

# Multi-Objective Fairness in Team Assembly

Rodrigo Borges, Otto Sahlgren, Sami Koivunen, Kostas Stefanidis, Thomas Olsson, and Arto Laitinen

Tampere University, Finland

{rodrigo.borges,otto.sahlgren,sami.koivunen}@tuni.fi  
{konstantinos.stefanidis,thomas.olsson,arto.laitinen}@tuni.fi

**Abstract.** Team assembly is a problem that demands trade-offs between multiple fairness criteria and computational optimization. We focus on four criteria: (i) fair distribution of workloads within the team, (ii) fair distribution of skills and expertise regarding project requirements, (iii) fair distribution of protected classes in the team, and (iv) fair distribution of the team cost among protected classes. For this problem, we propose a two-stage algorithmic solution. First, a multi-objective optimization procedure is executed and the Pareto candidates that satisfy the project requirements are selected. Second,  $N$  random groups are formed containing combinations of these candidates, and a second round of multi-objective optimization is executed, but this time for selecting the groups that optimize the team-assembly criteria.

## 1 Introduction

Given a set of optimization criteria and constraints, team assembly algorithms are designed to select, from a pool of candidates who each have a set of skills, a set of individuals that jointly fulfils the requirements of a predefined project. Decision-makers have to establish a clear understanding of project requirements and teams' envisioned tasks to translate them into computationally tractable formal requirements, and also choose an appropriate way of assigning candidates into teams [9]. Team assembly also requires attention to ethical and legal questions, especially when conducted in high-stakes domains, such as formal education or professional work contexts. Unfair bias, whether social or technical in origin [6], is a particularly salient concern in team assembly since it can lead to unfair treatment of candidates in the selection process and disadvantage members of protected groups (such as gender or ethnic groups).

Existing work has developed methods for improving team assembly algorithms in different respects, such as reducing the cost of team assembly [2], distributing workloads more equitably among candidates [1] and improving the representation of different groups in the resulting teams [2]. Whereas a large body of work is devoted to developing methods for identifying and mitigating wrongful bias and unfairness in rankings and recommendations [15], research addressing these issues in computational team assembly contexts remains scarce. Most existing approaches are designed for incremental solutions where teams

are formed by selecting one candidate after the other, and a single distributive desideratum is optimized. Our work addresses two problems. First, decision-makers often have multiple objectives that need to be balanced simultaneously [12]. Second, incremental approaches to team-assembly can be undesirable in certain cases, such as when choosing one candidate at time  $t_1$  closes off the possibility to choose another more suitable candidate at a later time  $t_2$ .

To address these issues, we formulate team assembly as a multi-objective optimization task where fairness-aware team assembly involves considering several competing objectives. We describe our framework and illustrate its benefits by employing four criteria for fairness-aware team-assembly. Ideally, a team assembly algorithm would compare every possible team-composition in light of the criteria and choose the one that minimizes a target objective. However, this approach can be expensive especially when the candidate pool is large. To address this issue, we propose a two-step team assembly procedure: First, a multi-objective optimization procedure is executed and the Pareto candidates that satisfy the project requirements are selected. Second,  $N$  random groups are formed containing combinations of these candidates, and a second round of multi-objective optimization is executed, but this time for selecting the groups that optimize the team assembly criteria. The choice between teams is determined according to a combination of all fairness criteria. This algorithm is not as cheap as selecting the best candidates incrementally, and not as expensive as testing all possible groups that can be formed. Instead, the proposed algorithm filters the best candidates among the ones that fulfill the project requirements, and forms several random groups containing these candidates. When selecting a group that is already formed one can directly access the fairness metrics, and it is easier for the algorithm to minimize a given criterion or a set of criteria.

## 2 Related Work

A team can be defined as a set of two or more individuals interacting in a dynamic and interdependent manner towards a common goal [7]. Depending on the context, teams may be required to complete tasks ranging from group-learning (e.g., writing group essays) to product development and innovation processes (e.g., generating ideas [10]). Various techniques and computational solutions are available for team assembly, ranging from software systems that enable users to control the selection process to systems that optimize the process using machine learning (see [7], [9]).

*Team assembly.* Research on computational team-assembly spans diverse application areas and different computational approaches. For example, [11] proposes a team recommender that groups individuals within a social network based on pre-defined skill requirements. Given a task, a pool of individuals with different skills, and a social network which captures whether individuals are compatible with one another, the goal is to define a set of users to perform the task, where the users both meet the skill requirements of the task and work effectively together as a team, using a communication cost indicator to measure effective-

ness. [14] presents an approach to form and recommend emergent teams based on how software artifacts are changed by developers. In [17], teams are formed based on the personality of the team members using a classifier to predict the performance of the constructed teams. [16] frames team assembly as a group recommendation task, where a team is formed from users with specific constraints, and then items are recommended to that team.

*Bias and fairness in team-assembly.* Research on software fairness has developed various metrics for identifying and addressing unfair bias. For example, Statistical Parity [5] requires that the protected groups (e.g., gender groups) experience equal selection rates, whereas Equalized Odds [8] requires that those groups experience equal error rates. Numerous techniques can be applied to mitigate bias in data, algorithms, or the output [15]. Fairness has been discussed less in computational team-assembly contexts, though there are some notable exceptions. For example, [3] formulates fair team assembly as a fair allocation task: students should be assigned to projects so that there are balanced workloads and tasks in the resulting teams. Another example is [2], which examines fair team-formation in an online labour marketplace. Here, each user possesses a set of skills and all users are divided into disjoint groups subject to a single fairness constraint: the formed team should have the same number of users belonging to each protected class, and the team should have all the skills needed to complete a given task. Lastly, [13] presents the most similar setting to our work by exploring a problem where teams have multidisciplinary requirements and the incremental selection of team members is based on the match of their skills and the requirements.

### 3 Motivation

We propose a multi-objective approach to fairness-aware team-assembly in a subset selection setting where a single team is formed with a one-shot (non-incremental) process given a pool of candidates and a set of project requirements. This approach addresses two limitations or gaps in existing work: the focus on incremental team-assembly tasks and the use of a single fairness objective.

*Team Assembly as One-Shot Subset Selection.* We address a gap in the research literature by considering a one-shot subset selection task where a team is formed by choosing an optimal set of individuals from a larger set of candidates, where project-to-team fit is evaluated by considering project requirements and candidates' skills. This is because (1) subset selection has received comparably less attention in research on software fairness (see, however, [4]) and (2) existing work on fairness-aware team assembly has largely focused on an incremental approach to selecting candidates, which can be undesirable or suboptimal in certain cases since the overall composition of the team can be known only after selecting all candidates. For example, to reduce the overall cost of the team, the incremental approach would intuitively start with selecting the cheapest candidate that satisfies one of the project requirements even though another candidate with a slightly higher cost would satisfy two other criteria with a higher score.

**Table 1.** Example of a team T formed according to project requirements P. The values in the table indicate the cost associated with each user/skill combination.

Members	Class	Req 1	Req 2	Req 3	Req 4
Member 1	0	0.000	0.035	0.000	0.000
Member 2	1	0.100	0.000	0.000	0.022
Member 3	1	0.000	0.000	0.090	0.081

*Multi-Objective Fairness in Team Assembly.* Using a single fairness objective (or constraint) can result in narrow evaluations of the team-compositions and preclude identification of trade-offs between objectives [12]. Our multi-objective approach to fair team assembly recognizes that team assembly procedures often distribute multiple distinct goods and opportunities simultaneously. E.g., fair team-assembly involves fair distribution of access to teams, as well as fair allocation of tasks and responsibilities between team-members, taking into account also the overall utility of the outcomes of the procedure. We specify four objectives for fair team assembly: (a) *Fair Representation*: The distribution of protected attributes within a team should be as equitable as possible. (b) *Fair Workload Distribution*: The distribution of tasks within a team should be as equal as possible. (c) *Fair Expertise Distribution*: The distribution of skills within a team should be as balanced as possible. (d) *Fair Cost Distribution*: The distribution of the total cost within a team should be as fair as possible considering candidates’ protected attributes. In Table 1, Member 1 presents a substantially lower cost compared to others, and Req 2 sums a substantially lower cost if compared to other requirements. The members associated with class 1 concentrate a higher cost, compared to the single member associated with class 0.

## 4 Problem Formulation

Let  $S = \{s_1, \dots, s_m\}$  be a set of skills,  $A$  be a binary sensitive attribute that can assume values  $A_0$  or  $A_1$ ,  $U = \{u_1, \dots, u_k\}$  be a set of individuals, i.e., the candidate pool, and  $P$  be a set of requirements for a project, i.e., a subset of the skill set ( $P \subset S$ ). An individual  $u \in U$  is represented as a combination of a cost profile ( $u^S$ ) containing the hiring cost associated with their skills, and a value ( $u^A$ ) associated with a sensitive attribute. The cost profile is obtained through a function  $\theta$  that returns the cost of a certain skill, for example,  $u^S = (\theta(s_1, u), \theta(s_2, u), \dots, \theta(s_m, u))$  represents user  $u$  according to their cost in skills  $\{s_1, s_2, \dots, s_m\}$ . We assume a user has a certain skill as long as the cost associated with that skill is greater than 0.

Any set of more than 2 and less than  $|U|$  individuals is considered a team  $T$ . And the number of project requirements that are fulfilled by a team is referred to as *coverage*. The aim of the team-assembly method is to select among all teams that cover the project requirements, the team that minimizes five objectives: team cost, workload uneven distribution, expertise uneven distribution, representation parity, and cost difference. These objectives are described next.

**Team Cost.** The total cost of hiring a team for a project is defined as:

$$Cost(T, P) = \sum_{u \in T} \sum_{j=1}^{|S \cap P|} \theta(s_j, u). \quad (1)$$

It can be described as the summation of the cost associated with each team member's skills that match the project requirements. Notably, one candidate can contribute to more than one task in the project when their skills coincide with more than one project requirement.

**Workload Uneven Distribution.** It is calculated as the standard deviation of the cost associated with each member of the team:

$$Workload(T, P) = \sqrt{\frac{1}{|T|} \sum_{u \in T} \left( \sum_{j=1}^{|S \cap P|} \theta(s_j, u) - \frac{Cost(T, P)}{|T|} \right)^2}. \quad (2)$$

A team in which the total cost is well distributed among members (low variance) is considered fair, and a team in which the cost is concentrated among a few members (high variance) is considered unfair.

**Expertise Uneven Distribution.** In addition to distributing the costs fairly among candidates, the costs should also be fairly distributed among project requirements. The unevenness of expertise distribution is calculated as:

$$Expertise(T, P) = \sqrt{\frac{1}{|P|} \sum_{j=1}^{|S \cap P|} \left( \sum_{u \in T} \theta(s_j, u) - \frac{Cost(T, P)}{|P|} \right)^2}. \quad (3)$$

Similar to workload distribution, this objective measures the standard deviation of the cost associated with each project requirement. A fair team is expected to distribute their cost among requirements as even as possible, thus resulting in a low standard deviation value.

**Representation Parity.** This measures the difference between the occurrences of  $A_0$  and  $A_1$  within a team as potential values for a sensitive attribute  $A$ . The objective is calculated as:

$$Representation(T, A) = \frac{\sqrt{(|f(T, A_0)| - |f(T, A_1)|)^2}}{|T|}, \quad (4)$$

where function  $f(T, A_0)$  returns a set containing the members of  $T$  associated with sensitive attributes  $A_0$ , as well as in the case of attribute  $A_1$ . A low Representation Parity indicates a fair distribution of attribute  $A$  whereas a high value indicates a majority of members associated with one of the classes,  $A_0$  or  $A_1$ .

**Cost Difference.** This measures the difference between the cost allocated to two categories,  $A_0$  and  $A_1$ , within a team. The total cost of team members associated with a certain sensitive attribute, named Cost Attribute (CA), is calculated as:

$$CA(T, P, A_0) = \sum_{u \in f(T, A_0)} \sum_{j=1}^{|S \cap P|} \theta(s_j, u), \quad (5)$$

in the case of attribute  $A_0$ . And the Cost Difference objective is calculated as:

$$CostDiff(T, A, P) = \frac{\sqrt{(CA(T, P, A_0) - CA(T, P, A_1))^2}}{Cost(T, P)}. \quad (6)$$

As mentioned before, our goal is to select a team  $T$  that fulfils the project  $P$  requirements and that minimizes the multi-objective condition:

$$\begin{aligned} argmin_T (Cost(T, P), Workload(T, P), Expertise(T, P), \\ Representation(T, A), CostDiff(T, A)). \end{aligned} \quad (7)$$

#### 4.1 Multi-Objective Fairness in Team Assembly

We propose a method designed for assembling teams with multiple fairness constraints. The method assumes a pool of candidates from which the team will be selected, a project and a sensitive attribute associated with each of the candidates. It formulates fairness-aware team-assembly as a multi-objective optimization problem performed in two stages. First, project requirements are considered as objectives, and the best candidates are selected for the next phase. Second, multiple teams are formed with these candidates and fairness constraints are calculated for each of the teams. This time the fairness constraints are assumed as objectives, and the team that minimizes these constraints while fulfilling the project requirements is considered the fairer one (Algorithm 1).

Given a candidate pool ( $U$ ) and set of project requirements ( $P$ ), the first action is to filter candidates with at least one skill required by the project. (*FilterCandidatesWithProjectSkills()* in Algorithm 1). The filtering process removes all users for which  $|u \cap P| = 0$ , and the remaining candidates are referred to as  $U_p$ . The candidates in  $U_p$  are then submitted to a multi-objective optimization step with the aim of selecting the best candidates for this specific project according to the Pareto dominance concept. According to this concept, a candidate dominates another if they perform better in at least one of the project requirements. A candidate is considered non-dominated if they are not dominated by any other candidate in the population, and the set of all non-dominated candidates compose the *Pareto candidates* subset.

At this point, the *paretoCandidates* subset contains the non-dominated candidates considered as the most suitable for the given project, but our notion of a fair team can only be assessed when having a formed team. The next step is to form a reasonable amount of teams containing a fixed number of Pareto candidates selected randomly. The number of random teams ( $N$ ) as well as the size of these teams ( $M$ ) are provided as parameters for the method. Once the teams are formed, it is possible to calculate their coverage, as well as their fairness objectives with Equations 1, 2, 3, 4 and 6. The teams that fulfil the project requirements are filtered out and considered in the following step. *Cost*, *Workload*, *Expertise*, *Representation* and *CostDifference* are then calculated, and each team end up being represented according to the values calculated for its objectives. A second round of multi-objective optimization is executed. This time,

---

**Algorithm 1:** Algorithm (in pseudo-Python) for multi-objective fairness-aware team assembly

---

**Data:** A pool of candidates ( $U$ ), project requirements ( $P$ ), team size ( $M$ ), number of random teams ( $N$ ), sensitive attribute ( $A$ )

**Result:** the selected team ( $T$ )

```
 $U_p = \text{FilterCandidatesWithProjectSkills}(U, P);$   
 $\text{paretoCandidates} = \text{MultiObjective}(U_p, \text{min});$   
  
 $\text{teams} = \text{FormNRandomTeams}(\text{paretoCandidates}, N, M)$   
  
 $\text{candTeams} = [];$   
for  $\text{team}$  in  $\text{teams}$  do  
    if  $\text{CalcTeamCoverage}(\text{team}) == |P|$  then  
         $\text{cost} = \text{CalcTeamCost}(\text{team}, P);$   
         $\text{workload} = \text{CalcTeamWorkload}(\text{team}, P);$   
         $\text{expertise} = \text{CalcTeamExpertise}(\text{team}, P);$   
         $\text{representation} = \text{CalcTeamRepresentation}(\text{team}, A);$   
         $\text{costDifference} = \text{CalcTeamCostDifference}(\text{team}, A);$   
  
         $\text{candTeams.append}([\text{team}, \text{cost}, \text{workload}, \text{expertise}, \text{representation},$   
                           $\text{costDifference}])$   
    end  
end  
  
 $\text{paretoTeams} = \text{MultiObjective}(\text{candTeams}, \text{min});$   
 $T = \text{argmax}(\text{paretoTeams})$ 
```

---

teams are being compared instead of candidates, with the same understanding as in from the first case. Given two arbitrary teams, two are the possibilities: (i) one team dominates the other if it has at least one objective with a lower value than the other team, or (ii) the two teams do not dominate each other. The non-dominated teams are referred to as *Pareto teams*. Finally, all objectives measured for Pareto teams are summed up as an indicator of an *overall unfairness*, and the method selects the team  $T$  with the lower unfairness value.

## 5 Experiments

**Dataset.** We evaluated the proposed method in a dataset obtained from the *freelancer*<sup>1</sup> website, in which candidates register themselves to be hired as freelance workers. The dataset was provided by the authors of [2], and it contains 1,211 candidates who self-declared their costs and their expertise in 175 skills. In the information available in the dataset, users are associated with skills in a

<sup>1</sup> <https://www.freelancer.com/>

binary fashion (having or not having expertise), but no information is provided about the cost of each skill separately. We decided that the cost declared by the users is the same for every skill in which they have the expertise, meaning that if user  $u$  declared a cost  $c$  and they are hired for a project in which they will contribute with two skills, then the total cost associated with this users is  $2 \times c$ .

The original dataset associates each user with a cost value, and it also associates the same user with a set of skills. We decided that the price paid for hiring one user is the result of multiplying their cost and the number of skills matching the project requirement. That is to say, if user  $u$  costs  $c$  and they are hired to work on a project whose requirements match with one of their skills, then they will be paid  $c$ , but the project requirement matches with two of their skills, then they will be paid  $2 \times c$ .

In [2], the authors attributes a hypothetical binary sensitive attribute to each candidate, and generates several versions of the same dataset, associating candidates with this attribute in different proportions. For example, if we consider this sensitive attribute as gender, we could imagine a candidate poll with 70% men, and 30% of women. The idea is to simulate biased candidate pools and evaluate the impact of those in the teams formed by team-assembly methods. We decided to use the dataset in which members are equally represented in the candidate pool (50/50), and the dataset in which members are more unevenly represented (10/90). The dataset contains also the requirements for 600 projects. No information to order the projects in any way is provided, and so we consider projects independently.

**Previous Approaches.** We applied two other team-assembly methods to the same task for the sake of comparison. The first method, named *Incremental*, was inspired in [13] and selects the most suitable candidates incrementally until the project requirements are fulfilled. The second method, named *Fair Allocation*, was inspired in [2] and operates in a similar fashion, but this time considering users associated with a sensitive attribute and forcing as much as possible that the formed team has a fair distribution of this attribute among its members.

## 6 Results

Multi-Objective, Incremental and Fair Allocation were evaluated for forming a team for each of the 600 projects in the freelancer dataset, and the average value obtained for each of the objectives presented in Section 4 was calculated. The results are reported in Table 2, separately for the two datasets containing different proportions of the sensitive attributes, 50/50 and 10/90. The results of Multi-Objective are presented according to five different optimization objectives, named configurations: Top-Cost, Top-Workload, Top-Expertise, Top-Representation and Top-Cost Difference, along with a Random selection variation. On average, the first round of multi-objective optimization reduced the number of candidates by 77%, meaning that the candidates dominated others with compatible skills representing 23% of the total. In the second round of multi-objective optimization, the teams were reduced by 98% on average. The

**Table 2.** Cost, Workload, Expertise, Representation, Cost Difference, and number of formed teams for Incremental, Fair Allocation and Multi-Objective methods. Multi-Objective can optimize different criteria, and its results are presented according to seven objectives: Random, Top-Cost, Top-Workload, Top-Expertise, Top-Representation, Top-Cost Difference and Top-Sum. The best results for each objective are in bold-face, and the second-best results are in underlining.

Classes Dist.	Algorithm	Cost	Workload	Expertise	Representation	Cost Difference	Teams
50/50	Incremental	<b>23.311 (15.548)</b>	0.035 (0.034)	0.025 (0.026)	0.480 (0.378)	0.610 (0.345)	486
	Fair Allocation	29.201 (20.453)	0.042 (0.040)	0.032 (0.032)	0.238 (0.268)	0.447 (0.286)	486
	Random	40.862 (23.391)	0.045 (0.045)	0.035 (0.036)	0.343 (0.339)	0.419 (0.351)	506
	Multi-Obj. Top-Cost	26.099 (16.876)	0.035 (0.038)	0.026 (0.027)	0.450 (0.360)	0.595 (0.333)	506
	Multi-Obj. Top-Workload	42.294 (27.074)	<b>0.018</b> (0.029)	0.032 (0.036)	0.434 (0.367)	0.471 (0.357)	506
	Multi-Obj. Top-Expertise	37.943 (23.646)	0.039 (0.039)	<b>0.011 (0.020)</b>	0.445 (0.363)	0.525 (0.359)	506
	Multi-Obj. Top-Repres.	28.503 (17.531)	0.036 (0.038)	0.028 (0.027)	<b>0.143 (0.170)</b>	0.393 (0.251)	505
	Multi-Obj. Top-Cost Diff.	41.458 (25.980)	0.040 (0.040)	0.037 (0.036)	0.188 (0.197)	<b>0.091 (0.196)</b>	506
	Multi-Obj. Top-Sum.	39.147 (22.969)	0.032 (0.035)	0.028 (0.029)	0.144 (0.174)	0.107 (0.206)	506
10/90	Incremental	<b>23.423 (15.693)</b>	0.035 (0.034)	0.025 (0.026)	0.758 (0.341)	0.875 (0.239)	486
	Fair Allocation	31.212 (20.802)	0.046 (0.040)	0.034 (0.030)	<b>0.413 (0.370)</b>	0.600 (0.328)	484
	Random	40.939 (25.097)	0.044 (0.047)	0.035 (0.038)	0.722 (0.354)	0.748 (0.350)	506
	Multi-Obj. Top-Cost	26.159 (17.047)	0.035 (0.038)	0.026 (0.027)	0.729 (0.347)	0.843 (0.253)	506
	Multi-Obj. Top-Workload	42.525 (27.463)	<b>0.018 (0.029)</b>	0.033 (0.036)	0.828 (0.297)	0.836 (0.285)	505
	Multi-Obj. Top-Expertise	38.140 (23.855)	0.039 (0.040)	<b>0.011 (0.020)</b>	0.795 (0.327)	0.842 (0.298)	506
	Multi-Obj. Top-Repres.	30.231 (19.608)	0.039 (0.041)	0.030 (0.032)	0.435 (0.390)	0.594 (0.353)	506
	Multi-Obj. Top-Cost Diff.	34.961 (22.256)	0.040 (0.039)	0.032 (0.031)	0.470 (0.377)	<b>0.459 (0.439)</b>	506
	Multi-Obj. Top-Sum.	37.661 (23.700)	0.035 (0.038)	0.027 (0.030)	0.432 (0.389)	0.460 (0.436)	505

teams that dominate the others represent only 2%. This reflects how much the teams formed randomly can be internally equivalent or redundant.

In general, the Incremental method was the most efficient in minimizing the total cost and the size of the formed teams. The Multi-Objective method was able to assemble a slightly higher number of teams than other methods, probably because of forming teams in one-shot instead of incrementally. When configured to minimize a specific objective, the Multi-Objective method was efficient in selecting the teams, except when the objective was the Cost, and when the objective was the Representation and the dataset contained an uneven distribution of classes (10/90). In the former case, Incremental was the most efficient method, and in the latter case, the Fair Allocation performed better.

In the context of Multi-Objective fairness, it is preferable that a team presents a good balance of objectives instead of an extremely low value for one objective despite the others. E.g., when configured to optimize the Expertise objective (Top-Expertise) in the uneven dataset (10/90), Multi-Objective selected, on average, teams with a fairly high Representation (0.795) and Cost Difference (0.842) values. The Top-Sum configuration, on the other hand, selected teams with the second-best average values (highlighted with underline in Table 2) for three out of five objectives, for both datasets. The two objectives in which the configuration did not perform well were Cost and Expertise, which leads us to the next step of analyzing potential conflicts between objectives in the results.

## 7 Conclusions

In this paper, we argued in favour of a wider notion of fairness in the context of team assembly by framing the task of forming teams as a multi-objective

optimization procedure. We have also proposed an algorithm for assembling teams with multiple fairness constraints that assembled teams in a one-shot fashion, as opposed to incremental methods proposed previously in the literature. Our method is flexible enough that it can be applied to situations when one single objective needs to be minimized (or maximized), as well as in situations when all objectives need to be optimized jointly.

## References

1. Anagnostopoulos, A., Becchetti, L., Castillo, C., Gionis, A., Leonardi, S.: Online team formation in social networks. In: WWW (2012)
2. Barnabò, G., Fazzino, A., Leonardi, S., Schwegelshohn, C.: Algorithms for fair team formation in online labour marketplaces. In: Companion of WWW (2019)
3. Bulmer, J., Fritter, M., Gao, Y., Hui, B.: FASTT: team formation using fair division. In: Canadian AI (2020)
4. Cachel, K., Rundensteiner, E.: Fins auditing framework: Group fairness for subset selections. In: AAAI/ACM Conference on AI, Ethics, and Society (2022)
5. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.: Fairness through awareness. In: Innovations in Theoretical Computer Science Conference (2012)
6. Friedler, S.A., Scheidegger, C., Venkatasubramanian, S.: The (im)possibility of fairness: Different value systems require different mechanisms for fair decision making. *Commun. ACM* **64**(4), 136–143 (2021)
7. Gómez-Zarà, D., DeChurch, L.A., Contractor, N.S.: A taxonomy of team-assembly systems: Understanding how people use technologies to form teams. *Proc. ACM Hum. Comput. Interact.* **4**(CSCW2), 181:1–181:36 (2020)
8. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. In: NIPS (2016)
9. Harris, A.M., Gómez-Zarà, D., DeChurch, L.A., Contractor, N.S.: Joining together online: The trajectory of CSCW scholarship on group formation. *Proc. ACM Hum. Comput. Interact.* **3**(CSCW), 148:1–148:27 (2019)
10. Koivunen, S., Olshannikova, E., Olsson, T.: Understanding matchmakers’ experiences, principles and practices of assembling innovation teams. *Comput. Support. Cooperative Work.* **30**(4), 589–616 (2021)
11. Lappas, T., Liu, K., Terzi, E.: Finding a team of experts in social networks. In: SIGKDD (2009)
12. Lee, M.S.A., Floridi, L.: Algorithmic fairness in mortgage lending: From absolute conditions to relational trade-offs. *Minds and Machines* **31**(1), 165–191 (2021)
13. Machado, L., Stefanidis, K.: Fair team recommendations for multidisciplinary projects. In: WI (2019)
14. Minto, S., Murphy, G.C.: Recommending emergent teams. In: MSR (2007)
15. Pitoura, E., Stefanidis, K., Koutrika, G.: Fairness in rankings and recommendations: An overview. *The VLDB Journal* (2021)
16. Stefanidis, K., Pitoura, E.: Finding the right set of users: Generalized constraints for group recommendations. In: PersDB (2012)
17. Yilmaz, M., Al-Taei, A., O’Connor, R.V.: A machine-based personality oriented team recommender for software development organizations. In: EuroSPI (2015)