- **Multiple rounds:** If included, allows for multiple rounds of back-and-forth communication before an agreement is made.
- **Preferences:** If included, consider preferences over meeting slots instead of simple ‘yes’ or ‘no’ answers on availability. none if this participant only answers ‘yes’/‘no’ to offered slots. on options if this participant provides preferences over offered slots. full if this participant provides access to their full preference profile.
- **Arguments:** if included, permits arguments to be added to answers regarding availability. Per participant, these are None: the participant neither has the capability or authority to add arguments to their answers, nor the ability to interpret arguments made by others. Restricted format: the participant can only add or interpret arguments of a pre-specified restricted format. Free text: the participant can add and interpret text arguments free of other restrictions. The arguments characteristic can be used, for example, to model that a participant provides a conditional answer, e.g., needs consultation with the human user, or that a ‘yes’ is only valid if an agreement can be found within a given time.

2.3 Performance criteria

We argue that the following abstract criteria should be considered when evaluating the performance of a solution to the MSP:

- **Obtained utility:** When considering preferences in the MSP, one can measure properties over the utilities attained by all

agents when used as a global performance measure. Examples are average utility, Pareto-optimality, distance to Nash product or Rawls point, [35, 37]. One can also look at the utility attained by an individual agent as a local measure.

- **Scheduling success:** The percentage of the meetings that could be scheduled. This measures the effectiveness of a meeting scheduling solution in finding common slots and aligning calendars. It should be easier to obtain a perfect score when the number of agents involved is lower or when the density of meetings is low, but becomes an interesting measure when the opposite is true.
- **Privacy preservation** [19, 21, 22]: When observability is (partially) enabled, which is likely true for real-world scenarios, then it becomes important not to reveal too much information, nor share that information with others.
- **Need for rescheduling:** This can be considered an efficiency measure. If a need for rescheduling meetings arises frequently, this could indicate that the agents are not good at estimating future conflicts and that they schedule meetings too easily.
- **Time investment of humans:** As we advocate to approach the MSP via a human-in-the-loop hybrid intelligent approach, humans must be included in the scheduling process, e.g., for preference elicitation or permission in exceptional situations. However, not bothering the human too much is essential for any system to achieve advantages over conventional meeting scheduling methods.
- **Trust and acceptance:** If humans do not trust that the agent will properly schedule their meetings, adoption will be compromised. We note that asking for too little input from humans might be detrimental to trust in the system and the quality of its solutions.
- **Computational cost:** Considering the complexity of the many variants of the MSP, it is important to pay attention to the computational cost incurred by systems for solving this problem.

3 NEGOTIATION IN MEETING SCHEDULING

As mentioned in Section 1, we can distinguish between market-based and negotiation-based approaches to solving the MSP from a multi-agent perspective. Market-based approaches assume that agents are self-interested [18] and are generally based on the ideal that fairness (e.g., maximum social welfare) can be guaranteed

through mechanism design, where the goal is to design a mechanism that satisfies both the incentive-compatible (IC) property (*i.e.*, agents are truthful about preferences) and the individually rational (IR) property (*i.e.*, you cannot receive a negative pay-off from the mechanism). Some success has been achieved using, for example, Clarke tax [6, 17] under simplified conditions, which do not hold up in real-world applications. Designing effective mechanisms for real-world multi-agent systems is theoretically challenging [9].

We argue that negotiation-based approaches are a good fit for the MSP. Firstly, a negotiation approach fits naturally with how humans agree on meeting times. Delegating the legwork to AI agents does not interfere with this and would enable an effective hybrid intelligent solution to this problem, where human capability is extended with AI. Secondly, negotiation is distributed in nature and does not *per se* require a trusted central authority. Thirdly, in negotiation, the practice is only to reveal information on a need-to-know basis, which promotes privacy and is part of the responsibility by-design approach we subscribe to [16].

If we do not require participants to reveal all their preferences and constraints and allow multiple scheduling attempts, then we are basically in a negotiation setting. This is how humans schedule meetings without tooling, often via email, which is cumbersome and inefficient. Tools like Doodle and When2meet can be considered single-shot negotiations [1] as they eliminate the multiple-round component while lowering participants' privacy. We believe negotiation methods make the most sense as we aim for multiple-round, privacy-preserving scheduling.

3.1 Negotiation protocols

Agents must communicate with each other to find agreements. Open communication with other agents in the form of "cheap talk" [11] or with humans in the form of natural language is possible but renders the problem more complex. We deem it more efficient to use negotiation protocols to aid the negotiation process in finding cooperative solutions. Such protocols restrict the type of messages and order in which they are sent [43].

We are not the first to propose negotiation for MSP; examples of proposed protocols for MSP are the single proposer mechanism (SPM) [29] and the distributed score-based multi-round (DSM) negotiation mechanism [19].

3.2 Human preferences

Agents representing humans in negotiations should attempt to optimise outcomes based on human preferences. We, therefore, consider preference elicitation and estimation ([5, 42, 47, 48]) as core components in negotiation-based approaches for solving the MSP.

In general, a preference model can be bootstrapped based on available historical data in the form of preference pairs through direct preference optimisation (DPO) [34]. In the case of MSP, it can be based on historical calendar data [27] and on the current state of the calendar [10]. Estimating the preferences of other humans can help in finding mutually beneficial outcomes. Opponent modelling techniques can be used to estimate these preferences while negotiating, see *e.g.*, [2].

3.3 Learning to negotiate

Given a protocol and preference profile, agents need to learn how to negotiate with other agents, focusing on maximising individual utility, optimising for cooperativeness (*e.g.*, social welfare), or a mixture of those, depending on the characteristics of the MSP at hand. Such agents can be trained using, *e.g.*, automated algorithm configuration [38] or reinforcement learning [3, 32, 41].

In MSP, one can assume that the environment is highly dynamic. New agents will be encountered, other agents will change their behaviour, human preferences over preferred slots will change, etc. Optimising performance means that continuous adaptation is required. An agent's policy can be retrained at fixed times based on historical interactions [31]. After training on a dataset, an agent can be guided by expert annotations to improve exploration online [28].

4 DISCUSSION

The MSP is a challenging problem. Finding optimal Nash equilibrium solutions in such cooperative AI problems is known to be NP-hard [7, 8, 24]. We therefore believe emphasis should be placed on finding solutions that are good enough, but not necessarily optimal, to avoid the need for exponential time solvers. Finding sufficiently good solutions also avoids problems caused by agents aiming to maximise utility deviating or rescheduling for minuscule improvements, which hurts mutually beneficial cooperation [36].

Another important point concerns interaction with humans that are not represented by agents. In the adoption of automated meeting scheduling systems, there will be a transition period during which some humans are represented by agents and others are not. Communication methods change and humans are likely to be less responsive, both in terms of the frequency of interaction and response time. Naturally, agents need to consider this in their scheduling behaviour.

We also have to ensure a degree of fairness in such systems. It cannot be the case that the calendar of some users will be inefficient or that they are being exploited by other agents, simply because they are not properly represented by their agent. Extra care must be taken when a group of agents is dealing with a single participant who is not represented by an agent. A lack of scheduling capabilities should not lead to a drastically less desirable outcome compared to other participants.

If we solve this problem, a societal implication is that humans might change their view on appointments as being somewhat more fluent than is currently the case. Whether this is net beneficial remains to be seen, but an effort should be made to be alert to potential negative side effects.

Finally, we reiterate that, in our view, the time is right to take on this challenge. Recent interest and advances in dealing with human preferences, aligning AI systems, cooperative AI and multi-agent systems can all come together within the domain of meeting scheduling. After a long period of off-and-on attention, the tools might now be available to tackle this problem in a manner that brings substantial benefits to the many individuals who have to regularly schedule meetings and to their organisations.

ACKNOWLEDGMENTS

This research was (partly) funded by the Hybrid Intelligence Center, a 10-year programme funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research, grant number 024.004.022 and by EU H2020 ICT48 project “Humane AI Net” under contract # 952026. This research was also partially supported by TAILOR, a project funded by the EU Horizon 2020 research and innovation programme under GA No 952215.

REFERENCES

- [1] D. Anbar and E. Kalai. 1978. A one-shot bargaining problem. *International Journal of Game Theory* 7, 1 (March 1978), 13–18. <https://doi.org/10.1007/BF01763116>
- [2] Tim Baarslag, Mark J. C. Hendriks, Koen V. Hindriks, and Catholijn M. Jonker. 2016. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems* 30, 5 (Sept. 2016), 849–898. <https://doi.org/10.1007/s10458-015-9309-1>
- [3] Jasper Bakker, Aron Hammond, Daan Bloembergen, and Tim Baarslag. 2019. RLBOA: A Modular Reinforcement Learning Framework for Autonomous Negotiating Agents. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 260–268. <https://www.ifaamas.org/Proceedings/aamas2019/pdfs/p260.pdf>
- [4] Florian Berger, Rolf Klein, Doron Nussbaum, Jörg-Rüdiger Sack, and Jiehua Yi. 2009. A meeting scheduling problem respecting time and space. *Geoinformatica* 13, 4 (Dec. 2009), 453–481. <https://doi.org/10.1007/s10707-008-0053-4>
- [5] Li Chen and Pearl Pu. 2004. *Survey of Preference Elicitation Methods*. Technical Report. EPFL. <https://infoscience.epfl.ch/record/52659>
- [6] Edward H. Clarke. 1971. Multipart Pricing of Public Goods. *Public Choice* 11 (1971), 17–33. <https://www.jstor.org/stable/30022651>
- [7] Vincent Conitzer and Caspar Oesterheld. 2023. Foundations of Cooperative AI. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 13 (Sept. 2023), 15359–15367. <https://doi.org/10.1609/aaai.v37i13.26791>
- [8] Vincent Conitzer and Tuomas Sandholm. 2008. New complexity results about Nash equilibria. *Games and Economic Behavior* 63, 2 (July 2008), 621–641. <https://doi.org/10.1016/j.geb.2008.02.015>
- [9] E. Crawford and M. Veloso. 2004. Mechanism design for multi-agent meeting scheduling including time preferences, availability, and value of presence. In *Proceedings. IEEE/WIC/ACM International Conference on Intelligent Agent Technology, 2004. (IAT 2004)*. IEEE, Beijing, China, 253–259. <https://doi.org/10.1109/IAT.2004.1342952>
- [10] E. Crawford and M. Veloso. 2005. Learning dynamic preferences in multi-agent meeting scheduling. In *IEEE/WIC/ACM International Conference on Intelligent Agent Technology*. IEEE, Compiegne, France, 487–490. <https://doi.org/10.1109/IAT.2005.94>
- [11] Vincent P. Crawford and Joel Sobel. 1982. Strategic Information Transmission. *Econometrica* 50, 6 (1982), 1431–1451. <https://doi.org/10.2307/1913390>
- [12] Allan Dafoe, Edward Hughes, Yoram Bachrach, Tantum Collins, Kevin R. McKee, Joel Z. Leibo, Kate Larson, and Thore Graepel. 2020. Open Problems in Cooperative AI. <https://doi.org/10.48550/arXiv.2012.08630>
- [13] Panayiotis Danassiss, Florian Wiedemair, and Boi Faltings. 2021. Improving Multi-agent Coordination by Learning to Estimate Contention. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Vol. 1. International Joint Conferences on Artificial Intelligence Organization, Montreal, Canada, 125–131. <https://doi.org/10.24963/ijcai.2021/18>
- [14] Andreas Darmann, Edith Elkind, Sascha Kurz, Jérôme Lang, Joachim Schauer, and Gerhard Woeginger. 2012. Group Activity Selection Problem. In *Internet and Network Economics (Lecture Notes in Computer Science)*, Paul W. Goldberg (Ed.). Springer, Berlin, Heidelberg, 156–169. https://doi.org/10.1007/978-3-642-35311-6_12
- [15] Andreas Darmann, Edith Elkind, Sascha Kurz, Jérôme Lang, Joachim Schauer, and Gerhard Woeginger. 2018. Group activity selection problem with approval preferences. *International Journal of Game Theory* 47, 3 (Sept. 2018), 767–796. <https://doi.org/10.1007/s00182-017-0596-4>
- [16] Virginia Dignum. 2019. *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way*. Springer International Publishing, Cham. <http://link.springer.com/10.1007/978-3-030-30371-6>
- [17] Eithan Ephrati and Jeffrey S. Rosenschein. 1991. The Clarke tax as a consensus mechanism among automated agents. In *Proceedings of the ninth National conference on Artificial intelligence - Volume 1 (AAAI'91)*. AAAI Press, Anaheim, California, 173–178.
- [18] Eithan Ephrati, Gilad Zlotkin, and Jeffrey S. Rosenschein. 1994. Meet your destiny: a non-manipulable meeting scheduler. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work (CSCW '94)*. Association for Computing Machinery, New York, NY, USA, 359–371. <https://doi.org/10.1145/192844.193049>
- [19] Farzaneh Farhadi and Nicholas R. Jennings. 2021. A Faithful Mechanism for Incremental Multi-Agent Agreement Problems with Self-Interested and Privacy-Preserving Agents. *SN Computer Science* 2, 4 (May 2021), 272. <https://doi.org/10.1007/s42979-021-00650-4>
- [20] Maria Sole Franzin, E C Freuder, F Rossi, and R Wallace. 2002. Multi-Agent Meeting Scheduling with Preferences: Efficiency, Privacy Loss, and Solution Quality. In *Proceedings of the AAAI 2002 Workshop on Preferences in AI and CP: Symbolic Approaches*. AAAI Press, Edmonton, Canada, 8. <https://cdn.aaai.org/Workshops/2002/WS-02-13/WS02-13-005.pdf>
- [21] M. S. Franzin, F. Rossi, E. C. Freuder, and R. Wallace. 2004. Multi-Agent Constraint Systems with Preferences: Efficiency, Solution Quality, and Privacy Loss. *Computational Intelligence* 20, 2 (2004), 264–286. <https://doi.org/10.1111/j.0824-7935.2004.00238.x>
- [22] Eugene C Freuder, Marius Minca, and Richard J Wallace. 2001. Privacy/Efficiency Tradeoffs in Distributed Meeting Scheduling by Constraint-Based Agents. In *Proceedings of the IJCAI-01 Workshop on Distributed Constraint Reasoning*. AAAI Press, Seattle, USA, 9.
- [23] Leonardo Garrido and Katia Sycara. 1996. Multi-Agent Meeting Scheduling: Preliminary Experimental Results. In *Proceeding of the Second International Conference on Multiagent Systems 1996*. AAAI Press, Kyoto, Japan, 8.
- [24] Itzhak Gilboa and Eitan Zemel. 1989. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior* 1, 1 (March 1989), 80–93. [https://doi.org/10.1016/0899-8256\(89\)90006-7](https://doi.org/10.1016/0899-8256(89)90006-7)
- [25] Irene Greif. 1982. *PCAL: A Personal Calendar*. Technical Report. MIT LCS. <https://apps.dtic.mil/sti/citations/ADA111706>
- [26] N. R. Jennings and A. J. Jackson. 1995. Agent-based Meeting Scheduling: A Design and Implementation. *Electronics Letters* 31, 5 (1995), 350–352. <https://eprints.soton.ac.uk/252097/>
- [27] Donghyeon Kim, Jinhyuk Lee, Donghee Choi, Jaehoon Choi, and Jaewoo Kang. 2018. Learning User Preferences and Understanding Calendar Contexts for Event Scheduling. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*. Association for Computing Machinery, New York, NY, USA, 337–346. <https://doi.org/10.1145/3269206.3271712>
- [28] Minae Kwon, Siddharth Karamcheti, Mariano-Florentino Cuellar, and Dorsa Sadigh. 2021. Targeted Data Acquisition for Evolving Negotiation Agents. In *Proceedings of the 38th International Conference on Machine Learning*. PMLR, Virtual, 5894–5904. <https://proceedings.mlr.press/v139/kwon21a.html>
- [29] Hooyeon Lee. 2014. Algorithmic and Game-theoretic Approaches to Group Scheduling. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*. IFAAMAS, Paris, France, 1709–1710.
- [30] Hooyeon Lee and Yoav Shoham. 2014. Stable group scheduling. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems (AAMAS '14)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1347–1348.
- [31] Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or No Deal? End-to-End Learning of Negotiation Dialogues. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Martha Palmer, Rebecca Hwa, and Sebastian Riedel (Eds.). Association for Computational Linguistics, Copenhagen, Denmark, 2443–2453. <https://doi.org/10.18653/v1/D17-1259>
- [32] Zun Li, Marc Lanctot, Kevin R. McKee, Luke Marris, Ian Gemp, Daniel Hennes, Paul Muller, Kate Larson, Yoram Bachrach, and Michael P. Wellman. 2023. Combining Tree-Search, Generative Models, and Nash Bargaining Concepts in Game-Theoretic Reinforcement Learning. <https://doi.org/10.48550/arXiv.2302.00797>
- [33] Thomas W Malone, Kenneth R Grant, Franklyn A Turbak, Stephen A Brobst, and Michael D Cohen. 1987. Intelligent information-sharing systems. *Commun. ACM* 30, 5 (May 1987), 390–402. <https://doi.org/10.1145/22899.22903>
- [34] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. <https://doi.org/10.48550/arXiv.2305.18290>
- [35] Howard Raiffa. 2007. *Negotiation Analysis: The Science and Art of Collaborative Decision Making*. Harvard University Press, USA.
- [36] Senthil Rajasekaran, Suguman Bansal, and Moshe Y. Vardi. 2023. Multi-Agent Systems with Quantitative Satisficing Goals. <https://doi.org/10.48550/arXiv.2305.00953>
- [37] John Rawls. 2020. *A Theory of Justice: Revised Edition*. Harvard University Press, USA.
- [38] Bram M. Renting, Holger H. Hoos, and Catholijn M. Jonker. 2020. Automated Configuration of Negotiation Strategies. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '20)*. International Foundation for Autonomous Agents and Multiagent Systems, Auckland, 1116–1124. <http://www.ifaamas.org/Proceedings/aamas2020/pdfs/p1116.pdf>
- [39] Hajer Salem and Ahlem Ben Hassine. 2015. Meeting Scheduling based on Swarm Intelligence. *Procedia Computer Science* 60 (Jan. 2015), 1081–1091. <https://doi.org/10.1016/j.procs.2015.08.153>

- [40] Sandip Sen and Edmund H. Durfee. 1991. A formal study of distributed meeting scheduling: preliminary results. In *Conference proceedings on Organizational computing systems - COCS '91*. ACM Press, Atlanta, Georgia, United States, 55–68. <https://doi.org/10.1145/122831.122837>
- [41] Ayan Sengupta, Yasser Mohammad, and Shinji Nakadai. 2021. An Autonomous Negotiating Agent Framework with Reinforcement Learning based Strategies and Adaptive Strategy Switching Mechanism. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '21)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1163–1172. <https://www.ifaamas.org/Proceedings/aamas2021/pdfs/p1163.pdf>
- [42] Burr Settles. 2009. *Active Learning Literature Survey*. Technical Report. University of Wisconsin-Madison Department of Computer Sciences. <https://minds.wisconsin.edu/handle/1793/60660>
- [43] Smith. 1980. The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver. *IEEE Trans. Comput.* C-29, 12 (Dec. 1980), 1104–1113. <https://doi.org/10.1109/TC.1980.1675516>
- [44] Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., Virtual, 3008–3021. <https://proceedings.neurips.cc/paper/2020/hash/1f89885d556929e98d3ef9b86448f951-Abstract.html>
- [45] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning, second edition: An Introduction*. MIT Press, USA.
- [46] Edward Tsang. 1993. *Foundations of Constraint Satisfaction*. Academic Press, UK. <https://doi.org/10.1016/C2013-0-07627-X>
- [47] Ivan Vendrov, Tyler Lu, Qingqing Huang, and Craig Boutilier. 2020. Gradient-Based Optimization for Bayesian Preference Elicitation. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 06 (April 2020), 10292–10301. <https://doi.org/10.1609/aaai.v34i06.6592>
- [48] Toby Walsh. 2007. Uncertainty in preference elicitation and aggregation. In *Proceedings of the 22nd national conference on Artificial intelligence - Volume 1 (AAAI'07)*. AAAI Press, Vancouver, British Columbia, Canada, 3–8.
- [49] Dujuan Wang, Yugang Yu, Yunqiang Yin, and Tai Chiu Edwin Cheng. 2021. Multi-agent scheduling problems under multitasking. *International Journal of Production Research* 59, 12 (June 2021), 3633–3663. <https://doi.org/10.1080/00207543.2020.1748908>
- [50] M. Yokoo, E.H. Durfee, T. Ishida, and K. Kuwabara. 1998. The distributed constraint satisfaction problem: formalization and algorithms. *IEEE Transactions on Knowledge and Data Engineering* 10, 5 (Sept. 1998), 673–685. <https://doi.org/10.1109/69.729707>
- [51] Alejandro Zunino and Marcelo Campo. 2009. Chronos: A multi-agent system for distributed automatic meeting scheduling. *Expert Systems with Applications* 36, 3, Part 2 (April 2009), 7011–7018. <https://doi.org/10.1016/j.eswa.2008.08.024>