

On Measuring the Diversity of Organizational Networks

Zeinab S. Jalali¹, Krishnaram Kenthapadi², and Sucheta Soundarajan¹

¹ Syracuse University, Syracuse, NY, USA {zsaghati, susounda}@syr.edu

² Amazon AWS AI, Sunnyvale, CA, USA kenthk@amazon.com

Abstract. The interaction patterns of employees in social and professional networks play an important role in the success of employees and organizations as a whole. However, in many fields there is a severe under-representation of minority groups; moreover, minority individuals may be segregated from the rest of the network or isolated from one another. While the problem of increasing the representation of minority groups in various fields has been well-studied, diversification in terms of numbers alone may not be sufficient: social relationships should also be considered. In this work, we consider the problem of assigning a set of employment candidates to positions in a social network so that diversity and overall fitness are maximized, and propose Fair Employee Assignment (FairEA), a novel algorithm for finding such a matching. The output from FairEA can be used as a benchmark by organizations wishing to evaluate their hiring and assignment practices. On real and synthetic networks, we demonstrate that FairEA does well at finding high-fitness, high-diversity matchings.

1 Introduction

In order for commercial and non-profit organizations to succeed, it is important for those organizations to recruit a workforce that is not only skilled, but also diverse, as diversity has been positively associated with performance [15]. However, diversity cannot be measured only in terms of numbers: it is known that negative effects may happen when a network is structured in a way that resources are not accessible through the social capital accessible to members of a minority group [9]. Social capital consists of bridging resources from outside of an individual’s group (inter-group connections) and bonding resources from internal group connections (intra-group connections) [13].

The literature contains a number of metrics for measuring network diversity/segregation, the most prominent being assortativity [17]. However, when dealing with dynamic networks where new nodes are being added, it is useful to know not only what the diversity of a specific network snapshot *is* after those nodes are added, but how good it *could have been*. In other words, if new nodes join a network, what is the best assortativity that one could possibly achieve, given pre-existing structure of the network and restrictions on where the new nodes can join?

Our work is motivated by the example of an organization that is evaluating their hiring and employee assignment practices with respect to the diversity (gender, race, etc.) of the organizational network. When positions are open, some set of candidates apply for those positions. Each candidate has some amount (possibly zero) of suitability for each of the open positions. If one’s goal is to minimize segregation in the network while ensuring that each position is filled by a candidate who is suitable for that position, which candidate should one hire for each position? If there are significant gender

disparities in applications across job categories (e.g., if software engineer candidates are disproportionately male), then it may not be possible to achieve perfect diversity in hiring and assignment; but nonetheless, it is useful to know how well one can do. One can imagine similar examples for, say, new graduate students joining an existing scientific collaboration network.

There has been a great deal of recent interest in fairness of hiring/assignment procedures (e.g., the Rooney Rule used by the American National Football League [5]). This is because one cannot simply eliminate an existing professional network and replace it with a diverse network; and moreover, at the hiring stage, the candidate pool may itself be non-diverse or exhibit correlations between protected attributes and skillsets.

In this work, we present `FairEmployeeAssignment` (`FairEA`), a novel algorithm for determining how to assign a set of attributed node *candidates* to *open positions* in a network such that both the *fitness* of the match between nodes and open positions and the *diversity* of the resulting network are maximized. We experimentally demonstrate that `FairEA` outperforms baseline strategies at achieving these goals.

We note that in the United States and other countries, it is illegal to make employment decisions based on protected attributes like race or religion.³ For this reason, although the output of `FairEA` is a matching of candidates to positions, it is *not* intended to be used directly to make assignments. Rather, the matching should be used as a baseline, and compared with the organization’s actual practices in employee hiring and assignment, to assess the quality of an organization’s assignment practices with respect to the diversity of the organizational network.

2 Related Work

The recent scientific literature contains many studies on modeling bias in human recruiting systems. For example, [20] examines strategies for hiring diverse faculty in universities, [2] shows the different likelihoods of hiring and promotion for candidates from different groups with equal skills, and [16] addresses the tradeoff between performance goals and company diversity. [15] shows a positive relationship between board member racial and gender diversity to performance of nonprofits. Unfortunately, automating recruiting systems will not necessarily solve problems of discrimination in hiring [8]. One solution is to make sure that protected attributes do not influence algorithmic decisions, but [7] shows that gender bias exists even after scrubbing gender indicators from a classifier. While traditional approaches measure diversity of organizations in terms of numbers [18], organizations are social networks, and network factors influence entrepreneurial success, mobility through occupational ladders, and access to employments [19].

3 Problem Formulation

We formulate this problem as a multi-objective problem in which the goal is to assign a set of newly-hired employees/employment candidates (without loss of generality, the ‘candidates’) to open positions so as to maximize (1) the *fitness* of employees to positions and (2) the *diversity* of the organizational network, under the constraint that all

³ <https://www.eeoc.gov/laws/practices/>

Table 1: Notation

Symbol	Definition
$\overline{G}(P, E)$	Unweighted, undirected attributed graph
$O = \{o_1, \dots, o_m\}$	Set of open/unfilled positions
$F = \{f_1, \dots, f_s\}$	Set of filled Positions
$P = \{p_1, \dots, p_n\}$	Set of positions (nodes of network G), ($F \cup O = P$)
$Q = \{q_1, \dots, q_s\}$	Set of current employees (q_1, \dots, q_s fill f_1, \dots, f_s)
$C = \{c_1, \dots, c_t\}$	Set of candidates to fill open positions
$\mathbf{W}_{m \times t}$	Fitness matrix, w_{ij} : fitness of candidate c_j for position o_i
$\mathbf{A}_{n \times n}$	Adjacency matrix of network G
$\mathbf{A}_{1s \times s}, \mathbf{A}_{2s \times m}, \mathbf{A}_{3m \times m}$	Sub-matrices of \mathbf{A} , (F to F , F to O , O to O edges)
$\mathbf{Y}_{n \times k}$	Membership matrix of current/future employees to k classes
$\mathbf{YF}_{s \times k}, \mathbf{YO}_{m \times k}$	Sub-matrices of \mathbf{Y} , Membership of current and future employees
$\mathbf{YC}_{t \times k}$	Membership matrix of candidates to k classes
$\mathbf{X}_{m \times t}$	binary output matrix, $x_{ij} = 1$ if c_j is assigned to o_i
$\mathbf{U}_{k \times k}$	u_{ij} : fraction of $class_i$ to $class_j$ edges

open positions must be filled. We compute diversity as assortativity, which measures the extent to which ‘like connect to like’. Figure ?? shows an overview of the problem.

The input for this problem consists of the following:

(1) An undirected network $G = (P, E)$, representing the professional network of an organization. Nodes represent positions (s filled, m open). Each edge (p_i, p_j) represents either a real or expected professional interaction between the employees who currently fill or will fill positions p_i and p_j (i.e., those employees do interact, or are expected to interact once the positions are filled).

(2) A set of t candidates ($t \geq m$). If $t = m$, this problem represents the case where new employees have already been hired and need to be assigned – e.g., newly hired software engineers are being assigned to teams. If $t > m$, this problem can be viewed as a combination of the hiring and assignment problems.

(3) The fitness of each candidate c_j for each position o_i (how well-qualified c_j is for o_i). We assume that it is possible to match candidates to open positions such that each open position is filled subject to having at least one candidate with greater than zero fitness for each open position.

(4) An attribute of interest, such as gender, that divides employees/candidates into k classes of attributes: $class_s, \dots, class_k$. We assume that this attribute is categorical and each node can be member of just one class (e.g., minority and majority).

The output is a matching of candidates to open positions. We refer to the input and output in the rest of the paper as described in Table 1.

3.1 Objectives

FairEA attempts to solve a multi-objective optimization problem with fitness and diversity-related objectives, described below.

Maximizing Fitness The first goal is to maximize the *fitness* of the assignment, subject to the constraint that all positions are filled, and each candidate fills at most one position. This objective corresponds to the organization’s primary goal of recruiting

employees with the required skill sets [6]. We compute the overall *fitness* of a matching by summing the fitness of each pair of matched open positions and candidates. Higher values indicate a matching of better qualified candidates for open positions. This goal can be formulated as an optimization problem: $\max f_1 = \sum_{ij} w_{ij} \cdot x_{ij}$, such that for each $1 \leq i \leq m$, $\sum_{1 \leq j \leq t} x_{ij} = 1$ and for each $1 \leq j \leq t$, $\sum_{1 \leq i \leq m} x_{ij} \leq 1$.

Maximizing Diversity We compute *diversity* as *assortativity*, which measures the extent to which a minority group is integrated into the larger network based on the number of inter- vs intra-group connections. If we consider matrix \mathbf{U} with u_{ij} to be the fraction of edges in the network that connect a node from class i to class j , the assortativity coefficient is equal to $\frac{\text{Trace}(\mathbf{U}) - \|\mathbf{U}^2\|}{1 - \|\mathbf{U}^2\|}$ where $\|\mathbf{U}^2\|$ means sum over all elements in \mathbf{U}^2 [17]. Positive values show more intra-group connections and negative values show more inter-group connections in network. Our goal is to minimize the absolute value of assortativity, so that groups are neither segregated nor isolated. Let \mathbf{A}_1 , \mathbf{A}_2 , and \mathbf{A}_3 be sub-matrices of adjacency matrix \mathbf{A} , where $\mathbf{A}_{1(s \times s)}$ shows the existence of the edges between nodes in F , $\mathbf{A}_{2(s \times m)}$ shows the existence of the edges from nodes in F to nodes in O , and $\mathbf{A}_{3(m \times m)}$ shows the existence of the edges between nodes in O . Let $\mathbf{Y}\mathbf{O}_{m \times k} = \mathbf{X} \cdot \mathbf{Y}\mathbf{C}$ be the binary matrix of the membership of candidates that will fill the open position to each attribute class. Then we have: $\mathbf{U} = \mathbf{Y}^T \cdot \mathbf{A} \cdot \mathbf{Y} = \mathbf{Y}\mathbf{F}^T \cdot (\mathbf{A}_1 \cdot \mathbf{Y}\mathbf{F} + \mathbf{A}_2 \cdot \mathbf{Y}\mathbf{O}) + \mathbf{Y}\mathbf{O}^T \cdot (\mathbf{A}_2^T \cdot \mathbf{Y}\mathbf{F} + \mathbf{A}_3 \cdot \mathbf{Y}\mathbf{O})$ and this goal can be formulated as an optimization problem as: $\min f_2 = \left| \frac{\text{Trace}(\mathbf{U}) - \|\mathbf{U}^2\|}{1 - \|\mathbf{U}^2\|} \right|$.

Other Constraints: In certain cases, there may be other constraints that one wishes to consider. For example, if a minority group is very small, a low-homophily network would indicate that members of that group are isolated. Organizations may want to avoid such an outcome, because such isolated individuals may be unable to find other members of a group and form a *community of support*, where peers share their experiences and discuss their problems and provide social support for each other [21]. In such cases, appropriate constraints (e.g., the minimum number of members from each protected group per team). We give an example of this in the next section.

3.2 Challenges

First, note that the problem considered here is NP-hard via a reduction from Unweighted Max Cut [11]. Second, the problem of minimizing *diversity* is not convex, and is neither sub-modular nor super-modular. Third, although the problem can be formulated as an integer program, this process is computationally slow. These challenges suggest that even a fast approximation algorithm may not exist. As such, we present a heuristic and demonstrate its strong performance experimentally.

4 Method

We propose FairEmployeeAssignment (FairEA), a method for assigning candidates to positions with the goals of maximizing *fitness* and *diversity*. Assume that we are given input and desire output as described in Section 3. In our initial discussions, we assume that candidates are divided into $k = 2$ classes, and in Section 4.2, we explain how FairEA can be generalized for $k > 2$.

FairEA consists of a sequence of iterations, where each iteration i consists of the following three steps:

1. Select two subsets O_i from O and C_i from C using the *FairEA selection process* described later in this section.
2. Assign C_i to O_i using the *FairEA Matching Process*, as described later in this section.
3. If $|O_i| < |O|$ then increment i and return to (1). Otherwise, terminate.

FairEA Selection Process For each pair of open positions $o_a \in O$ and candidates $c_b \in C$, where $w_{ab} > 0$, consider two scores: the *fitness score*, given by w_{ab} , and the *diversity score*. The diversity score is defined as 1 if $c_b \in class_j$ ($j \in \{1, 2\}$), and the number of positions adjacent to o_a filled by an employee from $class_j$ is less than the number of positions adjacent to o_a filled by an employee from the other class.

Using these two scores, we use the Pareto Optimality technique described in [3] to select subsets from O and C . At each iteration i that the *selection process* is called, the output contains all the open positions and candidates that are present in the pairs appearing in the top i Pareto front sets.

A *Pareto front set* consists of all points that are not dominated by any other point (a point (x_1, y_1) is dominated by (x_2, y_2) if $x_2 > x_1, y_2 > y_1$, or $x_2 \geq x_1, y_2 > y_1$, or $x_2 > x_1, y_2 \geq y_1$). To find the top i Pareto front sets, one finds the first Pareto front set as just described, removes all selected points, and repeats for i iterations.

FairEA Matching Process The FairEA matching process is based on the augmenting path approach from the Hungarian algorithm for weighted bipartite matching [10]:

1. Generate a bipartite graph B from C_i to O_i , where edges represent each pair of qualified candidates $o_a \in O$ and open positions $c_b \in C$ & $w_{ab} > 0$. To set weights:
 - Sort the fitness score and diversity score computed, as described earlier.
 - Set edge weights based on their position in the Pareto front sets levels. If (o_a, c_b) is present in level l , its score is equal to $\frac{1}{l}$.
2. Create a labeling l ($l[j] = 0$ for each node v_j in B), an empty matching M , and an empty bipartite graph B_l .
3. If all elements in O_i are matched and present in M , update matrix X based on matching M and stop. Otherwise, update labeling l and bipartite graph B_l .
 - For each unmatched open position o in graph B , set $l(o)$ to the maximum weight of edges connected to node o in B .
 - For each matched open positions o in M (if any), if o is matched to c , $l(o) = weight(c, o)$.
 - For each candidate c in graph B , set $l(c)=0$.
 - Graph B_l contains edges (o, c) where $o_a \in O_i, c \in C_i, l[o]+l[c] \leq weight(o, c)$.
4. Pick an unmatched open position $o_a \in O_i$. Let $S = \{o_a\}$ and $T = \{\}$ and
5. Let $N(S)$ be the set of neighbors of nodes from S in B_l .
6. If $N(S) = T$, update labels:
 - $\alpha = \min_{o \in S, c \notin T} l(o) + l(c) - w(o, c)$,
 - $l(o) = l(o) - \alpha$, if $o \in S$, $l(c) = l(c) + \alpha$, if $c \in T$.
7. If $N(S) \neq T$, pick $c \in N(S) - T$.

- If c is not matched find an augmenting path from o_a to c . Augment M and update G based on the new assignments. Update weights of edges in B using the same approach described in (1) and return to (2).
- If c is matched to o_b , extend the alternating tree. Add o_b to S and add c to T and return to (4).

In each iteration, the matching is improved and G is updated. For newly assigned positions, the attribute will be the attribute of the matched candidate.

4.1 Handling Constraints

FairEA can handle other diversity-related goals in the form of constraints. As an example, we consider the constraint that no minority individual should be isolated from other minority individuals. To address this (or any) constraint, add a step to FairEA.

Suppose a company has k' disjoint teams and wants to ensure that minorities from each team i are grouped with at least t_i other members of that minority group. In this step, FairEA assigns a set of best qualified candidates from each class j to open positions in the team i that has fewer than the threshold t_i employees from class i . (In practice, this can be accomplished via cluster hiring). A threshold of 0 indicates that avoiding isolation is not necessary. In this step, after each matching, FairEA ensures that all the remaining open positions can be filled (i.e., there is at least one distinct candidate with fitness function greater than zero for each remaining open position). It may sometimes not be possible to reach the threshold for a specific team (e.g., there are not enough open positions in the team). In such cases, so as to not *prevent* assignment of a qualified minority to such a team, the algorithm can perform the assignment, but notify the organization. This information can then be used by the organization: e.g., individuals who lack access to other minorities can be enrolled in a mentor/mentee program.

Specifically, FairEA performs the following: For each team i , denote the number of individuals from $class_j$, $j \in \{1, 2\}$ as $|c_{ji}|$, where $|c_{ji}| < t_i$. Sort all pairs of $\{(o_a, c_b), o_a \in O \text{ and } o_a \in class_j \text{ and } c_b \in C\}$ based on w_{ab} (fitness of c for o) in descending order. Next, iterate over all the elements (o_a, c_b) in the sorted set. If both o_a and c_b are not already matched (i.e., all elements of row a and column b in matrix \mathbf{X} are zero) and there is at least one possible complete matching from remaining candidates to remaining open positions, set $x_{ab} = 1$ and remove matched elements from O and C . Continue the iteration until $t_i - |c_{ji}| = 0$ (sufficient new matchings are established) or there are no elements in the set. In the end, if $|c_{ji}|$ is still less than t_i , return the team t_i for notifying the organization.

4.2 Variations on FairEA

It is easy to modify FairEA for other settings:

Non-binary attributes: If the protected attribute has more than two classes, the diversity score calculation has to be changed. For each pair of qualified candidates $o_a \in O$ ($o_b \in class_j$) and open positions $c_b \in C$, where $w_{ab} > 0$, the diversity score is defined as $\frac{ct-ca}{ct}$, where ct is the total number of filled positions adjacent to c_b and ca is the number of filled positions adjacent to c_b that filled by an employee from $class_j$.

Multiple attributes of interest: If there are multiple attributes of interest (e.g., race and gender), by combining them into one new attribute we can address the problem using FairEA for k classes. For instance, suppose we have k_1 classes for gender and k_2 classes for race, we can generate a new attribute called identity with at most $k_1 * k_2$ classes. Because this may be a large number of combinations, these intersectional classes may be merged as appropriate.

5 Experimental Setup

In our experiments, we first demonstrate that FairEA does well at matching employees to positions (with respect to fitness and diversity). Second, we provide an example of FairEA on a real-world organizational network.⁴

5.1 Datasets

It is difficult to get real data for this problem- in particular, on candidate pools and demographics. Thus, to test FairEA under a wide array of conditions, we use both real and synthetic network topologies and attributes and simulate candidate pools. Statistics of datasets are shown in Table 2. (L, M, H) notations indicate low, medium, and high values of assortativity (segregation) in the network.

Table 2: Dataset statistics.

Name	#nodes	#edges	Assort. Coeff.	Attributes (Maj, Min)
CC(M)	46	552	-0.02	(77%, 23%)
CC(H)	46	552	0.37	(57%, 43%)
RT(M)	77	1341	0.02	(88%, 12%)
RT(H)	77	1341	0.43	(65%, 35%)
Nor(L)	1522	4143	-0.19	(61%, 39%)
Nor(M)	1091	3418	.08	(90%, 10%)
Nor(H)	1421	3855	.29	(64%, 36%)
FO(H)	288	2602	.86	(70%, 30%)
DO(H)	265	921	0.92	(70%, 30%)
SF(L)	1000	4000	-0.30	(69%, 31%)
SF(M)	1000	4000	0.07	(69%, 31%)
SF(H)	1000	4000	0.39	(69%, 31%)

Real Network Topologies First, we use the Norwegian Interlocking Directorate network (Nor), which describes connections among directors of public companies in Norway⁵. This dataset includes the ‘gender’ attribute. We selected two snapshots of this network, the first from February 2003 (**Nor(M)**) and the second from October 2009 (**Nor(L)**). These networks have, respectively, the highest and lowest levels of gender assortativity among all snapshots. We also add synthetic attributes to the August 2011 to obtain a network with high assortativity, denoted as **Nor(H)**.

Next, we consider a set of intra-organizational networks showing interactions in a Consulting Company (CC) and a Research Team (RT),⁶ and consider different attributes to generate high and medium levels of assortativity. **CC(M)** is obtained from the ‘gender’ attribute, and **CC(H)** from the ‘location’ attribute (‘Europe’ vs. ‘USA’). **RT(L)** uses the ‘tenure’ attribute (less than one year vs. more than one year), and **RT(H)** uses the ‘location’ attribute (‘London’ vs. the rest of Europe). (In practice, not all of these attributes are things we care about in the context of network diversity; they were chosen for purposes of demonstrating the algorithm.) Additionally, these two sets of networks contain information about the the organisational level of employees, which we use in our second experiment.

⁴ For replication, we have posted our code and data at <https://github.com/SaraJalali/FairEmployeeAssignment>.

⁵ <http://www.boardsandgender.com/data.php>

⁶ <https://toreopsahl.com/datasets/>

Synthetic Networks To evaluate FairEA on additional network topologies, we construct synthetic datasets. First, **FO** (functional organization) has 6 teams and 12 sub-teams with equal number of nodes in each team, and follows the FedEx organizational chart pattern [14]. **DO** (divisional organization) has 3 divisions and 40 teams with an equal number of nodes in each team, and follows the Department of Energy organizational chart pattern [14]. We add a synthetic binary attribute so that members of each team are from one class. We next generated a set of power-law (SF) networks [12]. We generate three networks with the same topology, and assign attributes to have low **SF(L)**, medium **SF(M)**, and high **SF(H)** assortativity levels.

5.2 Open Positions, Teams, and Candidate Pool

For each network, we run 100 trials. In each trial, we sample 10%, 20%, and 30% of nodes randomly as open positions. To simulate the pool of candidates we consider two cases: (1) the candidate pool consists of the nodes set to open (with the same attributes), and (2) the candidate pool consists of two copies of each node set to open. The first setting corresponds to the case where a ‘batch’ of new employees has been hired, and now the employees need to be assigned to teams without considering the hiring process. The second setting corresponds to the case where we consider both hiring and assignment procedures. Moreover, the way we assign attributes to nodes ensures that changes in homophily are actually due to employee *assignment*, rather than changes in attributes.

5.3 Fitness Functions

The fitness function governs which candidates are suitable for which positions. For the first sets of experiments- the evaluation of FairEA- we consider two fitness functions. In F_1 , candidates are qualified for four randomly selected positions with fitness equal to a random number in $(0, 1)$. In F_2 , candidates are fit for the four open positions closest to the position that the candidate had previously filled with fitness equal to a random number in $(0, 1)$.

5.4 Baseline Methods

We use three baseline methods: (1) **Random**, which randomly assigns qualified candidates to each open position; (2) The weighted **Hungarian** algorithm, where the input is a bipartite graph whose two sides correspond to open positions and candidates. An edge (o_a, c_b) exists if $w_{ab} > 0$, and the weight of this edge is the sum of w_{ab} and the diversity score as described in section 4; and (3) **Optimization**, which uses the IPOPT solver in the GEKKO optimization suite [1] for solving the optimization problem with the two goals of maximizing *fitness* and *diversity*. This is the simplified version of the problem where *fitness* is maximized, as described in section 3.1 and *diversity* is optimized by decreasing the gap between number of neighbors from $class_i$ to number of neighbors from $class_j$ for each newly assigned position.

5.5 Metrics

We report results using the following metrics:

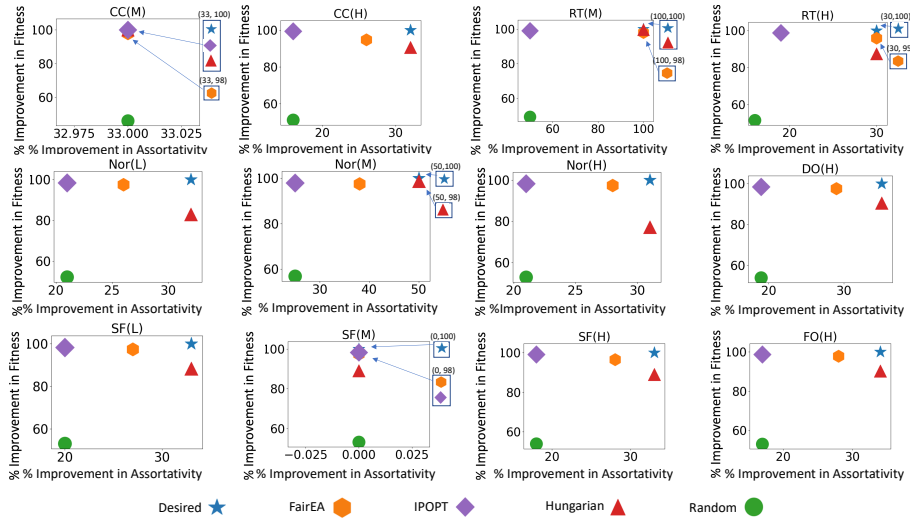


Fig. 1: Comparison results of percentage improvement in fitness and assortativity for FairEA and baseline methods. The ideal solution lies in the top right corner. Hungarian is good at increasing *diversity* and IPOPT is good at increasing *fitness*, but FairEA is good at increasing both. Hungarian performs well when assortativity is low (network is diverse).

- The *overall fit score* is the sum of the fitness scores for each matching. Let FS_h and FS_l be *overall fit score* of the best and worst possible matching in terms of fitness of employees for the open positions respectively and FS_a be the *overall fit score* of the network G after assignment using desired method. Then we define *Percentage Improvement in Fitness* = $\frac{FS - FS_l}{FS_h - FS_l} \cdot 100$.
- The *diversity* of the network is measured by the *assortativity coefficient* [17]. Let AC_b be the *assortativity coefficient* of G' , the subgraph of initial network G consisting only of filled positions, and AC_a be the (*assortativity coefficient* of the network G after assignments are made. Then the *Percentage Improvement in Assortativity* = $\frac{|AC_b| - |AC_a|}{|AC_b|} \cdot 100$.
- The *fraction of minorities* in team i is $FM_i = \frac{\min(|c1_i|, \dots, |ck_i|)}{|c1_i| + \dots + |ck_i|}$ where $|c_j|$ is the number of individuals from *class* j in team i . *Isolation Score* is the average fraction of minorities. *Isolation Score* = $\frac{1}{k} \cdot \sum_{1 \leq i \leq k} FM_i$.

6 Results and Analysis

We first compare FairEA to the baseline algorithms in order to evaluate its performance algorithmically. Next, we use real intra-organizational networks to demonstrate how FairEA might be used in practice.

6.1 FairEA Evaluation

Here, we compare FairEA and baseline algorithms with respect to *diversity* and *fitness*. Figure 1 shows results for *percentage improvement in fitness* and *percentage improvement in assortativity*, where the number of candidates is equal to the number of open




Fig. 2: Results of assortativity coefficient and isolation score over original network (Org) and networks after assignment with different *isolation* thresholds $t_i \in \{0, 2, 0.05 \cdot |k_i|, 0.1 \cdot |k_i|, 0.2 \cdot |k_i|\}$ where $|k_i|$ is size of team $team_i$. Both networks have the potential to become fair by at least 50%.

positions, with fitness function F_1 (candidates are qualified for positions across the network) and 10% open positions. The ideal solution (high fitness and diversity) is in the top right depicted as a star. Results for fitness functions F_2 (candidates are qualified for positions in a specific area) are similar.

In most cases, results for FairEA, IPOPT and Hungarian are in a non-dominated set, with results of FairEA having the lowest crowding distance. More simply, we see that Hungarian does very well with respect to *diversity* (especially when the network was already diverse), IPOPT does very well with respect to increasing *fitness*, and FairEA does well at increasing both.

To summarize results, we compute the *average percentage improvement in fitness* and *average percentage improvement in assortativity* over all datasets with high, medium and low levels of assortativity for each method. When the size of the candidate pool is equal to the number of open positions, FairEA achieves at least 97% of the maximum fitness score while improving the assortativity coefficient value by 39%, 56% and 67% for 10%, 20% and 30% of open positions respectively. (Results were similar for other experimental settings.) Overall, while IPOPT increases *fitness*, it performs poorly on diversity. This demonstrates that simply considering the number of neighbors of a node from each class for newly assigned candidates is not sufficient. Hungarian performs well when the number of open positions is small, but performance decreases as the number of open positions increases. In contrast, FairEA consistently does well.

6.2 Example Usage of FairEA

We next illustrate how FairEA can be used in practice- i.e., to evaluate an organization’s hiring/assignment practices- on the intra-organizational networks CC and RT, which contain position level-related annotations. In such a setting, the organization would identify the set of all positions that have been open in the recent past (whatever timespan is desired), and would use the actual applicants to those positions to form the candidate pool. Because we do not have access to this data, we mark a random $p\%$ of the positions as open and consider the employees that fill those position as candidates. We say that individuals are fit for positions at their level.

Recall that in addition to optimizing for diversity and fitness, FairEA can accommodate constraints related to isolation (ensuring that minority individuals are not too far away from other minority individuals, which can have a negative effect on effectiveness [4]), and such constraints may affect performance with respect to fitness and

diversity. Here, in addition to evaluating fitness and diversity, we also evaluate the effect of such a constraint. We compute the *Percentage Improvement in Fitness*, *Percentage Improvement in Assortativity* and *Percentage Isolation Score* of FairEA’s results when requiring that the number of minority group individuals in each group is at least $\{0, 2, 0.05 \cdot |k_i|, 0.1 \cdot |k_i|, 0.2 \cdot |k_i|\}$ where $|k_i|$ is size of team $team_i$.

We consider networks $\mathbf{CC}(\mathbf{H})$ and $\mathbf{RT}(\mathbf{H})$, both of which are extremely segregated: the original networks have (Assortativity Coefficient & Isolation Score) (0.38&0.07) and (0.43&0.02) respectively. Both of these networks are extremely segregated and have fewer than 10% minorities in each team.

When applying FairEA, 20% and 30% of the positions are open, with different threshold levels for isolation, we see huge improvements in segregation. Figure 2 shows the results of *Assortativity Coefficient* and *Isolation Score*: these large improvements in both assortativity and isolation indicate that both networks have great potential to become more fair.

7 Discussion, Limitations, and Conclusion

In this work, we proposed FairEA, a novel algorithm that can be used to gauge discrimination with respect to a protected attribute. Compared to baselines, FairEA does well at finding high-diversity, high-fitness matchings. While FairEA addresses an abstracted problem, it is a step towards a computational approach to create a diverse workplace in terms of social connections.

This work is intended as a step towards remedying segregation and isolation in organizational networks by providing a simple approach to assessing the quality of hiring/assignment practices. We acknowledge that there are substantially more considerations - such as the process used to generate the candidate pool- that go into evaluating hiring/assignment procedures than those described here, and hope that future work will build on what we have presented.

References

1. Logan DR Beal, Daniel C Hill, R Abraham Martin, and John D Hedengren. Gekko optimization suite. *Processes*, 6(8):106, 2018.
2. David Bjerk. Glass ceilings or sticky floors? Statistical discrimination in a dynamic model of hiring and promotion. *The Economic Journal*, 118(530):961–982, 2008.
3. Yair Censor. Pareto optimality in multiobjective problems. *Applied Mathematics and Optimization*, 4(1):41–59, 1977.
4. Susan G Cohen and Diane E Bailey. What makes teams work: Group effectiveness research from the shop floor to the executive suite. *Journal of management*, 23(3):239–290, 1997.
5. Brian W Collins. Tackling unconscious bias in hiring practices: The plight of the Rooney rule. *NYUL Review*, 82:870, 2007.
6. Michael Craig. Cost effectiveness of retaining top internal talent in contrast to recruiting top talent. In *Competition Forum*, volume 13. American Society for Competitiveness, 2015.
7. Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. Bias in bios: A case study of semantic representation bias in a high-stakes setting. In *Conference on Fairness, Accountability, and Transparency*, 2019.

8. Roel Dobbe, Sarah Dean, Thomas Gilbert, and Nitin Kohli. A broader view on bias in automated decision-making: reflecting on epistemology and dynamics. *arXiv preprint:1807.00553*, 2018.
9. Ka Yi Fung. Network diversity and educational attainment: A case study in China. *The Journal of Chinese Sociology*, 2(1):12, 2015.
10. Alex Grinman. The Hungarian algorithm for weighted bipartite graphs. *Massachusetts Institute of Technology*, 2015.
11. Martin Grötschel and William R Pulleyblank. Weakly bipartite graphs and the max-cut problem. *Operations Research Letters*, 1(1):23–27, 1981.
12. Petter Holme and Beom Jun Kim. Growing scale-free networks with tunable clustering. *Physical Review E*, 65(2), 2002.
13. Nan Lin. A network theory of social capital. *The Handbook of Social Capital*, 50(1):69, 2008.
14. lumenlearning. *Common organizational structures*, (accessed October, 2020).
15. Meg E Cotter Mazzola, Joseph L Pontacolon, Angel Claudio, Javier A Salguero, Marcelles James, and Robert Yawson. Recruiting for success. does boarddiversity matter? 2020.
16. Daniel A Newman and Julie S Lyon. Recruitment efforts to reduce adverse impact: Targeted recruiting for personality, cognitive ability, and diversity. *Journal of Applied Psychology*, 94(2):298, 2009.
17. Mark EJ Newman. Mixing patterns in networks. *Physical Review E*, 67(2), 2003.
18. Brian J OLeary and Bart L Weathington. Beyond the business case for diversity in organizations. *Employee Responsibilities and Rights Journal*, 18(4):283–292, 2006.
19. Alejandro Portes. Social capital: Its origins and applications in modern sociology. *Annual review of sociology*, 24(1):1–24, 1998.
20. Sandra Sgoutas-Emch, Lisa Baird, Perla Myers, Michelle Camacho, and Susan Lord. We’re not all white men: using a cohort/cluster approach to diversify stem faculty hiring. *Thought & Action*, 32(1):91–107, 2016.
21. Simon N Williams, Bhoomi K Thakore, and Richard McGee. Providing social support for underrepresented racial and ethnic minority phd students in the biomedical sciences: a career coaching model. *CBELife Sciences Education*, 16(4):ar64, 2017.