Fairness in Cooperative Multi-Agent Systems
MAS Course Project

Alex Periti
Master's Degree in Artificial Intelligence, University of Bologna
alex.periti@studio.unibo.it

Abstract

This systematic literature review (SLR) delves into the evolving landscape of fairness implementation and management within Multi-Agent Systems (MAS). Over the past decade, the integration of fairness considerations in MAS has gained prominence, necessitating a closer examination of methodologies and challenges associated with this integration. This review categorizes several approaches, highlighting the nuanced dimensions of fairness, including distributive fairness, procedural fairness and multi-objective optimization problem, and scrutinize strategies for managing fairness. By synthesizing these findings, this review offers a comprehensive understanding of the current state of fairness in MAS, contributing to the discourse on responsible and equitable multi-agent system development.

1 Introduction

Fairness has emerged as a critical consideration in the design and deployment of MAS, reflecting the growing awareness of ethical implications and societal impact associated with autonomous decision-making. This SLR aims to provide an overview of the current state of research on fairness within MAS, focusing on the implementation and management, providing example of articles that address this topic, crucially examples such as distributive fairness, procedural fairness, and fairness in multi-objective optimization problems.

"In MAS, distributive fairness pertains to the equitable allocation of resources and outcomes among autonomous agents. As MAS increasingly operate in dynamic and resource-constrained environments, understanding how distributive fairness is conceptualized and implemented becomes important. This review explores the diverse methodologies and models employed to ensure fairness in distribution of resources.

Procedural fairness, on the other hand, involves the fairness of the decision-making process itself within MAS. Examining the procedural aspect is crucial in understanding how agents arrive at decisions, ensuring transparency, accountability, and avoiding biased practices. This SLR investigates a refined cognitive architecture inspired by the beliefs-desires-intentions (BDI) proposed to enhance procedural fairness in multi-agent settings.

Moreover, this review delves into the intersection of fairness and multi-objective optimization problems within MAS. This paper shows how the challenging task of optimization can be addressed by means of a distributed algorithm to solve constrained optimization problems, aiming to find a fair solution among all efficient while upholding ethical standards.

This article is organized as follows: in section 2 the research objective and questions are proposed to guide the whole process of this SLR. Search Strategy, Study Selection and Inclusion Criteria and screening process are also discussed in section 2. Section 3 have the purpose to provide a "simplistic" theoretical background about the concept of fairness in a MAS. Section 4,5 and 6 presents main findings of the reviewed papers, that is the topics of distributive fairness, procedural fairness and multi-objective optimization. Section 7 exhibits the answer to the goal question. Finally section 8 concludes this SLR and also discusses future research.

2 Methodology

2.1 Research Question

This SLR is conducted to investigate the current state of research of the fairness in a MAS scenario. The research question hence is:

(G) How is fairness handled in a multiagent system?

To address this, is relevant to point out also subgoal question like:

- Is possible to define a concept of fairness that could be implemented in a MAS?
- What is the impact of fairness in the behaviour of agents?
- Is always possible to ensure fairness in agent's decisions?

2.2 Search Strategy

The search was performed using Scopus, a widely recognized academic search engine, which provides a vast collection of peer-reviewed journals, conference proceedings, and academic books. The results were reviewed in accordance with the inclusion criteria, and also by setting filters, such as the temporal window, set in the last 10 years (2013-2024) in order to gather only recent papers, english as language of the articles and conference paper or article as type of document.

2.3 Keywords

The following keywords were used to construct search queries: "fairness in multi agent/multiagent/multi-agent system(s)", "managing fairness in multi agent/multiagent/multi-agent system(s)", "handling fairness in multi agent/multiagent/multi-agent system(s)", "implement fairness in multi agent/multiagent/multi-agent system(s)". These queries were chosen to encompass the key themes related to this research topic. The queries is composed of several keywords since the topic is quite broad, thus it's needed to be enough precise in order to not get lost.

2.4 Inclusion and Exclusion Criteria

The inclusion and exclusion criteria were defined mainly for filtering the large amount of papers and also to ensure the relevance and significance of the selected articles.
The inclusion criteria are the following:

1. **Number of Citations**: articles with a substantial number of citations were given priority to ensure the inclusion of influential and widely recognized works.
2. **Pertinence with the Theme**: each article's pertinence to the research theme was assessed through an examination of the abstract. Only articles directly addressing implementation of fairness in MAS were included.

The exclusion criteria are the following:

1. **The paper is a SLR**: article that is itself a SLR is not included.
2. **Duplicate**: article that is a duplicate of one of the already analyzed paper is not included.
3. **Access Restriction**: article of which full text is accessible only previous payment is not included.
4. **Topic**: article unrelated to any of the queries is not included. Also, if the article, applies a specific framework to address a specific technical problem, is not included.

About the last exclusion criteria, i want to point out that i decided to focus more on general purpose approaches, not on the application on specific engineered problem, otherwise the amount of material would be overwhelming (also i don't feel to have the competence for judge such articles).

2.5 Screening Process

The initial screening involved reviewing the titles and, eventually, the abstracts of the retrieved articles, about 60 articles. Articles that did not align with the research theme or not meet the inclusion criteria were excluded. Finally the full texts of the remaining articles were then examined to ensure their suitability for inclusion in the review. Also in the last step some articles were excluded, since they don not add significant content with respect to the articles already read.
At the first step 13 papers was selected, then the abstract was analyzed and 2 papers were discarded. Among the 11 papers left, after reading them only 10 articles were considered meaningful.

3 Notion of fairness in multi-agent systems

The concept of fairness in MAS is crucial for ensuring that these systems are both efficient and stable. Fairness in MAS can be related to various aspects such as resource allocation, decision-making, and the distribution of rewards or tasks among agents. Usually, fairness often involves ensuring that no individual agent or group of agents is consistently disadvantaged or that the resources are allocated in a way that is considered equitable by some standard.

4 Fairness in resource distribution

In this section are discussed the examined articles that expose the topic of fairness in a resource distribution scenario.

4.1 Price of Fairness for allocating a bounded resource

The paper [8] discusses the concept of fairness in MAS, particularly in the context of resource allocation. The fairness is defined as the perception of an equitable distribution of resources among agents. An optimal solution may present a very unbalanced allocation of the resource to the agents and hence be perceived as unfair. Conversely, balanced allocations may be far from the optimum. The authors are interested in assessing the quality of fair solutions, i.e., in measuring the system efficiency loss under a fair allocation compared to the one that maximizes the total utility. This indicator is called the Price of Fairness (PoF)

For illustrating the PoF, the approach used, is to first formalize a number of properties holding for any general multi-agent problem without any special assumptions on the agent utilities. Then it's introducted an allocation problem, in which each agent can consume the available bounded resource in given discrete quantities (items). The utility of each agent is given by the sum of these quantities (weights of allocated items).

In major detail, the problem is called the Fair Subset Sum Problem (FSSP), where a common and bounded resource is to be shared among a set of agents. Each agent owns a number of indivisible items each of which requires a certain amount of the resource, called item weight, in order to be utilized. The problem consists in selecting, for each agent, a subset of items so that the sum of all selected item weights is not larger than a given upper bound expressing the resource capacity. The utility function of each agent consists of the sum of weights over all selected items of that agent. The authors provide upper and lower bounds on the Price of Fairness as functions of an upper bound on the items size for the case of two agents.

4.2 Equitable distribution and group-based fairness

In [6], the concept of fairness in cooperative MAS is presented, particularly focusing on the equitable distribution of rewards among agents. It is observed that naive maximization of team reward can lead to unfair outcomes for individual team members. To address this issue, the concept of team fairness is introduced as a group-based fairness measure for multi-agent learning. A method known as Fairness through Equivariance (Fair-E) is proposed to enforce team fairness during policy optimization by transforming the team's joint policy into an equivariant map. Additionally, a soft-constraint version of Fair-E, called Fairness through Equivariance Regularization (Fair-ER), is introduced, demonstrating higher levels of utility than Fair-E while achieving fairer outcomes than non-equivariant policies.

Novel findings regarding the fairness-utility tradeoff in multi-agent settings are presented showing that the magnitude of the trade-off depends on agent skill. When agent skill is high, fairness comes with no trade-off in utility. However, as skill decreases, gains in team fairness are increasingly offset by decreases in team utility. This suggests that fairness can be more expensive for lesser-skilled agents.

4.3 Auction-based resource allocation

The concept of fairness in multiagent systems through the lens of auction-based resource allocation is discussed in [10]. In this study, resources are allocated based on multiple attributes beyond just price, which is particularly relevant in scenarios where factors such as delivery time, quality, and ecological footprint influence resource allocation decisions. The consideration of multiple attributes aims to prevent undesirable outcomes, such as unbalanced workloads or production delays, which may arise from solely prioritizing price. A broader view of fairness in auctions, incorporating non-economic aspects to enhance social welfare by addressing issues like the bidder drop problem [1] and maintaining competitive market prices, is argued for.

To implement multi-dimensional fairness, a priority-based mechanism is proposed. Priorities are assigned to bidders based on their auction history, considering factors like the number of lost

[1]When using an utilitarian approach to measure the satisfaction, the aggregation does not consider the fact that there could be large differences among agents' payoffs and when auctions are repeated over time, this situation may lead to the dissatisfaction of certain participants who may eventually decide to leave the market. When this occurs only the most powerful bidders remain in the market, gaining the chance to create an oligopoly, control the market price and provoke a general fall in prices, which can bankrupt the auctioneers.

auctions, consecutive losses, and the fitness of their losing bids. These priorities influence auction outcomes, intending to sustain bidders' interest and promote a more equitable wealth distribution. Qualitative and quantitative priority methods, including deterministic and stochastic versions, are delineated to integrate multi-dimensional fairness into the auction mechanism.
The proposed approach undergoes evaluation using a simulator emulating an industrial production environment, where diverse agents compete for tasks. The evaluation focuses on three main aspects: price, delivery time, and fairness. Results indicate that multi-dimensional fairness encourages agent participation, improves wealth distribution equity, and maintains competitive market dynamics without compromising the quality of allocation attributes.

4.4 Distribution of indivisible items

The paper [5] specifically addresses the fair division of indivisible items among agents using the maximin share (MMS) fairness criterion. To discuss the importance of fairness, two well-studied notions are introduced: envy-freeness and proportionality. However, in the context of indivisible items, these notions cannot always be satisfied, which leads to the introduction of the maximin share (MMS) as an alternative fairness criterion.
The MMS concept is a generalization of the cut and choose protocol from the cake-cutting problem [2] and serves as a natural relaxation of proportionality. It is defined as the maximum value an agent can guarantee for himself if he partitions the items into bundles and is the last to choose a bundle. The MMS value represents a specific objective measure of fairness for each agent. The paper discusses the relevance of MMS in fair division and notes that while an MMS allocation may not always exist, approximations such as $\frac{2}{3}$ MMS allocations are always possible. In addiction it's presented a new approach that simplifies the process of showing the existence of a $\frac{3}{4}$ MMS allocation. The paper's main contributions include:

- A simple algorithm that demonstrates the existence of a $\frac{3}{4}$ MMS allocation.
- A strongly polynomial time algorithm to find a $\frac{3}{4}$ MMS allocation without needing to approximate MMS values.
- Proof is provided that there always exists a $(\frac{3}{4} + \frac{1}{12}n)$ MMS allocation, which improves the approximation guarantee, especially for a small number of agents.

The basic functioning of the algorithm, assumes known MMS values for all agents and scales valuations so each agent's MMS value is 1. It's adopted a greedy approach to assign high-value items to agents and then employ a bag filling procedure to distribute the remaining items. Moreover, the basic algorithm is extended to compute a $\frac{3}{4}$ MMS allocation in strongly polynomial time without computing actual MMS values. This involves tentative assignments and updating the MMS upper bound as necessary.

4.5 Trade-offs in repeated dynamic resource allocation

In [7] it is argued that while most resource allocation problems typically focus on optimizing resource usage, utility, or overall cost, fairness is a crucial factor that has often been overlooked. This holds true, especially in systems relying on the voluntary participation of autonomous entities and their provision of resources, where considerations include retaining as many participants as possible and ensuring the long-term endurance of the system itself.
The example of autonomous power management systems is used, where resource allocation determines which generators produce the power required to satisfy the overall load. It is noted that while a focus on utility can yield excellent results in terms of load distribution and overall cost, it may also result in situations where specific participants consistently win the bidding process, effectively driving others out of the system. This potential lack of robustness and the risk of systemic failure are highlighted.
To address this issue, a fair resource allocation scheme based on Ostrom's principles [9] [3] and Rescher's canons [11] of distributive justice is proposed. It is argued that fair allocations create value in repeated resource allocations, increasing the satisfaction of individual agents while reducing the risk of optimizing the problem in the short term at

[2] Fair cake-cutting is a kind of fair division problem. The problem involves a heterogeneous resource, such as a cake with different toppings, that is assumed to be divisible

[3] Ostrom's work is known for challenging the conventional wisdom that common-pool resources are poorly managed and prone to overuse. She identified several principles that characterize successful institutional arrangements for governing common-pool resources, emphasizing the importance of local knowledge, collective action, and institutional diversity

the cost of long-term robustness and stability.
The Linear Public Goods Game (LPG), a classical experiment in game theory demonstrating the contradiction of following the dominant individual strategy of withholding one's own resources to the disadvantage of all players, is discussed. A variant of the LPG, the LPG', is proposed, combining the principles put forward by Ostrom and the canons proposed by Rescher to establish a resource allocation system for open systems that ensures fairness and endurance. Furthermore, it's emphasized that a fair distribution of resources can increase the satisfaction of individual agents, improve system stability, and reduce the risk of systemic failure.

5 Fairness in procedural process

This section discusses the impact of fairness in the behaviour of agents in a decision-making task, exposed in the examined articles.

5.1 Impact of different belief facets

The concept of fairness, particularly within the context of agents' decision-making processes, is presented in the paper [1] through a conceptual refinement of agent cognitive architecture.
Fairness, identified as a parameter associated with the model presented, is defined as the perception of fairness within the system. In the context of this paper, fairness is concerned with how individual agents punish each other, not the organization of the system itself.
A refined cognitive architecture is presented, inspired by the beliefs-desires-intentions (BDI) and the theory of planned behavior (TPB) models. This architecture allows for the investigation of the impact of personality and the way an agent weighs its internal beliefs and social sanctions on actions. The study also incorporates the concept of cognitive dissonance associated with the fairness of institutions to examine behavior.
A variation of belief facets is proposed, considering three different punitive/rewarding components based on the agent's internal acceptance of the rule (internal belief), its perception of societal support for the rule (perceived norms), and its belief about the purpose of the rule (rule-understanding). These belief facets are used to model how institutional characteristics, including the fairness of institutions, impact decisions. An agent's decision utility function is described, modeling how an agent makes decisions. The utility function has four parts: the revenue earned for an action, how the agent is mentally punished for such an action, the agent's perception of societal support for the action, and the agent's understanding of the rule.
Different aspects of the decision-making process, such as learning method and belief weighting, are impacted by the agent's personality.
The paper emphasizes the importance of internal beliefs as a pivotal aspect for adhering to institutional rules.

5.2 Human inspired fairness

The concept of fairness in the context of social dilemmas is addressed in the paper [4]. The paper asserts that existing computational models of fairness are not explicitly tailored to handle social dilemmas, a category of tasks where choices involve a trade-off between group reward (social cooperation) and individual reward (defection). A novel computational model, termed human-inspired computational fairness, is proposed as a solution specifically designed for social dilemmas. The model draws inspiration from the way humans address social dilemmas, applying fully decentralized sanctioning mechanisms to ensure fair and optimal solutions emerge, even when some participants decide purely based on individual reward. Human-inspired computational fairness is posited as essential in Multi-Agent Systems (MAS), enabling agents to find improved solutions when confronted with social dilemmas. The model is characterized as fully decentralized, computationally tractable, and robust against agents attempting to exploit the fairness of others. The argument is put forth that this model surpasses the limitations of existing computational models, which often assume universal adherence to proposed fairness models—a presumption that does not consistently hold in open MAS.
The paper discusses the commonly used two-step approach in behavioral literature to model human fairness in social dilemmas. The first step involves controlled experiments with human subjects to understand their optimization of individual reward while considering fairness. The second step involves modeling proposed mechanisms in a computational context using multi-agent learning approaches, evolutionary-game-theoretic analysis, or statistical physics.
Regarding computational social choice, the paper delves into the limitations of existing definitions

and procedures from welfare economics and computational social choice, particularly their lack of explicit attention to social dilemmas and impracticality for large-scale, decentralized, and open multiagent systems. The authors contend that a pragmatic model, such as human-inspired computational fairness, is more suitable for such systems. The proposed human-inspired computational fairness model is designed to overcome the limitations of existing models, offering a decentralized, computationally tractable, and robust approach to ensuring fair and optimal solutions in open multiagent systems.

5.3 Multi-agent social choice

In [2], the concept of fairness in the context of recommender systems is explored. A framework is presented where the interests of various stakeholders, including providers, are represented as agents participating in the production of recommendations through a two-stage social choice mechanism. This approach allows for the representation of a wide variety of fairness concepts and can extend to multiple fairness concerns. The pursuit of algorithmic fairness in recommender systems requires rethinking the conventional definition of the "user". Traditionally, the user is the receiver of recommendations, but fairness, especially provider-side fairness, necessitates considering a broader array of system users and stakeholders. A post-userist user model is advocated for that accounts for the interests and preferences of all stakeholders, not just the recipients of recommendations.
An architecture for implementing multi-stakeholder fairness in recommender systems is introduced, where fairness concerns are represented as agents participating in a dynamic social choice environment. The framework is built to be dynamic, making decisions in the context of historical choices and results, and capable of integrating multiple fairness concepts that may be in tension with each other.
Fairness is defined as a contested concept, meaning that ideas about fairness are grounded in specific contexts and stakeholders and can be multiple and possibly conflicting. A fairness concern is defined as a specific type of fairness being sought, relative to a particular aspect of recommendation outcomes, evaluated in a particular way. This reification is crucial as it models the specific preferences that stakeholders may have concerning the outcomes produced by the system. Different logics of fairness as identified by welfare economists, including exogenous right, compensation, reward, and fitness, are discussed in the paper. These logics do not always align and can be the root of political disagreements. For example, in college admissions, the criteria for selection may pit reward for achievement against compensation for disadvantage.
Recommender systems often operate as two-sided platforms with consumers, providers, and the platform operator as direct stakeholders. Fairness concerns may derive from any stakeholder and may need to be balanced against each other. The platform may enforce fairness even when other stakeholders are not concerned with it. Additionally, fairness concerns can arise on behalf of indirect stakeholders who are impacted by recommendations but are not a party to them, such as representational fairness.

6 Fairness in optimization

In this section are discussed the examined articles that expose the topic of fairness in a problem optimization scenario.

6.1 Optimization algorithm for nash bargaining

The concept of fairness through the lens of cooperative game theory, specifically focusing on Nash bargaining is discussed in [3].
The goal is to address multi-objective optimization problem over networks where agents aim to maximize their own objective functions while adhering to both local and coupling constraints. This setup is common in real world applications such as smart electric grids, self-organized factories, or autonomous vehicles, where the goal is not just stabilization but optimal operation with respect to economic cost criteria.
The main idea for solving this problem arises from cooperative game theory. A bargaining problem, as defined by Nash, is a game where multiple players (agents) negotiate to share a surplus generated through cooperation. The solution to such a game is a unanimous agreement on how to split the surplus, with each agent aiming to maximize their share. The Nash bargaining solution is considered a fair solution within this context. The paper proposes a distributed algorithm to solve con-

vex constraint-coupled multi-objective optimization problems, aiming to find a fair solution among all efficient (Pareto optimal) solutions. The notion of fairness is crucial because the unweighted sum of local performance criteria, often used as the overall performance criterion, can lead to unfair outcomes where some systems are favored over others. The Nash bargaining solution is sought to ensure fairness in the optimal steady states for multi-agent control systems.

A negotiation mechanism among agents is designed to compute the Nash bargaining solution in a distributed manner. The problem is reformulated as the maximization of a weighted sum of the objective functions. The algorithm is a two-step procedure where local estimates of the Nash bargaining weights are updated online, and existing distributed optimization algorithms are applied. This method is formally analyzed, and numerical simulations are provided to demonstrate its efficacy. In conclusion, the proposed optimization algorithm yields fair solutions in multi-objective optimization problems, ensuring that all agents are treated equitably in the distribution of resources or benefits.

6.2 Reinforcement learning for optimizing offloading decisions

The comprehensive study on the formulation of computation offloading as a decentralized decision-making problem within autonomous agents is presented in the paper [12]. An interaction mechanism is introduced to balance competition and cooperation among agents, aiming to align private goals with the system's objectives. This mechanism is particularly relevant in contexts characterized by a large number of users, dynamic nature, and diverse Quality of Service (QoS), such as vehicular network (V2X) applications. A decentralized decision-making mechanism based on second-price sealed-bid auctions and a multi-agent reinforcement learning (MARL) algorithm is proposed to address challenges such as user autonomy, information delay, and computational intractability. The concept of fairness in this MAS is addressed through the design of the interaction mechanism and the MARL algorithm. The mechanism ensures that all agents have at least some access to resources, considered a form of fairness. The fairness constraint is mathematically defined, aiming to maximize overall resource allocation subject to this fairness constraint. The Nash Equilibrium (NE) of the game achieves optimal resource allocation, demonstrating the mechanism's ability to balance fairness with efficiency.

The MARL algorithm proposed in the paper learns with partial, delayed, and noisy state information, along with a reward signal that significantly reduces the need for extensive information sharing. This learning approach allows agents to improve both system and individual performance significantly, demonstrating reductions in offloading failure rate (40%), communication overhead (32%), and computation resource savings in low contention scenarios (38%). Furthermore, the algorithm enhances fairness among agents, ensuring a more equitable distribution of resources and opportunities for offloading tasks. Empirical results confirm the effectiveness of the proposed interaction mechanism and MARL algorithm in improving both system and individual performance, as well as fairness. The good convergence and generalization of the algorithm indicate its robustness and adaptability to different environments.

This work contributes significantly to the field of distributed autonomous systems, offering insights and methodologies applicable to various applications beyond V2X scenarios.

7 Discussion

The goal and subgoal questions for this SLR, as defined in 2, can now be addressed, summarizing the findings of the examined articles.

- **How is fairness handled in a multiagent system?** Fairness in MAS is addressed through various approaches, depending on the context. For resource allocation scenarios, we could use the metrics like the PoF, which evaluates system efficiency loss under fair allocation, or the metric MMS for indivisible item. In repeated dynamic resource allocation, fairness is crucial to prevent systemic failure and ensure the long-term endurance of systems. In cooperative MAS, Fairness through Equivariance and Fairness through Equivariance Regularization are a valid option to enforce team fairness during policy optimization. Auction-based resource allocation introduces multi-dimensional fairness, where priorities based on auction history are assigned to bidders to enhance wealth distribution equity.

  Fairness in procedural processes, particularly

in agents' decision-making, involves models that consider internal beliefs, societal support, and cognitive dissonance associated with institutional fairness.
In optimization scenarios, (ex:cooperative game theory) one way to achieve fairness is Nash bargaining, which ensure equitable distribution of resources among agents. Another example is the use of reinforcement learning to optimize offloading decisions, balancing competition, cooperation among agents, and ensuring fairness.

- **Is it possible to define a concept of fairness that could be implemented in a MAS?** Yes, fairness is defined and implemented in a MAS through various metrics and mechanisms. Look at the previous answer for a summary of several possible way to implement fairness.

- **What is the impact of fairness on the behavior of agents?** Fairness has a significant impact on the behavior of agents. In cooperative MAS, the concept of team fairness influences policy optimization. In resource allocation scenarios, the consideration of fairness in auctions affects agent participation and encourages equitable wealth distribution. The impact of fairness in procedural processes, particularly in agents' decision-making, is studied through cognitive architectures that incorporate beliefs, societal support, and cognitive dissonance. Fairness in optimization scenarios, such as Nash bargaining, influences agents' negotiation mechanisms and their pursuit of equitable resource allocation.

- **Is it always possible to ensure fairness in agents' decisions?** Ensuring fairness in agents' decisions depends on the context and the specific mechanisms implemented. However, the feasibility of ensuring fairness may vary. In some cases, the limitations of existing models are acknowledged, and pragmatic models like human-inspired computational fairness are proposed to address challenges in open multiagent systems. The impact of agent skill is also considered, suggesting that fairness can be more expensive for lesser-skilled agents.

8 Conclusion

This SLR provides a comprehensive examination of fairness within Multi-Agent Systems (MAS). The study categorizes various approaches to fairness, including distributive fairness, procedural fairness, and fairness in multi-objective optimization problems, and explores strategies for managing fairness in MAS. The review highlights the importance of fairness in the design and deployment of MAS, given the ethical implications and societal impact associated with autonomous decision-making. The research has shown that distributive fairness is crucial for the equitable allocation of resources among agents, especially in dynamic and resource-constrained environments. Procedural fairness ensures transparency and accountability in the decision-making process, avoiding biased practices. The intersection of fairness with multi-objective optimization problems is also addressed, showing how optimization tasks can be approached while maintaining ethical standards. It should be clear now that there isn't a good and broad solution to ensure fairness in every scenario, instead we should seek to establish relevant fairness criteria depending on the system that we have to deploy, by asking directly to the stakeholder and experts for that specific scenario. Future developments in fairness within MAS may involve refining models and algorithms to better address the complexities of fairness in various contexts. This could include enhancing the understanding of the fairness-utility trade-off, especially for lesser-skilled agents, and developing more sophisticated mechanisms for procedural fairness. Additionally, the continued evolution of fairness in MAS will likely involve interdisciplinary research to create more responsible and equitable systems.

References

[1] A. H. Afshar Sedigh, M. K. Purvis, B. T. R. Savarimuthu, C. K. Frantz, and M. A. Purvis. 2024. Impact of Different Belief Facets on Agents' Decision—A Refined Cognitive Architecture to Model the Interaction Between Organisations' Institutional Characteristics and Agents' Behaviour. *Lect. Notes Comput. Sci.*, 12298 LNAI:133–155.

[2] R. Burke, N. Mattei, V. Grozin, A. Voida, and N. Sonboli. *Multi-agent Social Choice for Dynamic Fairness-aware Recommendation*. Association for Computing Machinery, Inc.

[3] A. Camisa, P. N. Köhler, M. A. Müller, G. Notarstefano, and F. Allgöwer. 2024. A distributed optimization algorithm for Nash bargaining in multi-agent systems. *IFAC-PapersOnLine*, 53(2):2684–2689.

[4] S. de Jong and K. Tuyls. 2024. Human-inspired computational fairness. *Auton. Agents Multi-Agent Syst.*, 22(1):103–126.

[5] J. Garg and S. Taki. 2024. An improved approximation algorithm for maximin shares. *Artif. Intell.*, 300.

[6] N. A. Grupen, B. Selman, and D. D. Lee. 2024. Cooperative Multi-Agent Fairness and Equivariant Policies. [Online; accessed 3. Feb. 2024].

[7] T. Kohler, J.-P. Steghofer, D. Busquets, and J. Pitt. 2024. The Value of Fairness: Trade-offs in Repeated Dynamic Resource Allocation. *Int. Conf. Self-Adaptive and Self-Organizing Syst., SASO*, 2014-December(December):1–10.

[8] G. Nicosia, A. Pacifici, and U. Pferschy. 2024. Price of Fairness for allocating a bounded resource. *Eur. J. Oper. Res.*, 257(3):933–943.

[9] Elinor Ostrom. 1990. *Governing the commons: The evolution of institutions for collective action*. Cambridge university press.

[10] A. Pla, B. López, and J. Murillo. 2024. Multidimensional fairness for auction-based resource allocation. *Knowl. Based. Syst.*, 73:134–148.

[11] Nicholas Rescher. 1967. Distributive Justice: A Constructive Critique of the Utilitarian Theory of Distribution. [Online; accessed 6. Feb. 2024].

[12] J. Tan, R. Khalili, H. Karl, and A. Hecker. 2024. Multi-Agent Distributed Reinforcement Learning for Making Decentralized Offloading Decisions. *Proc. IEEE INFOCOM*, 2022-May:2098–2107.