Computer Science and Information Systems 21(1):269–289 https://doi.org/10.2298/CSIS230119075L

How to Fairly and Efficiently Assign Tasks in Individually Rational Agents' Coalitions? Models and Fairness Measures

Marin Lujak[1,⋆], Alessio Salvatore[3], Alberto Fernández[1], Stefano Giordani[2], and Kendal Cousy[1,2]

[1] CETINIA, University Rey Juan Carlos, Madrid, Spain
{marin.lujak, alberto.fernandez, kendal.cousy}@urjc.es
[2] Dipartimento di Ingegneria dell'Impresa, University of Rome "Tor Vergata", Rome, Italy
stefano.giordani@uniroma2.it
[3] Istituto di Analisi dei Sistemi ed Informatica "A. Ruberti", Consiglio Nazionale delle Ricerche, Rome, Italy
alessio.salvatore@iasi.cnr.it

Abstract. An individually rational agent will participate in a multi-agent coalition if the participation, given available information and knowledge, brings a payoff that is at least as high as the one achieved by not participating. Since agents' performance and skills may vary from task to task, the decisions about individual agent-task assignment will determine the overall performance of the coalition. Maximising the efficiency of the one-on-one assignment of tasks to agents corresponds to the conventional linear sum assignment problem, which considers efficiency as the sum of the costs or benefits of individual agent-task assignments obtained by the coalition as a whole. This approach may be unfair since it does not explicitly consider fairness and, thus, is unsuitable for individually rational agents' coalitions. In this paper, we propose two new assignment models that balance efficiency and fairness in task assignment and study the utilitarian, egalitarian, and Nash social welfare for task assignment in individually rational agents' coalitions. Since fairness is a relatively abstract term that can be difficult to quantify, we propose three new fairness measures based on equity and equality and use them to compare the newly proposed models. Through functional examples, we show that a reasonable trade-off between efficiency and fairness in task assignment is possible through the use of the proposed models.

Keywords: Task Assignment, Multi-Agent Systems, Fairness , Efficiency, Resource Allocation, Multi-Agent Coordination

1. Introduction

An individually rational agent operates within a decision-making context, driven by self-interest to ensure that it attains a payoff or utility that is at least as favourable as its best alternative, including the option of not participating. When a coalition is formed by individually rational agents, they share a common objective, even though the agents themselves may possess distinct, potentially conflicting interests. They collaborate by pooling

⋆ Corresponding author

their capabilities and/or resources to efficiently carry out designated tasks, often achieving synergistic outcomes that surpass what they could accomplish individually. In this context, agents strive to achieve superior performance through collective action. Instances of such coalitions can be observed in various domains, including emergency services (e.g., [26]), agricultural cooperatives (e.g., [28]), taxi and ride-sharing services (e.g., [5, 27]), as well as smart grids (e.g., [36]), among others.

Of our concern in this paper is the one-on-one agent-task assignment in a coalition composed of individually rational agents. We assume a centralised decision-making process (or an algorithm) that is in charge of deciding which agent is assigned to each task. From a utilitarian point of view, an optimal solution would be the assignment that produces the lowest overall cost (or the highest benefit) for the coalition as a whole. However, the most efficient solution for the overall system may create large differences among the agents in terms of their individually assigned costs (we refer to this as an "unfair" assignment). The perception of an unfair task assignment solution may motivate unsatisfied agents to leave the coalition, putting the survivability of the same at risk. Thus, assignment decisions should be made based on minimising overall assignment costs and considering social welfare and fairness.

The linear-sum assignment problem, related to the topic of this paper, is a largely studied generally computationally easy problem that maximizes the efficiency of a multi-agent system without considering fairness. Exact solutions for this problem can be produced efficiently even for problem instances with a very large number of agents and tasks. However, to the best of our knowledge, related work on balancing fairness and efficiency in task assignment is scarce. Therefore, in this work, we explore the means of balancing the overall cost and fairness in task assignment in agent coalitions. These two aspects are generally opposed, i.e., solution approaches focusing on cost minimisation are likely to produce unfair assignments for some agents, while fair assignments may be far from the minimum cost solution for the coalition as a whole. In this paper, we study the trade-off between these two requirements and focus on finding task assignment solutions that are as fair as possible while not overly penalising the overall coalition's assignment cost. This implies finding efficient and fair assignments considering the distribution of individual costs among coalition members.

The main contributions of this paper are twofold. First, we propose three new fairness measures for a multi-agent system composed of self-concerned individually rational agents: Egalitarian Fairness Measure (EFM), Relative All-to-all Fairness measure (RAF), and Overall Relative Opportunity Cost Fairness (OROCF) measure. Then, we present two new one-on-one task assignment models that maximise the social welfare of the system while balancing efficiency and fairness: an envy-free utilitarian model that uses the utilitarian social welfare function while constraining the differences in the costs between agents, and the Nash model that optimises the Nash product of assigned tasks' benefits or costs of individual agents composing the system. We choose Nash social welfare due to its structure (being a product of costs) that explicitly balances efficiency and fairness.

The rest of the paper is organised as follows. In Section 2, we give an overview of the state of the art. In Section 3, we give motivation for this work and define the general problem of one-on-one task assignment. We propose new equality and equity fairness measures in Section 4 and introduce a linearised model for the calculation of one of the more computationally complex fairness measures. The two new proposed mathematical models for efficient and fair task assignment are presented in Section 5. Section 6 presents simple functional tests and discusses how the presented models differ based on the proposed fairness measures. Following these functional tests, in Section 7, we compile the results of many experiments to have a complete overview of the strengths and weaknesses of each model. In Section 8, we conclude the paper by giving an overview of the results and discuss the potential of the new proposed models and fairness measures to make a fair and efficient task assignment. We also give lines of future work to improve the proposed assignment models and fairness measures.

2. State of the Art

The assignment of resources, chores, or tasks in a multi-agent coalition may vary when defining fairness and efficiency depending on agents' mutual inter-dependencies and their relation with the coalition's objectives (e.g., [8, 11, 29]).

In this paper, we study balancing fairness and efficiency in the allocation of indivisible goods (resources, chores, or tasks) (e.g., [12]), and, more specifically, in the one-on-one assignment of tasks to agents in a cooperative multi-agent coalition composed of individually rational agents.

Cooperative decision making considers working toward a shared goal even though its ownership is not shared [33], as opposed to collaborative decision making, which considers a goal that is shared and owned by all agents in a coalition. Thus, cooperative decision making results are generally differentially beneficial to different agents [30], while collaboration is generally about equally sharing efforts, costs, and benefits.

Cooperative multi-agent task allocation problem was studied in, e.g., [18, 23, 27]. This problem has many different real-world applications where fairness can be a challenge. For example, in spatial crowdsourcing [42], there is a need to minimise the payoff difference among workers while maximising the average worker payoff. Similarly, in ride-share platforms, it was shown in [32] that, during high-demand hours, lacking any consideration of fairness and seeking only an optimal number of trips could lead to increased societal biases in the choice of the clients. This problem is relevant for many other applications including manufacturing and scheduling, network routing and the fair and efficient exploitation of Earth Observation Satellites, among others (e.g., [11]).

Various fairness measures exist for different contexts, e.g., machine learning (e.g., [14]), neural networks (e.g., [34]) and algorithm development (e.g., [20]). Fairness is studied as well in other contexts, like the multi-winner voting problem and recommender systems (e.g., [40]), but also in decision making (e.g., [37]). The importance of the individual perception of fairness within a system to keep individual satisfaction high is emphasised in [38]. In the Machine Learning context, the potential contradiction between individual and group fairness is studied in [6]. Some works study more generally the concepts of distributive justice, equality and equity (e.g., [13]).

The two most well-known fairness measures for the allocation of indivisible goods include max-min and proportional allocation. The max-min fairness aims to maximize the utility of the agent who contributes the least to the overall utility of the system. In other words, it ensures that the agent with the smallest contribution still receives as much utility as possible, while proportional fairness dictates that each agent should receive at least one-nth of the utility they would have obtained if they were the sole recipient of the goods. Max-min fairness is generalised in the case of resource allocation for systems with different resource types in [17] while max-min fairness, proportional fairness and balanced fairness are compared in the setting of a communication network of processor-sharing queues in [7]. Additionally, in resource allocation, there can be agents desiring some tasks (resources) more than others, or there can even be agents desiring tasks given to other agents, creating envy in the system (e.g., [11]). Envy-freeness criterion implies that an allocation should leave no agent envious of the other (e.g., [9]). However, it is not always enough to achieve envy-freeness for a fair solution (e.g., [2, 22]).

Balancing fairness and efficiency in divisible resource sharing was studied in, e.g., [1], while the related work on the competitive counterpart of cooperative systems usually studies finding a Pareto-optimal and fair allocation of indivisible items aiming at maximising (computationally expensive) Nash welfare (e.g., [4, 19, 41]). Proving the existence of such an allocation for an arbitrary number of agents is still an open problem [3].

Various social welfare concepts play an important role in balancing between efficiency and fairness. Their modelling and importance in enhancing the quality of task allocation are studied in [11]. In this work, we study egalitarianism and utilitarianism in this regard. Egalitarianism is a trend of thought in political philosophy that favours equality among the individuals composing the coalition no matter what their circumstances are (e.g., [16]). Utilitarianism, on the other hand, is a theory of morality that advocates actions that maximise happiness or well-being for all individuals while opposing the actions that cause unhappiness or harm. When directed toward making social and economic decisions, a utilitarian philosophy aims at the improvement of the coalition as a whole (e.g., [31]).

The Nash social welfare combines efficiency and fairness considerations. This function, or variants of it, are studied in literature considering, e.g., fairness in the ambulance location problem [21], and in allocating indivisible goods [10]. The multi-agent resource allocation problem considering Nash social welfare (the product of the utilities of the individual agents) is studied in [35].

This paper is an extended version of [15], where we previously studied the problem of balancing efficiency and fairness in linear-sum one-on-one task assignment. To the best of our knowledge, most of the related works treat efficiency in terms of proportionality in competitive multi-agent systems and propose computationally expensive models suitable for instances with a relatively small number of agents and tasks. In this paper, we propose two scalable computationally efficient models for task assignment in cooperative multi-agent systems, Nash model and envy-free utilitarian model, both with quality of solution guarantees. The proposed models balance efficiency in terms of overall system cost and the proposed fairness measures.

3. Motivation and problem definition

Most of the state-of-the-art literature on task assignment generally focuses on the efficiency of the assignment in terms of cost minimisation or benefit maximisation and does not consider fairness in the process, thus optimising only the system's overall general assignment cost or benefit (e.g., time, distance, monetary value, etc.). It is noteworthy that, from an optimisation standpoint, the task of maximising the overall benefit of multi-agent task assignments can be effectively converted into a cost-minimisation problem by simply inverting the values associated with agent-task assignment benefits.

The overall cost-minimisation approach in multi-agent task assignment is equivalent to optimising utilitarian social welfare function, a concept from welfare economy that sums the utility of each individual to obtain society's overall welfare (see, e.g., [11, 39]). All agents are treated the same, regardless of their initial level of utility or cost distribution among the tasks. This strategy is admissible in the case of a single decision maker, but might be unacceptable when multiple self-concerned and individually rational agents must mutually decide on the assignment of tasks.

Let us introduce a simple example showing how unfair a task assignment optimising a utilitarian social welfare function can be. Let us consider 3 self-concerned, individually rational agents (a_1, a_2, a_3) that need to be assigned in a one-on-one manner to a set of 3 tasks (k_1, k_2, k_3) and vice versa. The cost matrix containing the assignment costs for these agents and tasks is shown in Table 1a.

Table 1. Example of a cost matrix and different one-on-one task assignment solutions with minimum overall cost (in bold)

(a) Cost Matrix

	k_1	k_2	k_3
a_1	50	60	70
a_2	30	40	50
a_3	10	50	30

(b) Solution s_1

	k_1	k_2	k_3
a_1	50	**60**	70
a_2	**30**	40	50
a_3	10	50	**30**

(c) Solution s_2

	k_1	k_2	k_3
a_1	**50**	60	70
a_2	30	**40**	50
a_3	10	50	**30**

(d) Solution s_3

	k_1	k_2	k_3
a_1	50	60	**70**
a_2	30	**40**	50
a_3	**10**	50	30

By applying the conventional (linear-sum) task assignment model (i.e., the utilitarian social welfare model) that minimises the overall assignment cost of the system without considering fairness (see, e.g., [24, 27]), we might get the assignments (called solution s_1) marked in bold in Table 1b where agent a_1 is assigned to task k_2, agent a_2 is assigned to task k_1 and agent a_3 to task k_3. The overall (minimum) assignment cost found by this model is 120. However, if we focus on its cost distribution over individual agents, we see large discrepancies. Indeed, the cost of agent a_1 is 60, while the cost of a_2 (and a_3) is only 30. Thus, a_1 is charged twice more than a_2 (and a_3). In Table 1d (i.e., solution s_3), this difference is even larger resulting in a 7 times larger cost of the worst-off concerning the best-off agent. Generally, an upper bound on the difference in the assignment cost is the maximum value in a given cost matrix. In centralised systems, where agents are owned and controlled by a single decision-maker, this would not cause any problems. However, in the case of decentralised systems composed of self-concerned and individually rational agents, such an unfair solution might result in the worst-off agents leaving the system due to the lack of fairness in the solution.

Table 1c shows a fairer solution (that we name s_2) where the costs of the agents are as close as possible, thus minimising the envy of agents. This is an ideal situation regarding fairness in this case where all agents are assigned tasks of similar costs. Notice that, here we did not sacrifice efficiency to achieve balanced individual costs. In case of repetitive task allocations, the assignments can be altered to further facilitate balance in the accumulated assignment costs.

Problem definition. Given are a set of agents $a \in A$ and a set of tasks $k \in K$ that form a weighted complete bipartite graph $G = (A \bigcup K, E)$ with edge set $E = A \times K$ and with given edge weights c_{ak} on each edge $(a, k) \in E$, where c_{ak} is the cost of assigning task k to agent a, for all $a \in A$ and $k \in K$. W.l.o.g, we assume that the cardinality of the two sets is equal, i.e., $|A| = |K|$. In the case of unequal cardinality, we add a sufficient number of dummy vertices and assume that $c_{ak} = 0$ where $a \in A$ or $k \in K$ are dummy vertices. Moreover, for simplicity, we assume that agents are indexed from 1 to $|A|$, i.e., $A = \{1, \ldots, |A|\}$. The objective is to assign agents $a \in A$ to tasks $k \in K$ in a one-on-one manner and, therefore, find a perfect matching among vertices in A and vertices in K considering both assignment efficiency and fairness. An edge (a, k) is matched if its (two) extreme vertices $a \in A$ and $k \in K$ are mutually matched, and matching is perfect if every vertex in A is matched (assigned) to exactly one vertex in K, and vice versa. The following is the mathematical formulation of these constraints.

$$\sum_{k \in K} x_{ak} = 1, \forall\, a \in A \quad (1) \qquad \sum_{a \in A} x_{ak} = 1, \forall\, k \in K \quad (2)$$

$$x_{ak} \in \{0, 1\}, \forall\, a \in A, \forall\, k \in K \tag{3}$$

where x_{ak} is a binary decision variable equal to 1 if agent $a \in A$ is assigned to task $k \in K$, and zero otherwise. Constraints (1) and (2) assure that there is one-on-one assignment for each agent $a \in A$ and task $k \in K$, respectively. Constraints (3) fix the ranges of the decision variables.

4. Proposed fairness measures

In this section, we propose three new fairness measures for quantifying the balance between fairness and efficiency in task assignment from the egalitarian and equity points of view. All the fairness measures are fractions ranging between 0 and 1. We avoid the division by 0 in some extreme cases by adding a very small number ϵ (e.g., $\epsilon = 1e^{-10}$) to both the numerator and the denominator of these fractions.

The proposed fairness measures should be computed only for non-dummy vertices in the bipartite graph that represents the agents and tasks to ensure that these measures can still reach either the value of 0 or 1.

4.1. Egalitarian Fairness Measure (EFM)

Egalitarian Fairness Measure (EFM) focuses on the assignment cost faced by the worst-off agent (i.e., the agent with the highest assignment cost in a given feasible solution). Given the assignments x_{ak}^{sol},with $a \in A$ and $k \in K$, of a feasible solution sol, EFM is computed as follows:

$$EFM(sol) = \frac{c_{max} - c_{sol}^{wo} + \epsilon}{c_{max} - c_{min}^{wo} + \epsilon} \tag{4}$$

where $c_{max} = max_{a \in A, k \in K}\{c_{ak}\}$ is the maximum value in the cost matrix, $c_{sol}^{wo} = \max_{a \in A}\{\sum_{k \in K} c_{ak} x_{ak}^{sol}\}$ is the cost paid by the worst-off agent in the given solution, and c^{wo}_{min} is the minimum (or the preferred) cost that the same agent could pay for task assignment. In particular, c^{wo}_{min} is the optimal solution of the given mathematical problem:

$$c^{wo}_{min} = \min \lambda \tag{5}$$

s.t. (1)–(3) and

$$\sum_{k \in K} c_{ak} x_{ak} \leq \lambda, \forall\, a \in A \tag{6}$$

$$\lambda \geq 0 \tag{7}$$

where constraints (6) impose the upper limit on the cost (λ) paid by the worst-off agent, and Constraints (7) fix the range of the additional variable λ. When the worst-off assigned cost c^{wo}_{sol} is equal to c_{max}, $EFM(sol)$ will equal zero (ignoring ϵ). On the other hand, when c^{wo}_{sol} is equal to c^{wo}_{min}, $EFM(sol)$ will equal one; moreover, this also occurs when there exists an agent $a \in A$ such that $c_{ak} = c_{max}$, for all $k \in K$.

For the cost matrix given in Table 1a, where $c_{max} = 70$ and $c^{wo}_{min} = 50$, we calculate the $EFM(sol)$ for each solution reported in Tables 1b–1d. All the solutions reported in Table 1, have a minimum overall assignment cost equal to 120, while the values of c^{wo}_{sol} are $c^{wo}_{s_1} = 60$, $c^{wo}_{s_2} = 50$, and $c^{wo}_{s_3} = 70$, for the solutions reported in Table 1b, Table 1c, and Table 1d, respectively. $EFM(sol)$ value for these solutions are: $EFM(s_1) = \frac{70-60}{70-50} = 0.5$, $EFM(s_2) = \frac{70-50}{70-50} = 1$, and $EFM(s_3) = \frac{70-70}{70-50} = 0$.

According to the EFM measure definition, solution s_2 is the fairest one. Note that the increase in EFM value in solution s_2 corresponds to a distribution of the costs that leaves the worst-off agent better off than in s_1, and that solution s_3 leaves the worst-off agent with the worst possible (highest) cost. Note that, generally, there may be multiple such distributions.

4.2. Relative All-to-all Fairness (RAF) measure

Relative All-to-all Fairness (RAF) measure evaluates fairness at a society level by taking into account every agent's assignment in comparison with the others. The measure considers the squared differences of the assignment costs of each agent concerning the costs of all the others, as seen in Equation (8).

$$w_{sol} = \sum_{a \in A} \sum_{a' \in A | a' > a} \Big(\sum_{k \in K} c_{ak} x^{sol}_{ak} - c_{a'k} x^{sol}_{a'k}\Big)^2 \tag{8}$$

Then, relative all-to-all fairness is computed as follows:

$$RAF(sol) = \frac{w_{max} - w_{sol} + \epsilon}{w_{max} - w_{min} + \epsilon}, \tag{9}$$

where w_{max} and w_{min} represent the maximum and the minimum value for the RAF fairness measure that should be calculated a priori for a specific data set. They are modelled as follows:

$$w_{min} = \min \sum_{a \in A} \sum_{a' \in A | a' > a} \Big(\sum_{k \in K} c_{ak} x_{ak} - c_{a'k} x_{a'k}\Big)^2, \tag{10}$$

and

$$w_{max} = \max \sum_{a \in A} \sum_{a' \in A | a' > a} \big(\sum_{k \in K} c_{ak} x_{ak} - c_{a'k} x_{a'k} \big)^2, \tag{11}$$

both s.t. (1)–(3).

For the cost matrix given in Table 1a, the two components of RAF that are independent of the assignment solution are $w_{max} = 5400$ and $w_{min} = 0$, related to solutions s_{max}, with $x_{13}^{s_{max}} = x_{22}^{s_{max}} = x_{31}^{s_{max}} = 1$, and s_{min}, with $x_{11}^{s_{min}} = x_{23}^{s_{min}} = x_{32}^{s_{min}} = 1$, respectively. The values w_{sol} for the solutions reported in Table 1b, 1c and 1d are $w_{s_1} = 1800$, $w_{s_2} = 600$, and $w_{s_3} = 5400$, respectively. Related RAF values are: $RAF(s_1) = \frac{5400-1800}{5400-0} = 0.67$, $RAF(s_2) = \frac{5400-600}{5400-0} = 0.89$, and $RAF(s_3) = \frac{5400-5400}{5400-0} = 0$.

Also according to the RAF measure, solution s_2 is the fairest one and the order of the three solutions is the same as for EFM. This is not surprising as both measures evaluate the equality of a solution. However, s_2 is not the absolute fairest solution which, concerning this indicator, is $x_{11} = x_{23} = x_{32} = 1$ where all the agents pay the same cost; in this case, the RAF value is equal to 1. This is also not surprising as this particular measure considers not only the worst-off agent, but all of them, therefore making it less likely that one of the solutions with minimum cost also has the highest fairness value.

Linearisation of the RAF measure. It is computationally expensive to find the maximum value w_{max} and the minimum value w_{min} for the RAF fairness measure since their models (10) and (11), respectively, are composed of quadratic terms. This is especially the case in larger problem instances.

To fix this issue, in this section, we simplify the RAF measure by making a linear model that will be easier to solve. We first modify the RAF formula by exchanging the quadratic expressions with the absolute ones and then linearise the latter ones.

RAF measure with absolute values. We present in the following the proposed modified models, replacing the quadratic terms with absolute ones, subject to the same constraints as before.

$$w'_{min} = \min \sum_{a \in A} \sum_{a' \in A | a' > a} \big| \sum_{k \in K} c_{ak} x_{ak} - c_{a'k} x_{a'k} \big|, \tag{12}$$

and

$$w'_{max} = \max \sum_{a \in A} \sum_{a' \in A | a' > a} \big| \sum_{k \in K} c_{ak} x_{ak} - c_{a'k} x_{a'k} \big|, \tag{13}$$

both s.t. (1)–(3).

Minimising the RAF measure with linearised absolute values. In order to determine the optimal value of w'_{min} in (12) subject to (1)–(3), we linearise objective function (12) by adding new constraints and continuous free variables $r_{aa'}$, for all $a, a' \in A$, with $a < a'$, as follows.

$$w'_{min} = \min \sum_{a \in A} \sum_{a' \in A | a' > a} r_{aa'}, \tag{14}$$

s.t. (1)–(3), and

$$r_{aa'} \geq \sum_{k \in K} [c_{ak}x_{ak} - c_{a'k}x_{a'k}], \forall\, a, a' \in A, \text{ with } a < a', \tag{15}$$

$$r_{aa'} \geq \sum_{k \in K} [c_{a'k}x_{a'k} - c_{ak}x_{ak}], \forall\, a, a' \in A, \text{ with } a < a', \tag{16}$$

Maximising the RAF measure. In order to determine the optimal value of w'_{max} in (13) subject to (1)–(3), next, we add other constraints and binary variables $y_{aa'}$, $\forall\, a, a' \in A$, with $a < a'$, in our problem. We let $M = 2 \cdot \max_{a \in A, k \in K}\{c_{ak}\}$ be a parameter that is introduced for constraints (20) and (21) to select the largest of the two possible terms. We present next the modified problem.

$$w'_{max} = \max \sum_{a \in A} \sum_{a' \in A | a' > a} r_{aa'}, \tag{17}$$

s.t. (1)–(3), and

$$r_{aa'} \geq \sum_{k \in K} [c_{ak}x_{ak} - c_{a'k}x_{a'k}], \forall\, a, a' \in A \text{ with } a < a', \tag{18}$$

$$r_{aa'} \geq \sum_{k \in K} [c_{a'k}x_{a'k} - c_{ak}x_{ak}], \forall\, a, a' \in A \text{ with } a < a', \tag{19}$$

$$r_{aa'} \leq \sum_{k \in K} [c_{ak}x_{ak} - c_{a'k}x_{a'k}] + M \times y_{aa'}, \forall\, a, a' \in A \text{ with } a < a', \tag{20}$$

$$r_{aa'} \leq \sum_{k \in K} [c_{a'k}x_{a'k} - c_{ak}x_{ak}] + M \times (1 - y_{aa'}), \forall\, a, a' \in A \text{ with } a < a', \tag{21}$$

$$y_{aa'} \in \{0, 1\}, \forall\, a, a' \in A \text{ with } a < a', \tag{22}$$

4.3. Overall Relative Opportunity Cost Fairness (OROCF)

Overall Relative Opportunity Cost Fairness (OROCF) focuses on evaluating equity among the agents by taking into account the missed opportunities in terms of the assignment cost for each agent. The opportunity cost (e.g., [25]) is the concept in the microeconomics of lost benefit that would have been derived by an agent from an option not chosen. As the reference value, we consider a task of the minimum cost and normalise the difference in the cost value between the assigned task and the best-off task (the task with minimum cost) over the amplitude of costs for each agent, as seen in Equation (23).

$$y_{sol} = \sum_{a \in A} \frac{\sum_{k \in K} c_{ak} x_{ak}^{sol} - \min_{k \in K} \{c_{ak}\} + \epsilon}{\max_{k \in K} \{c_{ak}\} - \min_{k \in K} \{c_{ak}\} + \epsilon} \tag{23}$$

$$OROCF(sol) = \frac{y_{max} - y_{sol} + \epsilon}{y_{max} - y_{min} + \epsilon}, \tag{24}$$

where y_{max}, y_{min} represents the maximum and the minimum value of Equation (23) given Constraints (1)–(3).

The values y_{sol} for the solutions reported in Tables 1b, 1c and 1d are $y_{s_1} = 1$, $y_{s_2} = 1$, and $y_{s_3} = 1.5$, respectively. For the cost matrix given in Table 1a, the two components of OROCF that are independent of the assignment solution are $y_{max} = 2$, for $x_{13}^{s_{max}} = x_{21}^{s_{max}} = x_{32}^{s_{max}} = 1$, and $y_{min} = 1$, for $x_{11}^{s_{min}} = x_{22}^{s_{min}} = x_{33}^{s_{min}} = 1$. Related $OROCF$ values are: $OROCF(s_1) = \frac{2-1}{2-1} = 1$, $OROCF(s_2) = \frac{2-1}{2-1} = 1$, and $OROCF(s_3) = \frac{2-1.5}{2-1} = 0.5$.

Note that $OROCF$ value is the highest both for s_1 and s_2, meaning that these solutions offer the lowest highest opportunity cost for the sum of all agents. The reader can verify that the solution $x_{11} = x_{23} = x_{32} = 1$ would be the worst choice for agents a_2 and a_3 and would give a value of $OROCF$ equal to 0.

5. Proposed models considering fairness and efficiency

In this section, we propose new models that mitigate the equity issues posed by the classical linear-sum assignment model (e.g., [8]) and achieve a solution that is as fair as possible while sacrificing as little as possible the overall system's efficiency.

5.1. Nash Model for task assignment

The proposed Nash Model is inspired by the Nash social welfare function, a well-studied social welfare function in which the goal is to maximize the product of the utility functions of the agents composing the system. The proposed model is given next:

$$\min \prod_{a \in A} \sum_{k \in K} c_{ak} x_{ak} \tag{25}$$

s.t. (1)–(3). Since Eq. (25) is a nonlinear objective function, solving the above problem is computationally expensive. We propose next its linearised equivalent, which is possible due to the one-on-one assignment constraints (1)–(3).

$$\max \sum_{a \in A} \sum_{k \in K} \log(M - c_{ak}) x_{ak} \tag{26}$$

s.t. (1)–(3), where $M > \max_{k \in K, a \in A} \{c_{ak}\}$.

5.2. Envy-free utilitarian model for task assignment

We propose the envy-free utilitarian model that focuses both on efficiency and fairness. We introduce the fairness variable f_u to ensure that all the costs for each agent are inside a certain interval that shrinks as f_u becomes smaller. The model is defined as follows:

$$\min \ \alpha f_u + (1-\alpha)\frac{\sum_{k \in K}\sum_{a \in A} c_{ak}x_{ak}}{|A|} \tag{27}$$

s.t. (1)–(3), and

$$\sum_{k \in K} c_{ak}x_{ak} - \frac{\sum_{k \in K}\sum_{a \in A} c_{ak}x_{ak}}{|A|} \leq f_u, \forall\, a \in A \tag{28}$$

$$f_u \geq 0, \tag{29}$$

where fairness weight α in objective function (27) ranges between 0 and 1 and is used to weigh the fairness f_u and the average cost paid by the agents' coalition $\sum_{k \in K}\sum_{a \in A} c_{ak}x_{ak}/|A|$; when $\alpha = 0$, the model only considers the cost without considering fairness and vice versa for its value equal to 1. Constraints (28) guarantee that, for each agent, the difference between the cost of its assigned task and the average of the costs of the assigned tasks for all the agents is less than the value f_u. Constraint (29) fixes the range of variable f_u.

6. Functional tests

To evaluate the performance of the Nash model and the Envy-free Utilitarian model, we randomly generate three cost matrices (Table 2) with costs ranging from 1 to 1000. The models were solved for each matrix using IBM ILOG CPLEX Optimization Studio 20.0.1.

Table 2. Example cost matrices

(a) Functional test 1

	k_1	k_2	k_3
a_1	382	816	366
a_2	846	544	175
a_3	578	824	526

(b) Functional test 2

	k_1	k_2	k_3
a_1	450	895	358
a_2	856	233	449
a_3	890	672	976

(c) Functional test 3

	k_1	k_2	k_3
a_1	683	170	699
a_2	943	364	894
a_3	557	741	127

To compare the efficiency of the models presented in Section 5, we calculate the following normalised efficiency indicator (Eff):

$$Eff(sol) = \frac{z_{max} - z_{sol} + \epsilon}{z_{max} - z_{min} + \epsilon} \tag{30}$$

where $z_{sol} = \sum_{k \in K}\sum_{a \in A} c_{ak}x_{ak}^{sol}$ with x_{ak}^{sol} being the solution returned by the considered model. The values z_{max} and z_{min} are, respectively, the maximum and the minimum values of $\sum_{k \in K} \sum_{a \in A} c_{ak} x_{ak}$ given Constraints (1)–(3). Table 3 shows the results of our experiments.

Table 3. Results and Comparison

Functional test 1

Model	Eff	EFM	RAF	$OROCF$
Nash	0.91	1	1	1
Envy-free ($\alpha = 0$)	1	0.07	0	0.68
Envy-free ($\alpha \geq 0.5$)	0.91	1	1	1

Functional test 2

Model	Eff	EFM	RAF	$OROCF$
Nash	0.93	1	0.91	1
Envy-free ($\alpha = 0$)	1	0.28	0.17	0.92
Envy-free ($\alpha = 0.5$)	0.93	1	0.91	1
Envy-free ($\alpha \geq 0.9$)	0	0	1	0

Functional test 3

Model	Eff	EFM	RAF	$OROCF$
Nash	0.93	1	0.73	1
Envy-free ($\alpha = 0$)	1	0	0	0.99
Envy-free ($\alpha = 0.5$)	0.93	1	0.73	1
Envy-free ($\alpha = 0.7$)	0.59	0.94	0.93	0.64
Envy-free ($\alpha \geq 0.9$)	0.05	0.19	1	0.06

The case when $\alpha = 0$ corresponds to the case when we are optimising the global cost only (utilitarian social welfare function). We get very low values of fairness for this case according to our prior assumptions. It is interesting to notice similarities when we set the α value to 0.5. Indeed, in that case, the Envy-free Utilitarian model and the Nash model have the same behaviour and give us the same solutions. These solutions for $\alpha = 0.5$ are ideal for the fairness indicators EFM and $OROCF$ in our three tests, while RAF also increases significantly. Moreover, the efficiency (Eff) is greater than 0.9. Equality and equity can be improved without a significant decrease in efficiency. We notice in tests 2 and 3 that, generally, the higher the value of α, the lower the overall system's efficiency. This shows that striving for too much equality can be highly detrimental to the system's efficiency and even equity. The results for the cost matrix in Table 1a also support this claim in case $\alpha = 1$. Here, allocation $x_{11} = x_{23} = x_{32} = 1$ is an egalitarian allocation that decreases efficiency and equity simultaneously since agents a_2 and a_3 are allocated to their worst-off tasks and the overall allocation cost is 150 instead of the minimum cost of 120.

7. Simulation experiments

In this section, we perform simulation experiments to evaluate how the proposed models behave from the scalability standpoint.

For these experiments, 20 cost matrices $C = \{c_{ak} | a \in A, k \in K\}$ where $|A| = |K|$ of size 5, 10 and 20 were created, with each cost c_{ak} randomly obtaining a value from 1 to

Table 4. Results depending on problem size

Size 5

Model	Eff	EFM	RAF	$OROCF$	z
Nash	1.00	0.95	0.77	0.99	1402
Envy-free ($\alpha = 0$)	1.00	0.89	0.75	1.00	1399
Envy-free ($\alpha = 0.25$)	0.99	0.96	0.77	0.99	1405
Envy-free ($\alpha = 0.5$)	0.94	1.00	0.82	0.94	1520
Envy-free ($\alpha = 0.75$)	0.73	0.89	0.91	0.73	1913
Envy-free ($\alpha = 1$)	0.41	0.62	0.94	0.41	2513

$z_{min} = 1339$; $z_{max} = 3408$

Size 10

Model	Eff	EFM	RAF	$OROCF$	z
Nash	1.00	0.95	0.87	1.00	1439
Envy-free ($\alpha = 0$)	1.00	0.93	0.87	1.00	1438
Envy-free ($\alpha = 0.25$)	1.00	0.99	0.88	1.00	1466
Envy-free ($\alpha = 0.5$)	0.97	1.00	0.89	0.98	1626
Envy-free ($\alpha = 0.75$)	0.87	0.98	0.96	0.88	2330
Envy-free ($\alpha = 1$)	0.47	0.59	0.99	0.48	5235

$z_{min} = 1438$; $z_{max} = 8602$

Size 20

Model	Eff	EFM	RAF	$OROCF$	z
Nash	1.00	0.95	0.94	1.00	1522
Envy-free ($\alpha = 0$)	1.00	0.95	0.94	1.00	1522
Envy-free ($\alpha = 0.25$)	1.00	1.00	0.96	1.00	1610
Envy-free ($\alpha = 0.5$)	0.98	1.00	0.95	0.99	1789
Envy-free ($\alpha = 0.75$)	0.92	0.99	0.98	0.92	2924
Envy-free ($\alpha = 1$)	0.46	0.52	0.97	0.46	10687

$z_{min} = 1522$; $z_{max} = 18473$

1000 following a uniform law of probability. This results in a total of 60 instances that we test in the experiments in Table 4. For the envy-free utilitarian model, α takes the values of 0 (the classical utilitarian model), 0.25, 0.5, 0.75 and 1 (considering only fairness f_u).

Table 4 lists the average results for size $n = 5, 10, 20$, respectively, where z is the average value of the sum of the assignment costs. It is to be noted that RAF this time is calculated using the RAF minimisation and maximisation models from section 4.2. In addition, the z_{min} and z_{max} values from the efficiency formula Eff (30) are given just after the sub-table for each problem size.

Experiments results considering only lower bound. Table 5 presents the values of our four fairness measures as a ratio of their absolute values to their lower bounds. The experiments in this section were conducted on randomly generated instances of varying sizes: 5, 10, 20, 50, and 100, with 20 instances for each size.

This allows for a comparison of the models based solely on their relative performance to the lower bound, rather than taking into account both bounds. In this way, the difference between the models is clearer compared to the results of Table 4.

The fairness measures are calculated as follows:

Table 5. Results depending on problem size with only lower bounds

Size 5

Model	Eff'	EFM'	RAF'	$OROCF'$
Nash	1.00	1.05	2.17	1.02
Envy-free ($\alpha = 0$)	1.00	1.09	2.25	1.02
Envy-free ($\alpha = 0.25$)	1.01	1.04	2.15	1.04
Envy-free ($\alpha = 0.5$)	1.11	1.00	1.90	1.67
Envy-free ($\alpha = 0.75$)	1.42	1.04	1.41	3.52
Envy-free ($\alpha = 1$)	1.92	1.30	1.21	12.72

Size 10

Model	Eff'	EFM'	RAF'	$OROCF'$
Nash	1.00	1.11	1.93	1.00
Envy-free ($\alpha = 0$)	1.00	1.15	1.96	1.01
Envy-free ($\alpha = 0.25$)	1.02	1.02	1.86	1.05
Envy-free ($\alpha = 0.5$)	1.13	1.00	1.76	1.35
Envy-free ($\alpha = 0.75$)	1.63	1.03	1.32	2.94
Envy-free ($\alpha = 1$)	3.87	2.00	1.10	10.37

Size 20

Model	Eff'	EFM'	RAF'	$OROCF'$
Nash	1.00	1.22	1.63	1.00
Envy-free ($\alpha = 0$)	1.00	1.22	1.63	1.00
Envy-free ($\alpha = 0.25$)	1.06	1.01	1.49	1.15
Envy-free ($\alpha = 0.5$)	1.17	1.00	1.52	1.51
Envy-free ($\alpha = 0.75$)	1.91	1.04	1.23	3.58
Envy-free ($\alpha = 1$)	7.23	3.02	1.39	19.89

Size 50

Model	Eff'	EFM'	RAF'	$OROCF'$
Nash	1.00	1.24	1.15	1.02
Envy-free ($\alpha = 0$)	1.00	1.30	1.16	1.02
Envy-free ($\alpha = 0.25$)	1.03	1.01	1.09	1.09
Envy-free ($\alpha = 0.5$)	1.33	1.00	1.23	2.05
Envy-free ($\alpha = 0.75$)	2.22	1.05	1.03	4.80
Envy-free ($\alpha = 1.00$)	30.50	11.57	1.03	86.73

Size 100

Model	Eff'	EFM'	RAF'	$OROCF'$
Nash	1.00	1.25	1.07	1.00
Envy-free ($\alpha = 0$)	1.00	1.15	1.08	1.00
Envy-free ($\alpha = 0.25$)	1.01	1.00	1.05	1.04
Envy-free ($\alpha = 0.5$)	1.53	1.00	1.48	2.53
Envy-free ($\alpha = 0.75$)	2.41	1.48	1.98	3.77
Envy-free ($\alpha = 1.00$)	56.67	16.84	1.02	156.86

- $EFM' = c^{wo}_{sol}/c^{wo}_{min}$
- $RAF' = w_{sol}/w_{min}$
- $OROCF' = y_{sol}/y_{min}$
- $Eff' = z_{sol}/z_{min}$,

where the used terminology is explained in Section 4. The values for these metrics are in the range $[1.0, +\infty]$, with 1.0 indicating the best performance. The table shows the average for the values obtained on 20 samples.

7.1. Interpretation of the results

The different models offer a good choice to pick from depending on the desired goal. In the case that the system is aiming at giving the agents a similar cost, then using the envy-free utilitarian model and increasing the value of α will considerately help. If however, the system is aiming at an all-round and fast solution for allocating the tasks with both equity and equality taken into account, then the Nash model will work very well. Indeed, if we compare the Nash model to the Utilitarian model ($\alpha = 0$), which gives us the best solution for the system taken as a whole in terms of the cost paid by the agents, for each size we can notice that all the fairness measures are better, closer to 1.0. What is even more noticeable is that while being fairer overall, the efficiency of the Nash model is really good, with the global cost of the solutions being roughly just as low as the Utilitarian model ($\alpha = 0$).

As we can see in Table 5, the Envy-Free Utilitarian Model (EFUM) decreases its *efficiency* as α increases, since a higher α gives more importance to the fairness (envy-free) part. If $\alpha = 0$, only the utilitarian part is considered, thus the optimal utilitarian solution is obtained ($Eff' = 1$). The Nash model turns out to be optimal, from the utilitarian point of view, in most experiments (on average, $Eff' = 1$).

The *egalitarian fairness metric (EFM)* measures how much the cost of the worst-off agent is kept as low as possible in a given agent-task assignment. Total envy-freeness ($\alpha = 1$) is not the best option because it is possible to obtain higher differences among assigned costs (lower envy-freeness) even at a lower cost for the worst-off agent. Conversely, the optimal utilitarian solution ($\alpha = 0$) does not consider the worst-off agent cost at all and obtains a much higher EFM' value. The best result from the egalitarian (EFM) point of view is obtained by the EFUM with α values near 0.5, i.e. with a balance between the fairness and utilitarian perspectives. The Nash model behaves quite similarly to the total envy-freeness ($\alpha = 1$) since it obtains the optimal solution in most of our experiments.

Regarding *RAF*, while it seems that increasing α for the EFUM gives better results with smaller problem sizes, it is not consistent with size 20 ($\alpha = 0.5$ and $\alpha = 1$), size 50 ($\alpha = 0.5$) and size 100 ($\alpha = 0.5$ and $\alpha = 0.75$). The Nash model, as mentioned above, behaves very similarly to EFUM with $\alpha = 0$, thus it obtains bad results from the *RAF* measure point of view.

With regards to *OROCF* measure, EFUM is better for smaller α, i.e. in case we give more importance to efficiency than to envy-freeness. We assume this is because *OROCF* measures the loss against the most efficient solution from each local viewpoint. Since the Nash model obtains the optimal solution in most of the cases, its $OROCF'$ value can be also very good (1).

7.2. Run time comparison

Table 6 shows CPLEX solving times (in seconds) we obtain when comparing the initial quadratic version of RAF and the absolute version of RAF, for both minimising and maximising the related evaluation functions. The values shown are an average of 20 samples for each size $n = 5, 8, 10$.

Table 6. RAF run time comparison depending on size

	Quadratic RAF run time (s)		Absolute RAF run time (s)	
size	Min	Max	Min	Max
5	0.066	0.113	0.035	0.110
8	1.329	2.575	0.075	2.323
10	354.104	456.932	0.368	57.183

We can see a huge improvement (reduction) in the time it takes for the solver to find the solution for both the minimisation and the maximisation of the modified RAF measure with absolute terms in comparison with the one with quadratic terms. This difference is more noticeable with matrices of bigger size.

Now, looking at the average solving time for our models on the graphs of Fig. 1, tested with four values of alpha for the envy-free model which are 0.25, 0.5, 0.75 and 1.0, we can see that using the Utilitarian and Nash models we get the fastest results, with roughly half a second of running time even for problems of size 200.

When solving the Envy-free model, we observe an exponential increase in the solving time, with problems not being solved optimally after one hour for even small sizes of the problem when alpha gets close to 1.

Stars in yellow in the graph indicate problem sizes for which CPLEX was unable to find an optimal solution within an hour. In such instances, we used the relative Mixed Integer Programming (MIP) gap to assess the proximity of the found solution to the optimal solution. The MIP gap serves as an indicator of the proximity of the current solution to the optimal solution. It is defined as the difference between the upper bound (representing the best-known upper bound on the objective value obtained thus far) and the lower bound (representing the best-known lower bound on the objective value obtained thus far), divided by the absolute value of the upper bound. This measure, expressed as a precentage, provides a relative value, offering insights into the potential further reduction in the objective value of the model after the prescribed one-hour run-time limit.

Here, when $\alpha = 0.75$, the gap is 52% on average for problems of size 100, and it is 76% on average for problems of size 200. When $\alpha = 1.0$, the gap is 100% even for problems of size 20, 50, 100 and 200. An MIP gap of 100% is a clear indication that the solver has not yet been able to find a feasible solution within the defined constraints and within the prescribed one-hour run-time limit.

8. Conclusions

In this paper, we studied the means of balancing efficiency and fairness in one-on-one agent-task assignment in agent coalitions composed of individually rational agents. Here, an agent decides to collaborate with other agents only if it brings an individual benefit that is at least as good as when not collaborating. In this regard, we studied the utilitarian, egalitarian and Nash social welfare, the concepts from economics and philosophy that may be applied in such multi-agent coalitions to tackle this issue. Since quantitative fairness measures for task assignment are scarce and/or missing, we proposed three new fairness measures: egalitarian fairness measure (EFM), relative all-to-all fairness measure (RAF), and overall relative opportunity cost fairness (OROCF) measure. Moreover, to improve

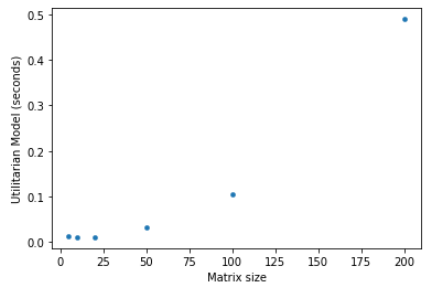

(a) Utilitarian Model

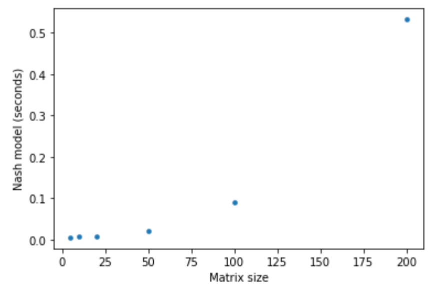

(b) Nash Model

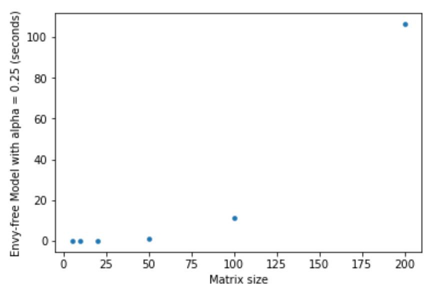

(c) Envy-free Model ($\alpha = 0.25$)

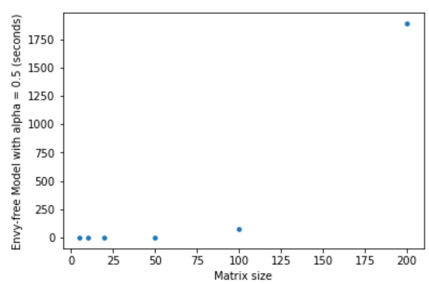

(d) Envy-free Model ($\alpha = 0.5$)

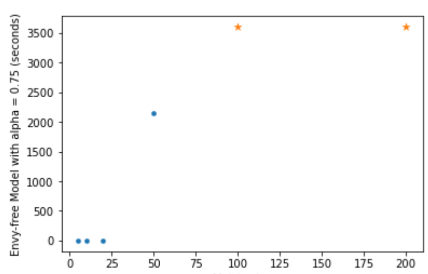

(e) Envy-free Model ($\alpha = 0.75$)

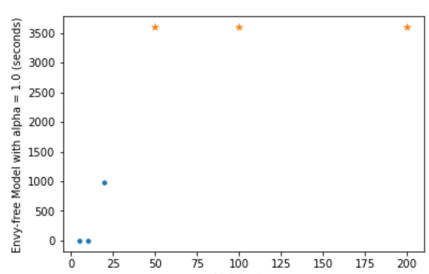

(f) Envy-free Model ($\alpha = 1.0$)

Fig. 1. Run time of the models depending on problem size

the performance of the conventional task assignment model, we proposed the Nash model for the one-on-one task assignment that minimises the product of the costs of each agent considering one-on-one assignment constraints and the envy-free utilitarian model that combines the envy-freeness concept, equity and the utilitarian social welfare measure. The performed simulation experiments show that by using our newly proposed models, we can achieve better fairness in terms of the proposed measures with little sacrifice in the overall efficiency and that the envy-free utilitarian model can be adjusted depending on the need for fairness in a coalition.

The potential impact of the proposed models and fairness measures in real-world applications is substantial. To implement the proposed system effectively, it is imperative to establish an a priori agreement or contract defining the efficiency and fairness measures for the multi-agent coalition.

In the domain of ride-sharing and delivery services, companies like Uber and Lyft can employ these models to optimize efficiency while ensuring fairness among their drivers. This approach can lead to improved driver satisfaction and retention rates. Additionally, in labour markets and the gig economy, these concepts can be harnessed to allocate freelance and short-term work equitably through digital platforms. This benefits both workers and employers by fostering a more engaged and content workforce. Furthermore, our proposed models and fairness measures may be invaluable in supply chain management, disaster response coordination, environmental conservation efforts, healthcare settings, academic research collaborations, and smart grid management. These models excel at promoting computationally efficient task allocation that strikes a balance between efficiency and fairness, potentially resulting in enhanced productivity, coordination, and cooperation across a diverse range of industries and sectors.

In future work, we plan to further study fairness measures, particularly the one encompassing both equality and equity to better support decision-making in collaborative and cooperative open societies where agents can enter and leave the system at any time based on their momentary interest. Moreover, we will focus on the three-index assignment problem where each agent requires a tool to perform a task. The assignment here is also performed in a one-on-one manner. Similarly, crafting a multi-objective model which considers equality, equity and fairness for such a problem is a challenge worth facing henceforth due to the importance of its impact in versatile real-world applications. These include emergency services, agriculture fleet task coordination, delivery services, waste management, construction and infrastructure projects, utility maintenance, and home healthcare, among others.

Acknowledgments. This work has been partially supported by grants PID2021 -123673OB-C32 and TED2021-131295B-C33 funded by MCIN/AEI/10.13039 /501100011033 and by "ERDF A way of making Europe" and the "European Union NextGenerationEU / PRTR", respectively; as well as the AGROBOTIX Project 2022-REGING-95993 funded by the IMPULSO own program of the University Rey Juan Carlos and AGROBOTS Project funded by the Community of Madrid, Spain.

References

1. Airiau, S., Aziz, H., Caragiannis, I., Kruger, J., Lang, J., Peters, D.: Portioning using ordinal preferences: Fairness and efficiency. Artificial Intelligence **314**, 103809 (2023)

2. Alkan, A., Demange, G., Gale, D.: Fair allocation of indivisible goods and criteria of justice. Econometrica: Journal of the Econ. Soc. **59(4)** (1991)
3. Aziz, H., Caragiannis, I., Igarashi, A., Walsh, T.: Fair allocation of indivisible goods and chores. Autonomous Agents and Multi-Agent Systems **36**, 1–21 (2022)
4. Aziz, H., Li, B., Moulin, H., Wu, X.: Algorithmic fair allocation of indivisible items: A survey and new questions. ACM SIGecom Exchanges **20**(1), 24–40 (2022)
5. Billhardt, H., Fernandez, A., Lujak, M., Ossowski, S., Julian, V., De Paz, J.F., Hernandez, J.Z.: Coordinating open fleets. a taxi assignment example. AI Communications **30**(1), 37–52 (2017)
6. Binns, R.: On the apparent conflict between individual and group fairness. In: Proc. of the ACM FAccT 2020. pp. 514–524 (2020)
7. Bonald, T., Massoulié, L., Proutiere, A., Virtamo, J.: A queueing analysis of max-min fairness, proportional fairness and balanced fairness. Queueing Systems **53**(1), 65–84 (2006)
8. Burkard, R., Dell'Amico, M., Martello, S.: Assignment problems: revised reprint. SIAM (2012)
9. Cappelen, A.W., Tungodden, B.: Fairness and the proportionality principle. Social Choice and Welfare **49**(3), 709–719 (2017)
10. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., et al.: The unreasonable fairness of maximum nash welfare. ACM-TEAC **7**(3), 1–32 (2019)
11. Chevaleyre, Y., Dunne, P.E., Endriss, U.e.a.: Issues in multiagent resource allocation. Informatica **30**(1), 3–32 (2006)
12. Conitzer, V., Freeman, R., Shah, N., Vaughan, J.W.: Group fairness for the allocation of indivisible goods. In: Proc. of the AAAI Conference on AI. vol. 33(1), pp. 1853–1860 (2019)
13. Cook, K.S., Hegtvedt, K.A.: Distributive justice, equity, and equality. Annual review of sociology **9**(1), 217–241 (1983)
14. Corbett-Davies, S., Goel, S.: The measure and mismeasure of fairness: A critical review of fair machine learning. arXiv preprint arXiv:1808.00023 (2018)
15. Cousy, K., Lujak, M., Salvatore, A., Fernández, A., Giordani, S.: On balancing fairness and efficiency of task assignment in agent societies. In: Highlights in Practical Applications of Agents, Multi-Agent Systems, and Complex Systems Simulation. The PAAMS Collection. Communications in Computer and Information Science, vol. 1678, pp. 95–107. Springer International Publishing (2022)
16. Fleurbaey, M., Maniquet, F.: A theory of fairness and social welfare, vol. 48. Cambridge U. Press (2011)
17. Ghodsi, A., Zaharia, M., Hindman, B.e.a.: Dominant resource fairness: Fair allocation of multiple resource types. In: 8th USENIX Symposium NSDI 11 (2011)
18. Giordani, S., Lujak, M., Martinelli, F.: A distributed multi-agent production planning and scheduling framework for mobile robots. Computers & Industrial Engineering **64**(1), 19–30 (2013)
19. Guo, H., Li, W., Deng, B.: A survey on fair allocation of chores. Mathematics **11**(16), 3616 (2023)
20. Hellman, D.: Measuring algorithmic fairness. VLR **106**(4), 811–866 (2020)
21. Jagtenberg, C., Mason, A.: Fairness in the ambulance location problem: maximizing the bernoulli-nash social welfare. Available at SSRN 3536707 (2020)
22. Kranich, L.: Resource-envy-free and efficient allocations: A new solution for production economies with dedicated factors. Journal of Math. Econ. **89**, 1–7 (2020)
23. Kraus, S., Plotkin, T.: Algorithms of distributed task allocation for cooperative agents. Theoretical Computer Science **242**(1-2), 1–27 (2000)
24. Kuhn, H.W.: The hungarian method for the assignment problem. Naval research logistics quarterly **2**(1-2), 83–97 (1955)
25. Kurzban, R., Duckworth, A., Kable, J.W., Myers, J.: An opportunity cost model of subjective effort and task performance. Behavioral and brain sciences **36**(6), 661–679 (2013)
26. Lujak, M., Billhardt, H., Ossowski, S.: Distributed coordination of emergency medical service for angioplasty patients. Annals of Mathematics and Artificial Intelligence **78**, 73–100 (2016)

27. Lujak, M., Giordani, S., Omicini, A., Ossowski, S.: Decentralizing coordination in open vehicle fleets for scalable and dynamic task allocation. Complexity **2020**(1047369), Article ID 1047369 (2020)
28. Lujak, M., Sklar, E., Semet, F.: On multi-agent coordination of agri-robot fleets. In: CEUR Workshop Proc. of The 11th International Workshop on Agents in Traffic and Transportation (ATT 2020) held in conjunction with the 24th European Conference on Artificial Intelligence (ECAI 2020). vol. 2701(12) (2020)
29. Moulin, H.: Fair division and collective welfare. MIT press (2004)
30. Moulin, H.: Axioms of cooperative decision making. No. 15 in Econometric Society Monographs, Cambridge U. P. (1991)
31. Mulgan, T.: Understanding utilitarianism. Routledge (2014)
32. Nanda, V., Xu, P.e.a.: Balancing the tradeoff between profit and fairness in rideshare platforms during high-demand hours. In: Proc. of the AAAI Conference on Artificial Intelligence. vol. 34(2), pp. 2210–2217 (2020)
33. Nelson, L.M.: Collaborative problem solving. Instructional design theories and models: A new paradigm of instructional theory **2**(1999), 241–267 (1999)
34. Padala, M., Gujar, S.: FNNC: achieving fairness through neural networks. In: Proc. of the Twenty-Ninth Int. Joint Conf. on Artif. Int. IJCAI-20 (2020)
35. Ramezani, S., Endriss, U.: Nash social welfare in multiagent resource allocation. In: Agent-mediated electronic commerce. Designing trading strategies and mechanisms for electronic markets, pp. 117–131. Springer (2009)
36. Rohbogner, G., Hahnel, U.J., Benoit, P., Fey, S.: Multi-agent systems' asset for smart grid applications. Computer Science and Information Systems **10**(4), 1799–1822 (2013)
37. Sampat, A.M., Zavala, V.M.: Fairness measures for decision-making and conflict resolution. Optimization and Engineering **20**(4), 1249–1272 (2019)
38. Schappe, S.P.: Understanding employee job satisfaction: The importance of procedural and distributive justice. J. of Business and Psyc. **12**(4), 493–503 (1998)
39. Sen, A.: Collective choice and social welfare. In: Collective Choice and Social Welfare. Harvard University Press (2017)
40. Shrestha, Y.R., Yang, Y.: Fairness in algorithmic decision-making: Applications in multi-winner voting, machine learning, and recommender systems. Algorithms **12**(9), 199 (2019)
41. Suksompong, W.: Constraints in fair division. ACM SIGecom Exchanges **19**(2), 46–61 (2021)
42. Zhao, Y., Zheng, K., Guo, J., Yang, B., Pedersen, T.B., Jensen, C.S.: Fairness-aware task assignment in spatial crowdsourcing: Game-theoretic approaches. In: 2021 IEEE 37th ICDE. pp. 265–276. IEEE (2021)

Marin Lujak is a Distinguished Researcher Lecturer at the University Rey Juan Carlos. His research focuses on AI and distributed optimization and coordination of multi-agent and cyber-physical systems. In more than 70 peer-reviewed publications, he addresses societal challenges like smart transport, emergency management, and multi-robot coordination.

Alessio Salvatore is a Postdoctoral Researcher at Institute for Systems Analysis and Computer Science - Italian National Research Council. His main research interests are Mixed Integer Linear Programming and the development of heuristic algorithms with applications in transportation, logistics and scheduling.

Alberto Fernández, a Full Professor at the University Rey Juan Carlos, has over 20 years of research experience in the field of Artificial Intelligence. With more than 120 peer-

reviewed publications, his research activities are centered particularly in areas such as multi-agent systems and knowledge representation.

Stefano Giordani, a Full Professor of Operations Research at the University of Rome Tor Vergata, is renowned for his extensive research contributions in combinatorial optimization and its applications in scheduling, transportation, and logistics. With over 100 peer-reviewed publications, his work has significantly advanced the field.

Kendal Cousy is a PhD student in multi-agent coordination for agricultural robotics. He pursued research through a cotutelle program between the University Rey Juan Carlos and the University of Rome Tor Vergata.

Received: January 19, 2023; Accepted: October 10, 2023.