Fairness in Cooperative Multi-Agent Systems

MAS course project

Alma Mater Studiorum - Università di Bologna

February 2024

Introduction

This SLR provides an overview of fairness in MAS, emphasizing implementation and management

The scenarios discussed includes

Distributive Fairness in MAS
Focus on equitable resource allocation among autonomous agents, by means of diverse methodologies and models for ensuring fairness in resource distribution, both in dynamic and resource-constrained environments.

Procedural Fairness in MAS
Examining the fairness of decision-making processes, which are crucial for transparency, accountability, and avoiding biased practices.

Fairness and Multi-Objective Optimization
Exploration of the intersection between fairness and multi-objective optimization.
Addressing optimization challenges aiming to find a fair solution while upholding ethical standards among efficient alternatives.

Methodology

(G) How is fairness handled in MAS?

- 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?

Research Strategy

fairness

managing fairness

handling fairness

implement fairness

in multi agent/multiagent/multi-agent system(s)

10 papers

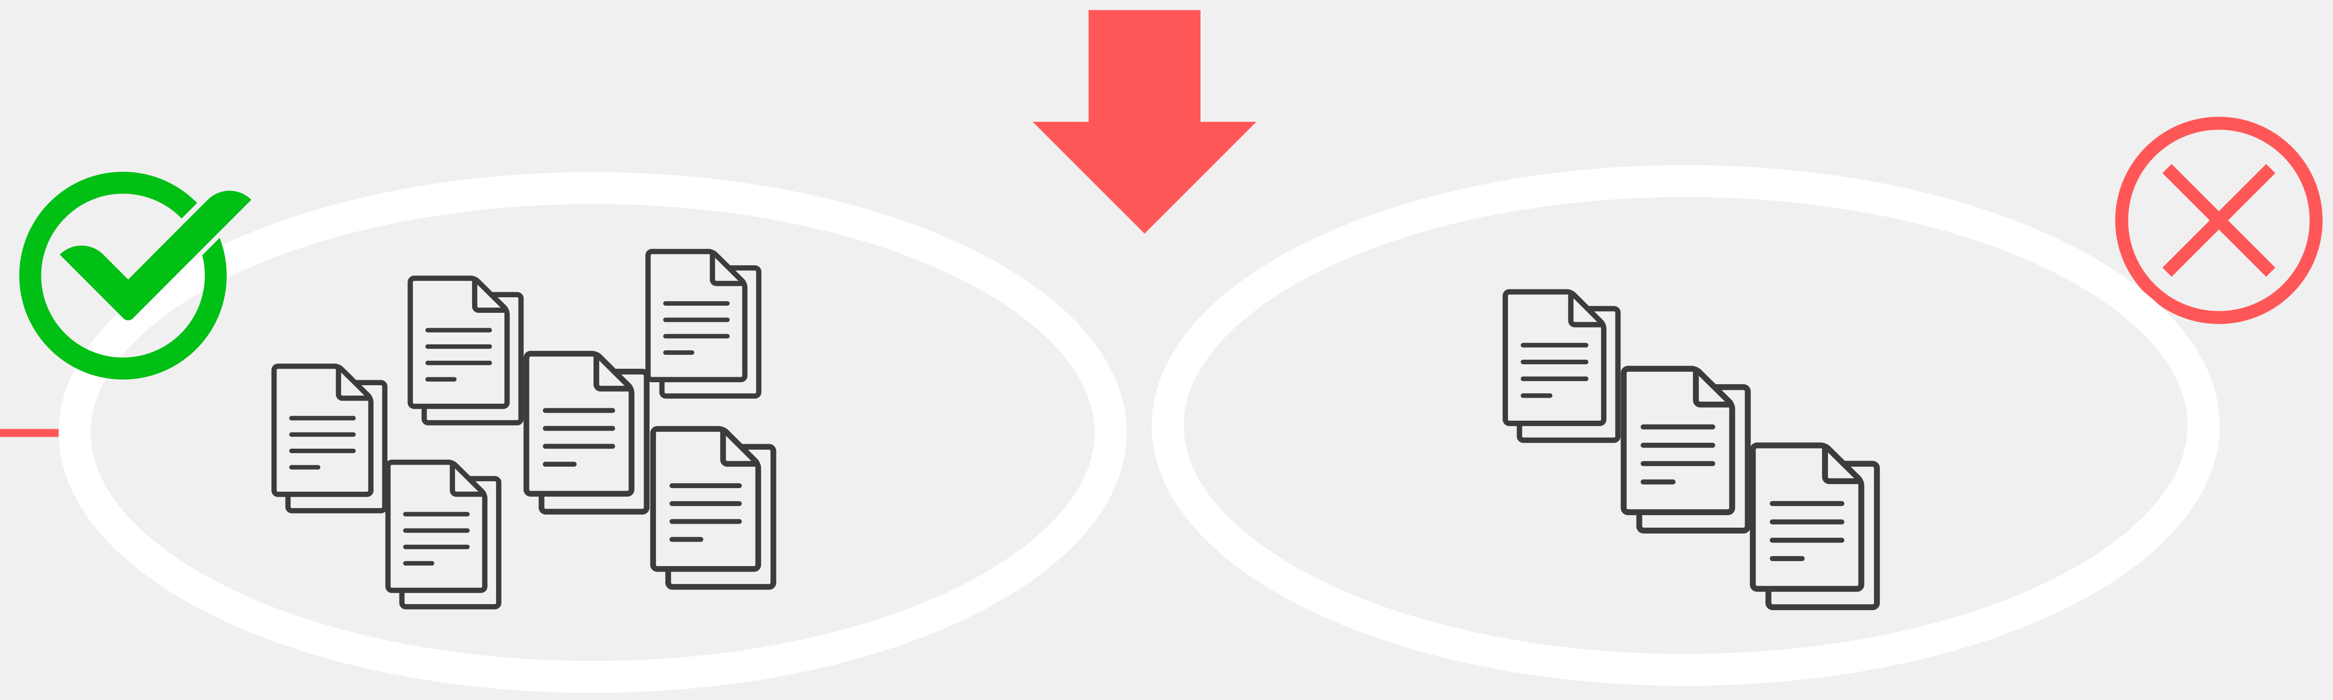

Notion of fairness within MAS

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.

Fairness in resource distribution

Price of Fairness for allocating a bounded resource

The paper 1 investigates the **Price of Fairness** (PoF), evaluating efficiency loss under fair allocation versus maximizing total utility. Using the **Fair Subset Sum Problem** (FSSP) as an illustration, the authors formalize properties for general multi-agent problems. In the FSSP, agents share a common bounded resource, selecting subsets of indivisible items to stay within capacity limits.

Equitable distribution and group-based fairness

Paper 2 introduces the concept of team fairness, highlighting the limitations of naive reward maximization for individual team members. The **Fairness through Equivariance** (Fair-E) method is proposed to enforce team fairness during policy optimization, with **Fairness through Equivariance Regularization** (Fair-ER) demonstrating improved utility and fair outcomes. Novel insights reveal that the trade-off's magnitude depends on agent skill.

Auction-based resource allocation

The paper 3 explores fairness through **auction-based** resource allocation. By means of a priority-based mechanism is addressed the bidder drop problem, enhancing social welfare, and maintain competitive market prices. Evaluation indicates that **multi-dimensional fairness** encourages agent participation and improves wealth distribution equity, without compromising allocation attribute quality.

Distribution of indivisible items

The paper 4 focuses on fair division of **indivisible items** among agents, using the **maximin share** (MMS) fairness criterion, which represents the maximum value an agent can guarantee when choosing last, offering a specific objective measure of fairness.
Not always the best MMS share always exist, but 2/3 MMS allocations are always possible, furthermore it' possible to improve till reaching a 3/4 MMS allocation.

Trade-offs in repeated dynamic resource allocation

In 5 is discussed a fair resource allocation scheme based on Ostrom's principles and Rescher's canons of distributive justice is proposed for managing fairness in system relying on **voluntary participation**. A variant of **Linear Public Goods Game** (LPG) is introduced to establish a fair resource allocation system for **open systems**, emphasizing that fairness enhances individual satisfaction, system stability, and reduces the risk of systemic failure.

Fairness in decision making

Impact of different belief facets

The paper 6 propose a refined cognitive architecture, incorporating **beliefs-desires-intentions** (BDI) and **theory of planned behavior** (TPB) models. Fairness, as a parameter, focuses on how individual agents punish each other. Belief facets, such as internal acceptance and perceived norms model how institutional characteristics influence decisions. Personality has an high influence, highlighting the pivotal role of internal beliefs in adhering to institutional rules.

Human inspired fairness

In 7 is exposed a sanctioning mechanisms inspired by human behavior, designed explicitly for **social dilemmas**. It aims to address the trade-off between group and individual rewards. The model is characterized as decentralized and robust against exploitation. It **surpasses limitations** of existing models by acknowledging the variability in adherence to fairness in open MAS,

Multi-agent social choice

In 8 the conventional "user" definition is expanded to include various stakeholders, and is introduced a dynamic architecture representing a **social choice environment**, capable of integrating multiple, possibly conflicting fairness concepts. It’s investigated the Recommender systems, as two-sided platforms, involve consumers, providers, and operators as direct stakeholders, with fairness concerns often requiring a balance among them.

Fairness in optimization

Optimization algorithm for nash bargaining

In 9, focusing on a **Nash bargaining** solutions, is addressed a **multi-objective optimization** problems over networks. It's discussed a distributed algorithm to find a fair solution among Pareto optimal solutions. The Nash bargaining solution is crucial for fairness, preventing unweighted sums from favoring some systems. The algorithm involves a negotiation mechanism and is demonstrated its efficacy through numerical simulations.

Reinforcement learning for optimizing offloading decisions

In 10, the paper presents a interaction mechanism based on sealed-bid auctions and **multi-agent reinforcement learning** (MARL), crucial in dynamic contexts like vehicular networks. The mechanism achieves optimal resource allocation, balancing fairness with efficiency. The MARL algorithm **enhances** system and individual **performance**, reducing offloading failure rate. Results confirm the effectiveness of the system beyond V2X scenarios, contributing to distributed autonomous systems.

Discussion

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

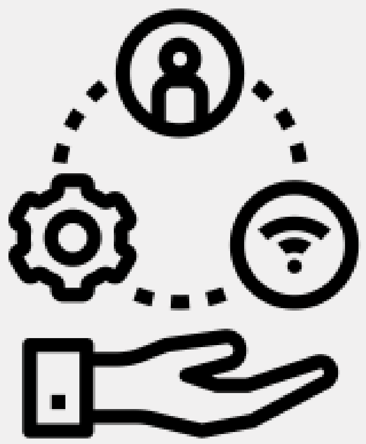

Metrics like PoF and MMS, methods like Fair-E and Fair-ER. Auction-based resource allocation. LPG variant and distributive justice.

Consider internal beliefs, rules and societal support.

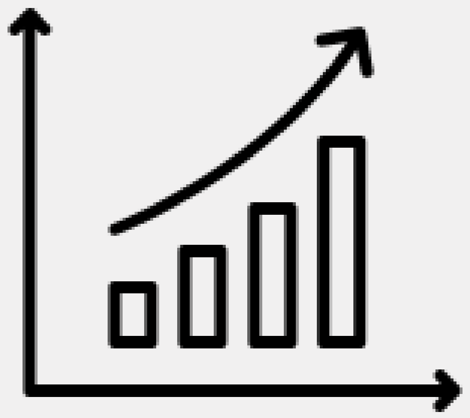

Cooperative game theory, specially Nash bargaining, and Reinforcement Learning

Is possible to define a concept of fairness that could be implemented in a MAS?

Fairness is defined and implemented in a MAS through various metrics and mechanisms as already mentioned in the goal answer. However is not possible to define a comprehensive concept of fairness that can be universally accepted, due to possible conflict between the goal of the individual and the goal of the group

What is the impact of fairness in the behaviour of agents?

Fairness plays a crucial role by highlighting how team fairness influences policy optimization in cooperative MAS, fair considerations impact agent participation and wealth distribution in resource allocation auctions, and the study of fairness in procedural processes (decision-making), incorporates cognitive elements like beliefs and societal support. Additionally, the impact of fairness in optimization scenarios, such as Nash bargaining, affects agents' negotiation mechanisms and their pursuit of equitable resource allocation.

Is always possible to ensure fairness in agent's 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 MAS. The impact of agent skill should also be considered, since fairness can be more expensive for lesser-skilled agents.

Bibliography

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

[2] N. A. Grupen, B. Selman, and D. D. Lee. 2024. Cooperative Multi-Agent Fairness and Equivariant Policies.

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

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

[5] 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

[6] 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.

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

[8] 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.

[9] 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.

[10] 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.