Journal Pre-proof

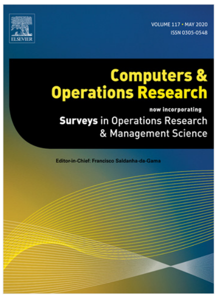

Comparison of a rule-based heuristic and a linear programming model for assigning mentees and mentors in a women in technology mentoring programme

Sarah E. Marshall, Mahsa Mohaghegh

PII: S0305-0548(25)00030-9
DOI: https://doi.org/10.1016/j.cor.2025.107002
Reference: CAOR 107002

To appear in: *Computers and Operations Research*

Received date : 11 February 2024
Revised date : 1 December 2024
Accepted date : 28 January 2025

Please cite this article as: S.E. Marshall and M. Mohaghegh, Comparison of a rule-based heuristic and a linear programming model for assigning mentees and mentors in a women in technology mentoring programme. *Computers and Operations Research* (2025), doi: https://doi.org/10.1016/j.cor.2025.107002.

Highlights

Comparison of a rule-based heuristic and a linear programming model for assigning mentees and mentors in a women in technology mentoring programme

Sarah E. Marshall, Mahsa Mohaghegh

- Defines a rule-based heuristic that provides a simple and easily understandable way to match mentees and mentors
- Demonstrates that the proposed rule-based heuristic performs almost as well and faster than the linear programming model
- Demonstrates that when applied to data from a real-life women in technology mentoring programme, the proposed approaches outperform the original manual approaches to matching

Comparison of a rule-based heuristic and a linear programming model for assigning mentees and mentors in a women in technology mentoring programme

Sarah E. Marshall*, Mahsa Mohaghegh

[a] *Department of Information Systems and Operations Management, Business School, University of Auckland, Private Bag 92019, Auckland, 1142, New Zealand*
[b] *Engineering, Computer & Mathematical Sciences, Auckland University of Technology, Private Bag 92006, Auckland, 1142, New Zealand*

Abstract

Women remain significantly underrepresented in the fields of science, technology, engineering and mathematics (STEM), but positive mentoring relationships can help mitigate the challenges they face when studying and working in these areas. To support female university students in STEM, the Auckland University of Technology (AUT) established the Women in Tech mentorship programme in 2019. Initially, the matching of mentees and mentors was achieved manually, but as the programme's popularity grew, this process became increasingly time consuming. This study addresses the challenges associated with assigning mentees to mentors by automating the matching process based on mentee and mentor attributes. A rule-based heuristic is proposed and compared with a linear programming (LP) approach. Numerical experiments were conducted to evaluate the performance of these algorithms across various scenarios. The rule-based heuristic provides a simple and easily understandable way to allocate mentees and mentors that performs nearly as well as an optimal matching provided by the LP approach. Applying these algorithms to real data from the AUT Women in Tech mentorship programme, it was found that they outperformed manual matching in several performance metrics.

Keywords: Mentoring, Linear Programming, Heuristics, Women in Technology

---

1. Introduction

Universities and other academic institutions have begun developing projects and seminars to encourage more female undergraduate students to pursue science, technology, engineering, and mathematics (STEM) careers. Despite these efforts, there continues to be a shortage of qualified workers in the STEM sector (O'Rourke, 2021). The secret to innovation and economic advancement is diversity (Bokova, 2017), with over 85% of respondents in a Forbes survey of 321 executives from large corporations agreeing that diversity in the workforce is crucial for creativity (Weathers, 2011). Additionally, it has been found that teams with an equal mix of men and women typically generate 40% more income (Weathers, 2011). Therefore, encouraging young women to work in STEM fields

---

*Corresponding author

Email addresses: `sarah.marshall@auckland.ac.nz` (Sarah E. Marshall), `mahsa.mohaghegh@aut.ac.nz` (Mahsa Mohaghegh)

is beneficial to society as a whole. Finding the factors that motivate young women to pursue STEM careers is the first step to ensuring that they enter these fields. Further, recognising the challenges that young women face in STEM education and careers allows a better understanding of where there are limitations in STEM programmes and systems and how to inspire young girls and women to pursue careers in STEM.

In the contemporary world of business, the establishment of strong mentoring relationships is considered indispensable because they provide and lead to valuable personal and professional growth (Bjursell and Florin Sädbom, 2018). Acknowledging the inherent worth of mentorship for both mentors and mentees holds great significance in a dynamic work environment. Fostering a mentoring culture enables organisations to actively encourage collaboration, knowledge sharing and continuous learning among workers and teams. Further, mentoring programmes play a crucial role in the development and success of high school students, university students and individuals in the workplace. While mentoring programmes for each of these different groups differ in focus, they all serve the purpose of guiding and supporting mentees towards their goals.

It is with these factors in mind that the Auckland University of Technology (AUT) Women in Tech mentorship programme was established in 2019. This programme aims to support students at AUT who identify as women to study STEM subjects. At the beginning of each semester, prospective mentees and mentors are asked to complete an online expression of interest form. For mentees, this includes providing information such as their year, programme and subject of study, preferred gender of mentor and their reasons for wanting to participate in the programme. For mentors, this includes providing information such as their gender, programmes and subjects that are aligned with their expertise, whether they are working in academia or industry, a brief biography and their reasons for wanting to participate in the programme. After this initial expression of interest, each mentee is matched with a mentor. Ideally, all mentees are assigned a mentor, and each mentor is assigned a maximum of one mentee. Currently, the process of assigning mentors and mentees is performed manually based on three principles (listed in order of priority): 1) students are assigned a mentor of their preferred gender; 2) first-year students are assigned academic mentors and second-year students and above are assigned industry mentors; and 3) students are assigned a mentor who is aligned with their programme and subject of study. Other information is considered on a case-by-case basis. When the programme began in 2019, there were 15 mentees and 10 academic mentors. In semester 2, 2021 this had increased to 31 mentees and 32 mentors (16 from academia and 16 from industry). While still relatively small, the programme is projected to increase over the coming years. The current matching process is undertaken by an administrative staff member using a series of spreadsheets and is very time consuming, frustrating and prone to human error.

This study aims to improve the efficiency and effectiveness of the mentee/mentor assignment processes by automating the matching of mentors and mentees. Two approaches are presented for matching mentees and mentors: a rule-based heuristic, and a linear programming (LP) model. The rule-based heuristic attempts to replicate the current manual process and provide an easily understandable method for allocating matches. The LP approach, which provides an optimal solution, is primarily used in this study to assess the performance of our proposed rule-based heuristic. The motivation for introducing a rule-based heuristic, rather than relying on an optimal solution from the LP model was the need to provide an easily interpretable and implementable solution.

Both approaches were designed with the AUT Women in Tech mentorship programme in mind, addressing features particular to this program, such as the treatment of gender preferences. However, to maintain the relevance of this work to other mentoring scenarios, the models are presented in this paper in general terms initially, with problem-specific details added when the models are applied to the AUT Women in Tech mentorship programme.

The remainder of this article is structured as follows. Section 2 discusses the relevant literature. Section 3 states the problem that the LP model and rule-based heuristic will address, and introduces the notation used in subsequent sections. Section 4 discusses the LP model, and Section 5 discusses the rule-based heuristic. Section refsec:example presents an example to illustrate the rule-based heuristic and LP approaches, and then Section 7 compares the two approaches using a series of numerical experiments. Section 8 applies the two approaches to real data from the AUT Women in Tech mentorship programme and Section 9 concludes the paper.

2. Literature Review

2.1. Women in STEM

Of all bachelor's degrees in STEM subjects in the United States in 1980, 37.2% were awarded to women (Sassler et al., 2017). By 2010, this increased to 50.3% (Sassler et al., 2017), indicating an increasing trend in women studying STEM subjects. However, the same trend has not been observed in the workforce, with women still being underrepresented in most STEM fields, particularly in computer science and engineering (Sassler et al., 2017). Figure 1 (taken from an article written for the United States Census Bureau by Martinez and Christnacht (2021)) presents the percentage of women in the workforce in various STEM fields from 1970 to 2019. In addition to the low percentages of women in computing and engineering (Figure 1), it is also concerning that there has been declining participation of women in the computing workforce since the 1990s.

Previous studies have attempted to identify the reasons for this gender disparity, with social and cultural expectations (Patterson et al., 2021) and a lack of role models, particularly in senior positions (Onuma et al., 2022; Patterson et al., 2021), being identified as contributing factors. Given the importance of diversity in the workforce (see Section 1), increasing the representation of women in STEM is crucial for the benefit of society in general.

2.2. Mentoring

Mentoring programmes provide an opportunity to promote the engagement and retention of women in STEM. Research consistently highlights the multifaceted benefits of mentoring relationships, particularly for women in STEM (Amelink, 2009); for example, it improves career development (Cross et al., 2019), increases a sense of belonging and self-efficacy (McClain et al., 2021), and decreases drop-out rates of university students in STEM subjects (Kricorian et al., 2020; Zavala et al., 2022).

Mentoring programmes are integral to the growth and success of individuals at different stages of education and career development. At high school level, mentorship initiatives play a crucial role in aiding students with their post-secondary planning, offering support for decisions related to further education or vocational training and personal development (Ahmed et al., 2021). For university students, mentorship schemes focus on academic and career progression, assisting with course and major selection, and securing internships, thereby facilitating informed decision-making, networking, and skills development (Avecilla et al., 2023; Rehman et al., 2022 Jul-Aug). Workplace mentoring programmes are designed to bolster career development in the professional realm by

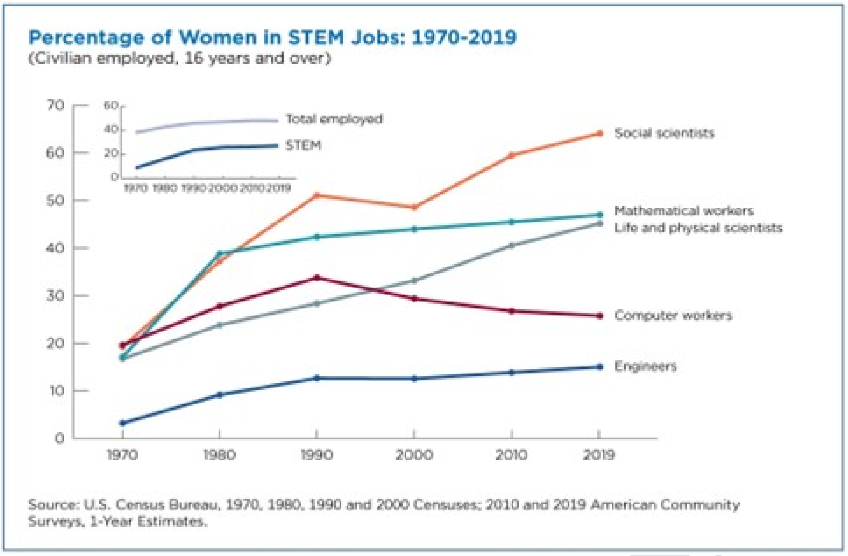

Figure 1: Percentage of women in STEM jobs from 1970 to 2019 (Martinez and Christnacht, 2021)

offering guidance on role-specific skills, industry knowledge, and advancement strategies (Barrett et al., 2017). Despite these differences in focus, all types of mentoring programmes share the common goal of guiding and supporting mentees towards achieving their unique goals.

The success of mentoring programmes hinges on effective pairing between mentors and mentees. Alignment in areas such as shared interests, professional goals, personality traits, and teaching/learning styles is crucial for fostering productive mentoring relationships (Bielczyk et al., 2019). The importance of compatibility in the mentoring relationship highlights the need for a structured and deliberate matching process.

Deng et al. (2022) conducted a qualitative study that identified effective strategies for practitioners to enhance the efficacy of formal mentoring programmes. The study provides three practical recommendations for matching mentors and mentees: 1) prioritise deep-level similarities; 2) consider the developmental needs of the mentee; and 3) seek input from both mentors and mentees. Hee et al. (2020) developed a framework for matching mentors and mentees based on criteria for deep-level similarity. Poulsen (2013) highlights the potential learning opportunities from mentoring programmes for mentors, organisations and society. The article emphasises the importance of establishing clear expectations, setting goals and establishing rules of engagement for collaboration. Barrett et al. (2017) identify the following three key attributes for positive mentoring relationships: 1) active engagement from both mentor and mentee; 2) effective communication between the mentor and mentee; and shared interests between the mentor and mentee. Using regression analysis and considering factors such as academic field, prior research productivity and academic experience, Hoenen and Kolympiris (2020) explored the effect of informal mentoring relationships on the research productivity of early-career scientists at the National Science Foundation. Their findings suggest that early-career scientists who have connections to academic insiders (academics on a two-year secondment with the National Science Foundation) are more likely to secure research funding and achieve higher publication output.

Hagler and Rhodes (2018) studied the benefits to young people of a natural mentoring connection, that is, a relationship with a caring, non-parental figure who provides support and guidance. Employing counterfactual analysis, a quasi-experimental design and controlling for potential confounding variables, Hagler and Rhodes (2018) estimated the causal effect of natural mentoring connections, highlighting the importance and lasting influence of such a connection. Their analysis demonstrates that individuals who had a natural mentor during their adolescence or emerging adulthood reported higher educational attainment, increased engagement in volunteer work, and a larger number of close friendships compared with their peers who did not have mentors. Counterfactual analysis, a quasi-experimental design, is employed to estimate the causal effect of the mentoring intervention while controlling for potential confounding variables.

2.3. *Assigning matches*

The well-known and well-studied issue of allocating tasks or people in one group to individuals in another group is known as the assignment problem (Burkard et al., 2009). The assignment problem can be formulated as a weighted maximum matching problem for bipartite graphs and can be solved using the Hungarian method in strongly polynomial time (Kuhn, 1955). Open-source implementations of the Hungarian method are available via the LEMON Graph Library (2018).

The assignment problem is also a well-known special case of the transportation problem, and can be formulated as a linear programme (Burkard et al., 2009). In this approach, there is a reward (cost) associated with allocating each pair of individuals, and the objective is to select the matches that maximise (minimise) the total reward (cost) obtained. Typical constraints involve the maximum number of tasks assigned to each individual. The assignment problem and the broader class of matching problems have been well-studied in the literature (for further details see Lovász and Plummer (1986); Burkard et al. (2009); Moghaddam (2016)).

The concept of stability, introduced by Gale and Shapley (1962), is an important consideration in matching problems and is widely discussed in the literature. The stable marriage problem involves finding matching between two sets of individuals, in which there is no pair in which both prefer members of another pair (Gale and Shapley, 1962). Recent work in this area has explored improved solution time (Pettersson et al., 2021); quotas (Gale and Shapley, 2013; Ágoston et al., 2022); ties (Ágoston et al., 2022; Pettersson et al., 2021) and applications of the stable marriage problem to other disciplines (Fenoaltea et al., 2021). The issue of stability is not relevant in our application because in our motivating case, matches are assigned rather than chosen, and furthermore, mentees and mentors do not have the visibility of all available mentees/mentors. Thus two participants would have no way of knowing whether there is someone else they would prefer. The short-time frame of the mentoring programme (several months) also reduces the potential practical implications of an unstable match.

Several previous studies have used similar LP models to assign mentees and mentors. For example, Biró and Gyetvai (2022) used a mixed integer linear programming (ILP) model to assign mentee/mentor pairs and study-groups to support students studying online during the early stages of COVID-19. Similar to our study, Haas and Hall (2019) examined a university-based mentoring programme that aimed to support women seeking to enter the computing industry. The problem of assigning mentors to mentees is closely related to other matching problems, such as assigning marriage partners (Cao et al., 2010); referees to conference papers (Garg et al., 2010); hospitals and residents (Pettersson et al., 2021); and colleges to students (Gale and Shapley, 1962; Ágoston et al., 2022; Bobbio et al., 2023).

Other approaches have been considered for assigning matches in a mentoring context. For example, Dutta et al. (2022) study an online question platform and use a neural network to build a recommender system to assign students' (mentees') questions to relevant professionals (mentors).

3. Preliminaries

3.1. Problem description

There is a set of mentees, S and a set of mentors M. In our motivating example, the mentees are students, hence the use of the 'S' to denote mentees. The aim is to achieve a one-to-one matching of mentees and mentors. It is assumed that it is better to provide a mentee with a mentor that does not meet their preferences, than to provide no mentor at all.

Each mentee and mentor has certain attributes. A set of attributes, A, on which to base the matching is selected. For attribute $a \in A$, mentee $s \in S$ has a value of $v_s^a \in V_S^a$ and mentor $m \in M$ has a value of $v_m^a \in V_M^a$. Mentors are assumed to have a broad set of skills and experience, and they can thus be classified using multiple values for some attributes. Therefore, v_m^a is a set of values. The most important value of each attribute is denoted this by $v_m^{a'}$, where $v_m^{a'} \in v_m^a$. In contrast, mentees on the other hand, are assumed to be at an earlier stage in their education and career, and thus have a narrower set of skills and experience. For this reason, mentees can be classified by a single value for each attribute, that is, v_s^a is a single value not a set of values. For an attribute a representing a *preference* of mentee s, then mentees may indicate 'no preference', in which case v_s^a = 'no preference'.

Let $v_{s,m}^a = f(v_s^a, v_m^a)$ denote the score related to a matching of mentee s and mentor m on attribute a. Three types of binary scoring functions are used in the present study:

- one-to-one match: $v_{s,m}^a = \begin{cases} 1 & \text{if } v_s^a = v_m^a \\ 0 & \text{otherwise} \end{cases}$
- flexible one-to-one match: $v_{s,m}^a = \begin{cases} 1 & \text{if } v_s^a = v_m^a, \text{ or} \\ & v_s^a = \text{no preference} \\ 0 & \text{otherwise} \end{cases}$
- one-to-many match: $v_{s,m}^a = \begin{cases} 1 & \text{if } v_s^a \in v_m^a \\ 0 & \text{otherwise} \end{cases}$

The total score or reward, $r_{s,m}$, associated with matching mentee s and mentor m across all attributes is a function of the vector of attribute scores $\mathbf{v}_{s,m} = (v_{s,m}^1, v_{s,m}^2, \ldots, v_{s,m}^{|A|})$ and a vector of weights representing the importance of each attribute $\mathbf{w} = (w_1, w_2, \ldots, w_{|A|})$:

$$r_{s,m}^{\mathbf{w}} = f(\mathbf{v}_{s,m}, \mathbf{w}).$$

A weighted sum is used in the present study, however this could easily be generalised without affecting the methods described in the following sections. Thus the total score is defined as:

$$r_{s,m}^{\mathbf{w}} = f(\mathbf{v}_{s,m}, \mathbf{w}) = \sum_{a \in A} w_a v_{s,m}^a, \tag{1}$$

where $0 \leq w_a \leq 1$ for $a \in A$, and $\sum_{a \in A} w_a = 1$.

Let $x_{s,m} = 1$, if mentee s is matched with mentor m, and 0 otherwise. Then, across all mentees and mentors in a cohort the total score of a matching is defined as

$$R^{\mathbf{w}} = \sum_{s \in S} \sum_{m \in M} x_{s,m} r^{\mathbf{w}}_{s,m} \tag{2}$$

and a cohort's mean score per mentee is

$$\bar{R}^{\mathbf{w}} = \frac{R^{\mathbf{w}}}{|S|}.$$

The minimum cohort score of the matching, that is, the score of the 'worst' match in a cohort is:

$$R^{\mathbf{w}}_{\min} = \min_{s \in S; m \in M : x_{s,m}=1} \{r^{\mathbf{w}}_{s,m}\}.$$

For each attribute $a \in A$, the mean score across all mentees and mentors in a cohort is

$$P_a = \sum_{s \in S} \sum_{m \in M} \frac{x_{s,m} v^a_{s,m}}{|S|}.$$

If the score for each attribute ($v^a_{s,m}$) is binary and mentees receive at most one mentor, then P_a also represents the proportion of mentees receiving a mentor who matches with them on attribute a.

3.2. Application to AUT Women in Tech mentorship programme

In this section the problem described in Section 3.1 will be applied to the motivating example: the AUT Women in Tech mentorship programme. The model could easily be customised for other programmes of a similar nature. In this study, the set of attributes is $A = \{$gender, programme, subject, type$\}$. The values associated with each attribute are provided in Table 1. Note that for the subject attribute, the values provided in Table 1 are for illustrative purposes only and do not represent a full list of subjects available to the students.

In this mentoring programme, all mentees identify as female, so the attribute gender refers to the mentor's gender. For the mentor, gender attribute is the mentor's self-reported gender, and can take the values male, female or other. For the mentee, gender represents the gender of mentor that the mentee prefers, either female, male, other, or no preference. This could be extended to include other values for gender, but it has been restricted to the options provided to the students in the 2021 cohort of the AUT Women in Tech mentorship programme. Given that gender represents a preference, it is scored with the flexible one-to-one matching function.

For the mentee, *programme* and *subject* represent the programme and subject of study corresponding to the mentee's current enrolment. For the mentor, these attributes represent all programmes or subjects, respectively, in which a mentor has expertise. These two attributes are scored using the one-to-many matching function.

For mentees, *type* categorises whether the mentee is a first-year undergraduate (and thus should be matched with an academic mentor) or a second-year or higher student (and thus should be matched with an industry mentor). While alternative allocations could be beneficial in certain scenarios – such as matching a student interested in graduate studies with an academic mentor – the current structure was prioritised to reflect the programme's objectives for the semester 2, 2021 cohort. This approach is aligned with the programme organisers' understanding of typical post-study goals, where earlier-stage students benefit from academic guidance while advanced students gain insights relevant to entering the workforce. For mentors, *type* denotes whether they are affiliated with academia or industry. Type is scored using a one-to-one matching function.

All information is provided in advance via an 'expression of interest' form. Complete information is available for mentees. However, some information may be missing for mentors. If information about attribute $a \in A$ for mentor $m \in M$ is not available, then a score of 0 is used, that is, $v^a_{s,m} = 0$ for all $s \in S$.

Table 1: Attributes and scoring functions

Attribute	**Mentee values V^a_S**	**Mentor values V^a_M**	**Scoring function $v^a_{s,m}$**
Gender	Male, female, other, no preference	Male, female, other	$v^a_{s,m} = \begin{cases} 1 & \text{if } v^a_s = v^a_m, \text{ or} \\ & v^a_s = \text{no preference} \\ 0 & \text{otherwise} \end{cases}$
Programme	Science, engineering, computer and information sciences	Science, engineering, computer and information sciences	$v^a_{s,m} = \begin{cases} 1 & \text{if } v^a_s \in v^a_m \\ 0 & \text{otherwise} \end{cases}$
Subject	For example, analytics, computer science, electrical engineering	For example, analytics, computer science, electrical engineering	$v^a_{s,m} = \begin{cases} 1 & \text{if } v^a_s \in v^a_m \\ 0 & \text{otherwise} \end{cases}$
Type	1 (first year), 0 (not first year)	1 (academic), 0 (industry)	$v^a_{s,m} = \begin{cases} 1 & \text{if } v^a_s = v^a_m \\ 0 & \text{otherwise} \end{cases}$

4. LP formulation

The assignment of mentors to mentees is an example of the well-known (linear sum) assignment problem (Kuhn, 1955). The assignment problem is also a well-known special case of the transportation problem, and can be formulated as an LP problem (Burkard et al., 2009). A standard formulation is presented below.

Let the decision variable $x_{i,j}$ be defined as follows:

$$x_{i,j} = \begin{cases} 1 & \text{if mentee } i \text{ is assigned to mentor } j \\ 0 & \text{otherwise,} \end{cases}$$

and let $r_{i,j}$ denote the reward provided by assigning mentee i to mentor j. If there are equal numbers of mentees and mentors, then the LP problem can be specified as follows:

$$\max \quad Z = \sum_{i=1}^{|S|} \sum_{j=1}^{|M|} r_{i,j} x_{i,j} \tag{3}$$

$$\text{subject to:} \quad \sum_{j=1}^{|M|} x_{i,j} = 1 \quad i = 1, 2, \ldots |S| \tag{4}$$

$$\sum_{i=1}^{|S|} x_{i,j} = 1 \quad j = 1, 2, \ldots |M| \tag{5}$$

$$x_{i,j} \geq 0 \quad \text{for all } i \text{ and } j \tag{6}$$

In this formulation, it is assumed that there are equal numbers of mentees and mentors and that one mentee at the most can be assigned to each mentor, and vice-versa. This formulation can be easily modified to represent alternative scenarios by altering the right hand side values, and/or the signs of equations (4) and (5), as appropriate. Several variations are detailed below:

- more mentees than mentors: equation (4) becomes $\sum_{j=1}^{|M|} x_{i,j} \leq 1, \quad i = 1, 2, \ldots |S|$
- more mentors than mentees: equation (5) becomes $\sum_{i=1}^{|S|} x_{i,j} \leq 1, \quad j = 1, 2, \ldots |M|$
- mentors can be assigned up to n mentees: equation (5) becomes $\sum_{i=1}^{|S|} x_{i,j} \leq n, \quad j = 1, 2, \ldots |M|$

If $r_{i,j}$ in the objective function (equation (3)) is calculated the same way as equation (1), and if the same weights are used, then, the objective function will equal the total score in equation (2), that is, $Z = R^{\mathbf{w}}$. However, there may be cases in which the one set of weights is used within the LP formulation and another set of weights is used to compute the total score; for example, when comparing the performance of different approaches.

The LP can be solved using one of many commercial or open-source packages that are currently available for solving LPs, transportation problems or assignment problems.

5. Rule-based Heuristic

5.1. Overview

This section describes a rule-based heuristic for assigning mentees and mentors. The heuristic is simple enough that it could be implemented manually using a spreadsheet; however, it would be more efficient if the heuristic were automated using a programming language. The heuristic is described here using the specific characteristics of the motivating example; however, it could easily be generalised to include characteristics relevant to other mentoring programmes.

The basic idea is that in each iteration a set of matching criteria are chosen, and all unmatched mentees are assigned an unmatched mentor who matches them on all the selected criteria. The matching criteria are chosen from the mentee/mentor attributes. In subsequent iterations, the matching criteria are less restrictive enabling less desirable matches to occur. In this sense, the heuristic can be described as ‘greedy’, with the ‘best’ matches occurring first. Furthermore, given that the mentees have the choice of indicating ‘no preference’ for the preferred gender of the mentor in our current application, a filtering option is provided that enables only mentees who selected ‘no preference’ to be matched in a particular iteration.

5.2. Detailed description

Let D_S^0 (D_M^0) denote a dataset of mentees (mentors), in which the rows represent unique mentees (mentors) and the columns represent their attributes. A general description of the heuristic is provided below:

0. **Initialise:** set $D_S = D_S^0$ and $D_M = D_M^0$ and initialise $D = \emptyset$
1. **Filter:** apply current filtering criteria (if applicable) to remove rows from D_S (D_M) and update D_S (D_M) with resulting dataset
2. **Sort, select, number mentees:**
    sort D_S (unmatched, filtered mentees) by the current matching criteria
    add a group variable representing each unique combination of the matching criteria
    add an index variable numbering the mentees within each group
    update D_S to include the group and index variables
3. **Sort, select, number mentors:**
    sort D_M (unmatched, filtered mentors) by the current matching criteria
    add a group variable representing each unique combination of the matching criteria
    add an index variable numbering the mentors within each group
    update D_M to include the group and index variables
4. **Identify new matches:** inner join D_S and D_M by the group and index variables and add to D
5. **Identify unmatched mentees:** D_S = all mentees in D_S^0 that do not appear in D
6. **Identify unmatched mentors:** D_M = all mentees in D_M^0 that do not appear in D
7. **Return or stop:** if $D_S = \emptyset$ or $D_M = \emptyset$ then, stop as matching is complete; otherwise return to step 1.

In our application, the heuristic assigns mentees and mentors using matching criteria based on the attributes presented in Table 1. For the mentor, if multiple programmes and subjects were selected, then only the one assigned as most relevant, that is, $v_m^{a'}$ is used in the heuristic. Note, that this contrasts with the LP model, which retains all values of each attribute. Multiple versions of this heuristic can be obtained by selecting different matching and filtering criteria. An example of some selection criteria are presented in Table 2.

Table 2: Example of selection criteria for a rule-based heuristic with six iterations

Iteration	Matching criteria	Filtering criteria
1	Gender, Type, Programme	
2	Type, Programme	v_s^a =‘No preference’
3	Gender, Type	
4	Type	v_s^a =‘No preference’
5	Gender	
6	No criteria	

5.3. Additional explanation

The heuristic will now be described in more detail, referring to the example criteria in Table 2. During the first iteration, mentees and mentors are matched based on gender, type and programme. To do this, the mentees dataset D_S and mentors dataset D_M are sorted based on gender, type and programme. Using the problem description in Section 3.2, there are four gender categories and two type categories. Assume that $V_s^{programme} = V_m^{programme}$ and thus that $|V_s^{programme}| = |V_m^{programme}| = n_p$. This creates at most 4 (gender) × 2 (type) × n_p (programme) = $8n_p$ groups of mentees, and at most 3 (gender) × 2 (type) × n_p (programme) = $6n_p$ groups of mentors. Some groups may not have any mentees or mentors. Groups are numbered so that mentee group k represents the same attribute values as mentor group k. Within each group, the mentors and mentees can be sorted arbitrarily. A new index variable is added to both the mentees and mentors datasets, and this index variable numbers the people in each group of these groups. This is added to provide a unique identifier of members within each group and to provide a joining variable to be used in the next step. An inner join is performed between the mentees dataset D_S and mentors dataset D_M, using the newly created group and index variables. All matched mentees and mentors are added to a master matches dataset D. The mentees dataset D_S and mentors dataset D_M are updated to contain only unmatched mentees and mentors, respectively, and then the heuristic moves to the next iteration.

During the second iteration, filtering is performed so that only unmatched mentees with no preference for mentor gender are selected. These mentees are matched with the unmatched mentors based on type and programme. The remaining steps are the same as in the first iteration. During the third iteration, unmatched mentees and unmatched mentors are matched based on gender and type. In the fourth iteration, unmatched mentees with no preference for mentor gender are selected and matched with unmatched mentors based on type. In the fifth iteration, mentees are matched with mentors of their preferred gender. The final iteration assigns all remaining unmatched mentees with unmatched mentors.

6. Simple example

To demonstrate the algorithms described in Sections 4 and 5, a simple example with six mentees and seven mentors will be used. The attribute values for mentees and mentors are presented in Table 3. In this example, there are two programmes (science (Sc) and engineering (Eng)) and five different subjects (denoted as a, b, c, d and e). Recall that type refers to whether a mentee is a first-year student (1) or not (0), and whether a mentor is based in academia (1) or industry (0). To replicate a scenario observed in the AUT Women in Tech mentorship programme, occasionally mentors do not select any relevant programmes or subjects in the expression of interest form. We assume that mentees complete all fields.

In this example, we compare the two approaches described above (the rule-based heuristic and the LP model). Two variations of each approach are considered, giving a total of four algorithms.

Recall that the total score corresponding to assigning mentee s with mentor m, denoted by $r_{s,m}^{\mathbf{w}}$ (see equation (1)) is a weighted sum. A value $r_{s,m}^{\mathbf{w}} = 1$ indicates that mentee s matches with mentor m on all included attributes. To enable a fair comparison between the algorithms, a consistent set of weights is used to compute the total score. In particular, the weights are 0.5 for gender, 0.25 for type, 0.15 for programme and 0.1 for subject. These weights are used in computing the overall scores, and may differ from the weights used in the LP model objective function.

Table 3: Simple example: attributes of mentees and mentors

Mentee, s					Mentor, m					
Name	Gender pref.	Prog.	Type	Subj.	Name	Gender	All prog.	Prog.	Type	Subj.
MEE1	Male	Eng	0	e	MOR1	Male	Eng;Sc	Eng	1	c
MEE2	Female	Eng	1	c	MOR2	Male	Eng;Sc	Sc	1	
MEE3	Female	Sc	0	b	MOR3	Male	Eng;Sc	Eng	1	c
MEE4	No pref	Eng	1	c	MOR4	Male	Eng;Sc	Sc	0	b;c
MEE5	No pref	Sc	1	c	MOR5	Female	Eng	Eng	0	c
MEE6	Female	Eng	0	d	MOR6	Female			1	
					MOR7	Female	Eng;Sc	Sc	0	a;c

pref. = preference; Prog. = Programme; All prog. = All programmes; Subj. = Subjects

6.1. Application of the LP model to the simple example

Two variations of the LP were computed, each with a different set of objective function weights, as presented in Table 4. The first variation, denoted LP1, places a priority on gender, with less emphasis on the remaining attributes. The assignment of mentees and mentors for the first variation, along with the score for each match and the cohort total score, are presented in Table 5. In the first variation of the LP model, mentee MEE1 was matched with mentor MOR4. The score for this match was 0.9 because the pair matched on gender, type and programme. Recall, that in the LP model, all programmes selected by a mentor are used to determine the score of a pair. Mentor MOR2 was not matched with a mentee. The second variation, denoted LP2, places equal priority on gender and type, and no weight on programme or subject. The assignment of mentees and mentors for the second variation, along with the score for each match and the cohort total score, are presented in Table 6. Additional summary statistics and matching percentages for both variations are presented in Table 10.

Table 4: Simple example: weights used in LP objective function

Algorithm	Gender preference	Type	Programme	Subject	Description
LP1	0.50	0.25	0.15	0.10	Gender priority
LP2	0.50	0.50	0.00	0.00	Gender and type only

Table 5: Simple example: mentee/mentor matches obtained using algorithm LP1

Mentee, s					Mentor, m						
Name	Gender pref.	Prog.	Type	Subj.	Name	Gender	All prog.	Prog.	Type	Subj.	Score $r^{\mathbf{w}}_{s,m}$
MEE1	Male	Eng	0	e	MOR4	Male	Eng;Sc	Sc	0	b;c	0.90
MEE2	Female	Eng	1	c	MOR6	Female			1		0.75
MEE3	Female	Sc	0	b	MOR7	Female	Eng;Sc	Sc	0	a;c	0.90
MEE4	No pref	Eng	1	c	MOR1	Male	Eng;Sc	Eng	1	c	1.00
MEE5	No pref	Sc	1	c	MOR3	Male	Eng;Sc	Eng	1	c	1.00
MEE6	Female	Eng	0	d	MOR5	Female	Eng	Eng	0	c	0.90
					MOR2	Male	Eng;Sc	Sc	1		
Cohort total score $R^{\mathbf{w}}$											5.45

pref. = preference; Prog. = Programme; All prog. = All programmes; Subj. = Subjects

Table 6: Simple example: mentee/mentor matches obtained using algorithm LP2

Mentee, s					Mentor, m						
Name	Gender pref.	Prog.	Type	Subj.	Name	Gender	All prog.	Prog.	Type	Subj.	Score $r^{\mathbf{w}}_{s,m}$
MEE1	Male	Eng	0	e	MOR4	Male	Eng;Sc	Sc	0	b;c	0.90
MEE2	Female	Eng	1	c	MOR6	Female			1		0.75
MEE3	Female	Sc	0	b	MOR5	Female	Eng	Eng	0	c	0.75
MEE4	No pref	Eng	1	c	MOR1	Male	Eng;Sc	Eng	1	c	1.00
MEE5	No pref	Sc	1	c	MOR2	Male	Eng;Sc	Sc	1		0.90
MEE6	Female	Eng	0	d	MOR7	Female	Eng;Sc	Sc	0	a;c	0.90
					MOR3	Male	Eng;Sc	Eng	1	c	
Cohort total score $R^{\mathbf{w}}$											5.2

pref. = preference; Prog. = Programme; All prog. = All programmes; Subj. = Subjects

6.2. Application of the rule-based heuristic to the simple example

Two variations of the rule-based heuristic are used in this simple example, a version with six iterations presented in Table 2 (we refer to this algorithm as RBi) and a simpler version with two iterations presented in Table 7 (we refer to this algorithm as RBii).

Table 7: Example of selection criteria for rule-based heuristic with two iterations (RBii)

Iteration	Matching criteria	Filtering criteria
1	Gender, Type, Programme	
2	No criteria	

Table 8 presents the assignment of mentees and mentors, along with the score for each match and the cohort total score, under the six-iteration variation rule-based heuristic, RBi. The column 'iteration' indicates the iteration of the algorithm in which the match was made. For example, mentee MEE1 was matched with mentor MOR4 in iteration 3. Recall that the rule-based heuristic uses only the most relevant programme, so this pair was not matched in iteration 1. However, the score is calculated based on all programmes so the match for this pair is 0.9 compared with 0.75 for MEE2 and MOR6 who also matched in iteration 3. Mentor MOR3 was not matched with any mentees.

Table 8: Simple example: mentee/mentor matches obtained using algorithm RBi (six-iterations)

Mentee, s					Mentor, m							
Name	Gender pref.	Prog.	Type	Subj.	Name	Gender	All prog.	Prog.	Type	Subj.	It.	Score $r^{\mathbf{w}}_{s,m}$
MEE1	Male	Eng	0	e	MOR4	Male	Eng;Sc	Sc	0	b;c	3	0.90
MEE2	Female	Eng	1	c	MOR6	Female			1		3	0.75
MEE3	Female	Sc	0	b	MOR7	Female	Eng;Sc	Sc	0	a;c	1	0.90
MEE4	No pref	Eng	1	c	MOR1	Male	Eng;Sc	Eng	1	c	2	1.00
MEE5	No pref	Sc	1	c	MOR2	Male	Eng;Sc	Sc	1		2	0.90
MEE6	Female	Eng	0	d	MOR5	Female	Eng	Eng	0	c	1	0.90
					MOR3	Male	Eng;Sc	Eng	1	c		
Cohort total score $R^{\mathbf{w}}$												5.35

pref. = preference; Prog. = Programme; All prog. = All programmes; Subj. = Subjects; It. = Iteration

Table 9 presents assignment of mentees and mentors, along with the score for each match and the cohort total score, under the two-iteration rule-based heuristic, RBii. In this version of the rule-based heuristic mentee MEE1 was matched with mentor MOR1 in iteration 2. Mentor MOR6 was not matched with any mentees.

Table 9: Simple example: mentee/mentor matches obtained using algorithm RBii (two-iterations)

Mentee, s					Mentor, m							
Name	Gender pref.	Prog.	Type	Subj.	Name	Gender	All prog.	Prog.	Type	Subj.	It.	Score $r^{\mathbf{w}}_{s,m}$
MEE1	Male	Eng	0	e	MOR1	Male	Eng;Sc	Eng	1	c	2	0.65
MEE2	Female	Eng	1	c	MOR2	Male	Eng;Sc	Sc	1		2	0.40
MEE3	Female	Sc	0	b	MOR7	Female	Eng;Sc	Sc	0	a;c	1	0.90
MEE4	No pref	Eng	1	c	MOR3	Male	Eng;Sc	Eng	1	c	2	1.00
MEE5	No pref	Sc	1	c	MOR4	Male	Eng;Sc	Sc	0	b;c	2	0.75
MEE6	Female	Eng	0	d	MOR5	Female	Eng	Eng	0	c	1	0.90
					MOR6	Female			1			
Cohort total score $R^{\mathbf{w}}$												4.6

pref. = preference; Prog. = Programme; All prog. = All programmes; Subj. = Subjects; It. = Iteration

6.3. Comparison of algorithm performance on the simple example

Table 10 provides a summary of the performance of the two rule-based heuristic and the two LP models. The table provides the total cohort score of the matching $R^{\mathbf{w}}$, the mean cohort score of the matching $\bar{R}^{\mathbf{w}}$, and the minimum cohort score of the matching $R^{\mathbf{w}}_{\text{min}}$. These represent the sum, mean and minimum, respectively, of the score column in Tables 5, 6, 8 and 9. If all mentees received a mentor that matched based on gender, programme, type and subject, then the mean match score $\bar{R}^{\mathbf{w}}$ would be 1 and the total match score $R^{\mathbf{w}}$ would be equal to the number of mentees, which in this example is six. Notice that the total score is the same as the objective function value for LP1, the LP approach with the same weights used in scoring, but differs for LP2. The objective function is not available for the rule-based heuristic approaches. The final four columns of Table 10 provide the percentage of mentees who were assigned a mentor that matched them on each attribute. For example, all mentees were assigned a mentor of their preferred gender under LP1, LP2 and RBi, but only 83.3% (5/6) under RBii.

For this simple example, LP1 performs at least as well as the other methods on all scoring criteria, and on three out of four of the percentage matching criteria. This is closely followed by RBi, which has the second highest total score and mean score, and performs as well as LP1 on three of the four percentage matching criteria.

Table 10: Simple example: summary results for all algorithms

					% of mentees with matching mentor			
Algorithm	Total score $R^{\mathbf{w}}$	Obj val Z	Mean score $\bar{R}^{\mathbf{w}}$	Min score $R^{\mathbf{w}}_{\min}$	Gender pref.	Type	Prog.	Subj.
LP1	5.45	5.45	0.91	0.75	100.0	100.0	83.3	33.3
LP2	5.20	6.00	0.87	0.75	100.0	100.0	66.7	16.7
RBi	5.35		0.89	0.75	100.0	100.0	83.3	16.7
RBii	4.60		0.77	0.40	83.3	66.7	100.0	33.3

7. Numerical experiments

7.1. Scenarios

Numerical experiments were conducted to investigate the effectiveness of the LP models and the rule-based heuristics. In these experiments, a range of scenarios were considered, in which the nature of the groups of mentees and mentors was varied in relation to gender and type. For example, in some scenarios there are, on average, equal numbers of mentors and mentees, and in some there are not. In some scenarios there are approximately equal proportions of male and female mentors and in other scenarios there are not. These scenarios were chosen to mirror a variety of scenarios that might be encountered when running a Women in Tech mentorship programme.

For each scenario, a set of mentees and a set of mentors are randomly generated using the properties in Table 11. To capture the variability associated with this approach, 1000 replications of each scenario were performed, each with a randomly generated set of mentees and mentors.

As presented in Table 11, three types of programmes were considered: computer and information science (CIS), science (Sc) and engineering (Eng). Subjects are labelled by letters of the alphabet, and are chosen to allow some overlap between programmes (e.g. suppose the f represents software engineering and that this programme may be studied under either an engineering or a computer and information sciences programme). The numbers in the columns labelled 'Scenario 1' – 'Scenario 3' represent the average proportion of mentees/mentors with each of the values listed in the column 'Values'. For example, under mentee scenario 1, when the mentees are randomly generated, a mentee prefers a female mentor with probability 0.5, a male mentor with probability 0.4 and has no preference with probability 0.1. Thus, across all replications, on average 50% of mentees will prefer a female mentor in this scenario. On average, programme and subjects are evenly represented across the replications of each scenario. However, the proportions of each programme and subject may vary in individual replications. Overall, three scenarios for mentees and three for mentors were used, giving a total of nine scenario combinations.

Table 11: Numerical experiments: parameter values and combinations for creating scenarios

Mentee				
Attribute	Values	Scenario 1	Scenario 2	Scenario 3
Gender pref.	Female,male,other,no pref	0.5,0.4,0,0.1	0.8,0.1,0.01,0.09	0.8,0.1,0.01,0.09
Type	1,0	0.5,0.5	0.5,0.5	0.7,0.3
Prog.	CIS,Sc,Eng	0.333,0.333,0.333	0.333,0.333,0.333	0.333,0.333,0.333
Subject: CIS	d,e,f	0.333,0.333,0.333	0.333,0.333,0.333	0.333,0.333,0.333
Subject: Eng	f,g,h	0.333,0.333,0.333	0.333,0.333,0.333	0.333,0.333,0.333
Subject: Sc	a,b,c,d	0.25,0.25,0.25,0.25	0.25,0.25,0.25,0.25	0.25,0.25,0.25,0.25
Mentor				
Attribute	Values	Scenario 1	Scenario 2	Scenario 3
Gender	Female,male,other	0.5,0.5,0	0.7,0.29,0.01	0.7,0.29,0.01
Type	1,0	0.5,0.5	0.5,0.5	0.7,0.3
Prog.	CIS,Sc,Eng	0.333,0.333,0.333	0.333,0.333,0.333	0.333,0.333,0.333
Subject: CIS	d,e,f	0.333,0.333,0.333	0.333,0.333,0.333	0.333,0.333,0.333
Subject: Eng	f,g,h	0.333,0.333,0.333	0.333,0.333,0.333	0.333,0.333,0.333
Subject: Sc	a,b,c,d	0.25,0.25,0.25,0.25	0.25,0.25,0.25,0.25	0.25,0.25,0.25,0.25

Three different mentoring programme sizes are considered: small, $|S|$=20; medium, $|S|$=100; and large, $|S|$=250. For each size, three variations are considered: the number of mentors is the same as the number of mentees, $|S|$=$|M|$; the number of mentors is 10% more than the number of mentees, $|M| = 1.1\ |S|$; and the number of mentors is 10% less than the number of mentees $|M| = 0.9\ |S|$. This gives a nine size combinations. Thus, in total there are $9 \times 9 = 81$ scenario and size combinations.

7.2. Algorithms

Six variations of the LP model and the rule-based heuristic were applied to the randomly generated datasets, giving a total of 12 algorithms. The variations of the LP approach were defined using a different set of weights in the objective function, as presented in Table 12. The variations of the rule-based heuristic were defined using a different set of matching and filtering criteria, as presented in Table 13. Both the LP and rule-based approaches were implemented in R (version 4.3.1)(R Core Team, 2023). The LP implementation used the `lp.transport` function from the `lpSolve` package (version 5.5) (Berkelaar et al., 2023).

7.3. Performance metrics

To enable consistent scoring across all algorithms, and to reflect the overall intentions of this mentoring programme, $r^{\mathbf{w}}_{s,m}$, the score for matching mentee s with mentor m, was computed using weights corresponding to the 'gender priority' scenario (row 1 of Table 12, i.e. 0.5 for gender, 0.25 for type, 0.15 for programme and 0.1 for subject). For each cohort and algorithm, the total $R^{\mathbf{w}}$, mean

Table 12: Numerical experiments: weights used in LP objective function

Algorithm	Gender preference	Type	Programme	Subject	Description
LP1	0.50	0.25	0.15	0.10	Gender priority
LP2	0.25	0.25	0.25	0.25	Equal
LP3	0.50	0.50	0.00	0.00	Type gender
LP4	1.00	0.00	0.00	0.00	Gender only
LP5	0.00	1.00	0.00	0.00	Type only
LP6	0.00	0.00	1.00	0.00	Programme only

$\bar{R}^{\mathbf{w}}$ and minimum $R^{\mathbf{w}}_{\mathrm{min}}$ match scores were computed, as in Table 10. The percentage of mentees per cohort matching with each mentor attribute are computed based on the total number of *matched* mentees. Thus, even when there are fewer mentors than mentees, the maximum percentage possible is 100, even though there will be some mentees who are not assigned a mentor.

7.4. Results

For each of the 81 scenarios, the algorithms were applied to 1000 randomly-generated cohorts of mentees and mentors. The performance metrics were computed for each replication of each scenario. Summary statistics were computed for each cohort to give an overall indication of the performance of an algorithm for a particular cohort. The cohort scores are discussed in Section 7.4.1 and the attribute match percentages are discussed in Section 7.4.2.

7.4.1. Analysis of cohort scores

The distribution of mean cohort total score $\bar{R}^{\mathbf{w}}$ across 1000 replications is presented in Figure 2 for mentee scenario 1 and mentor scenario 1, for small-, medium- and large-sized mentoring programmes with equal numbers of mentees and mentors. As presented in Figure 2, for all three programme sizes, LP1 performed the best, followed by LP2, RB1 and RB2. The LP methods are listed in Table 12 (and thus numbered) in the order in which they were expected to perform, on average. This was confirmed by the numerical experiments, as presented in Figure 2. Similarly for the rule-based heuristic, the order in which the variations are presented in Table 13 also reflected the expected and actual performance.

Table 13: Numerical experiments: criteria used in rule-based heuristics

Algorithm	Iteration	Matching Criteria	Filtering Criteria
RB1	1	Gender, type, programme, subject	
	2	Type, programme, subject	Gender preference = ‘no preference’
	3	Gender, type, programme	
	4	Type, programme	Gender preference = ‘no preference’
	5	Gender, type	
	6	Type	Gender preference = ‘no preference’
	7	Gender	
	8	Type, programme	
	9	Type	
	10	Programme	
	11		
RB2	1	Gender, type, programme	
	2	Type, programme	Gender preference = ‘no preference’
	3	Gender, type	
	4	Type	Gender preference = ‘no preference’
	5	Gender	
	6	Type, programme	
	7	Type	
	8	Programme	
	9		
RB3	1	Gender, type, programme	
	2	Type, programme	Gender preference = ‘no preference’
	3	Gender, type	
	4	Type	Gender preference = ‘no preference’
	5	Gender	
	6		
RB4	1	Gender, type, programme	
	2	Type, programme	Gender preference = ‘no preference’
	3	Gender, type	
	4		
RB5	1	Gender, type, programme	
	2	Gender	
	3		
RB6	1	Gender, type	
	2		

Table 14 presents the summary statistics for the mean match score across 1000 replications for mentee scenario 1 and mentor scenario 1 with equal numbers of mentees and mentors. The numbers in this table support results in Figure 2. It can be observed that the mean of the mean match scores increases with the size of the programme, suggesting that the average quality of the matches improves with larger programme sizes. This is expected because within a larger pool of mentees and mentors, there is a greater chance of finding a good match. For the same reason, variation in the mean score was much lower in the large programme of 250 mentees and mentors, compared with the small programme with 20 mentees and mentors, as demonstrated by the standard

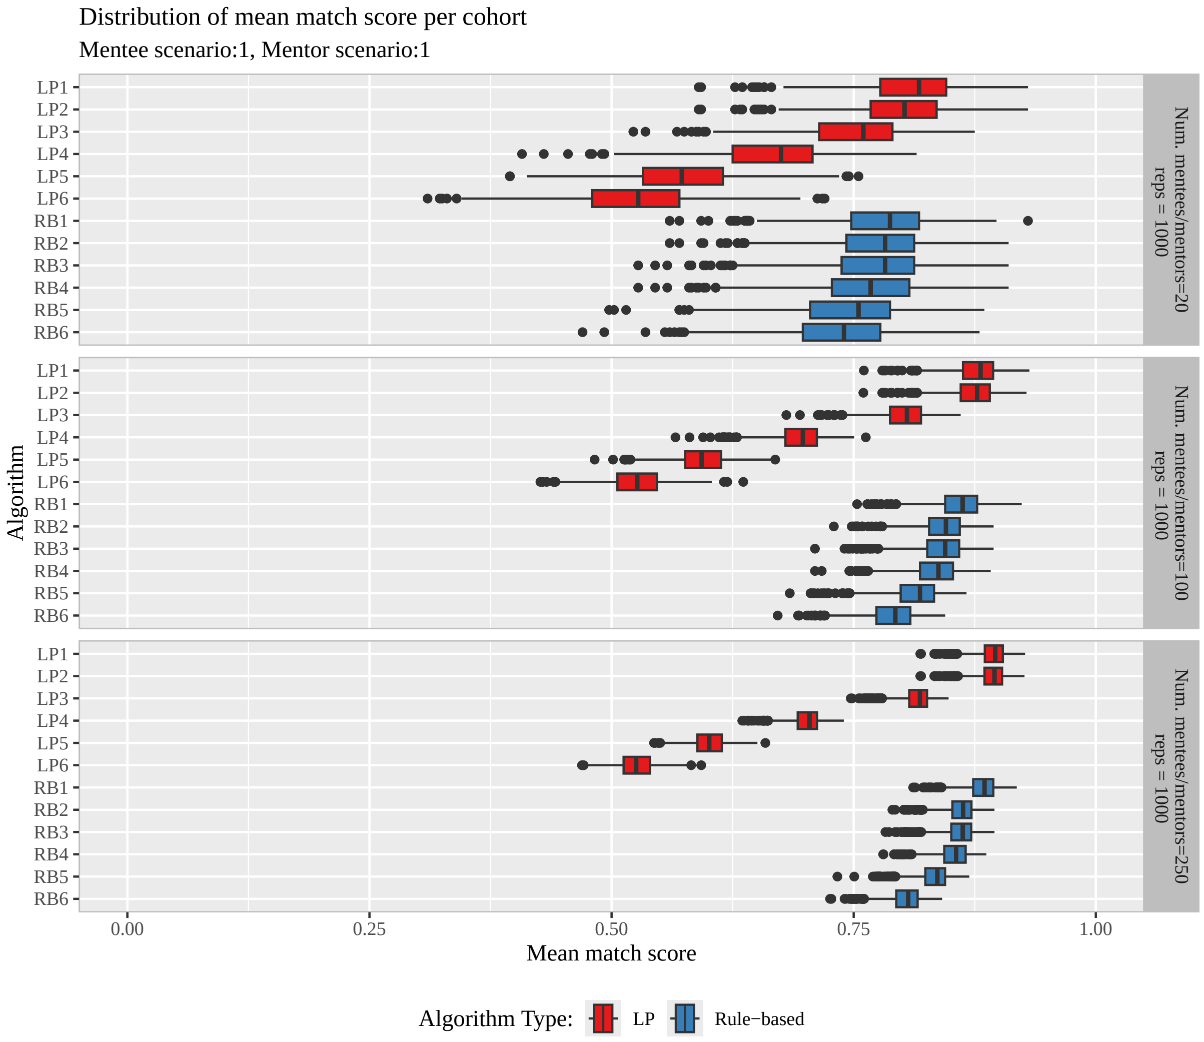

Figure 2: Numerical experiments: distribution of mean match scores $\bar{R}^{\mathrm{w}}$ for all algorithms across 1000 cohorts of mentees (randomly-generated from mentee scenario 1) and mentors (randomly-generated from mentor scenario 1), with equal numbers of mentees and mentors

deviations in Table 14 and the interquartile range in Figure 2.

Similar plots were created and examined for the remainder of the 81 scenarios and similar trends were found. As presented in Table 15, in all 81 scenarios, LP1, LP2, RB1 were the first, second, third place, respectively. RB2 was fourth and RB3 was fifth in all but one scenario.

While the best two LP approaches (LP1 and LP2) perform better than the rule-based heuristics, there is not much difference between these two LP approaches and the best four rule-based heuristics (RB1 – RB4). As presented in Table 14, the maximum difference between LP1 and RB4 is 0.05, suggesting the average ‘quality’ of the matches provided is comparable. However, it is important to also consider the quality of the worst matches in a cohort. For example, it is important to consider whether some approaches provide good matches, on average, but provide some mentees with poorly

Table 14: Numerical experiments: summary statistics of mean match scores $\bar{R}^{\mathbf{w}}$ for all algorithms for 1000 cohorts of mentees (randomly generated from mentee scenario 1) and mentors (randomly generated from mentor scenario 1)

		$\|S\| = \|M\| = 20$			$\|S\| = \|M\| = 100$			$\|S\| = \|M\| = 250$		
Alg.	n	mean	median	s.d.	mean	median	s.d.	mean	median	s.d.
LP1	1000	0.808	0.818	0.052	0.877	0.881	0.024	0.893	0.896	0.016
LP2	1000	0.798	0.802	0.050	0.874	0.878	0.024	0.893	0.895	0.016
LP3	1000	0.750	0.760	0.055	0.801	0.805	0.024	0.815	0.818	0.016
LP4	1000	0.664	0.675	0.061	0.694	0.698	0.026	0.701	0.704	0.016
LP5	1000	0.573	0.573	0.063	0.594	0.593	0.028	0.601	0.601	0.018
LP6	1000	0.524	0.527	0.065	0.525	0.526	0.030	0.526	0.525	0.019
RB1	1000	0.780	0.787	0.054	0.859	0.863	0.025	0.882	0.885	0.017
RB2	1000	0.775	0.782	0.053	0.842	0.845	0.024	0.860	0.863	0.016
RB3	1000	0.772	0.782	0.057	0.840	0.845	0.026	0.859	0.863	0.017
RB4	1000	0.763	0.767	0.059	0.834	0.838	0.027	0.853	0.856	0.017
RB5	1000	0.743	0.755	0.060	0.814	0.818	0.028	0.832	0.836	0.018
RB6	1000	0.735	0.740	0.059	0.789	0.793	0.027	0.803	0.806	0.018

matched mentors. To investigate such as possibility, the distribution of the worst match (i.e. the score of the mentee/mentor pair with the minimum match score $R^{\mathbf{w}}_{\text{min}}$) in each cohort was examined. As presented in Figure 3 and Table 16, under the best LP approaches, the worst matches are, on average, not as bad as the worst matches under the rule-based heuristics. This is not surprising because the rule-based heuristics have a greedy approach, assigning the best matches first, and then assigning whoever is left at the end. In contrast, the LP models take a more holistic view in that they have the match scores from all pairs included in the objective function.

Table 15: Numerical experiments: placing of algorithms for all scenarios for all programme sizes across 1000 replications, based on mean match score per cohort $\bar{R}^{\mathbf{w}}$.

	Place											
Algorithm	1	2	3	4	5	6	7	8	9	10	11	12
LP1	81	0	0	0	0	0	0	0	0	0	0	0
LP2	0	81	0	0	0	0	0	0	0	0	0	0
LP3	0	0	0	0	0	4	32	45	0	0	0	0
LP4	0	0	0	0	0	0	0	0	0	68	13	0
LP5	0	0	0	0	0	0	0	0	0	13	68	0
LP6	0	0	0	0	0	0	0	0	0	0	0	81
RB1	0	0	81	0	0	0	0	0	0	0	0	0
RB2	0	0	0	80	1	0	0	0	0	0	0	0
RB3	0	0	0	1	80	0	0	0	0	0	0	0
RB4	0	0	0	0	0	72	9	0	0	0	0	0
RB5	0	0	0	0	0	5	40	36	0	0	0	0
RB6	0	0	0	0	0	0	0	0	81	0	0	0

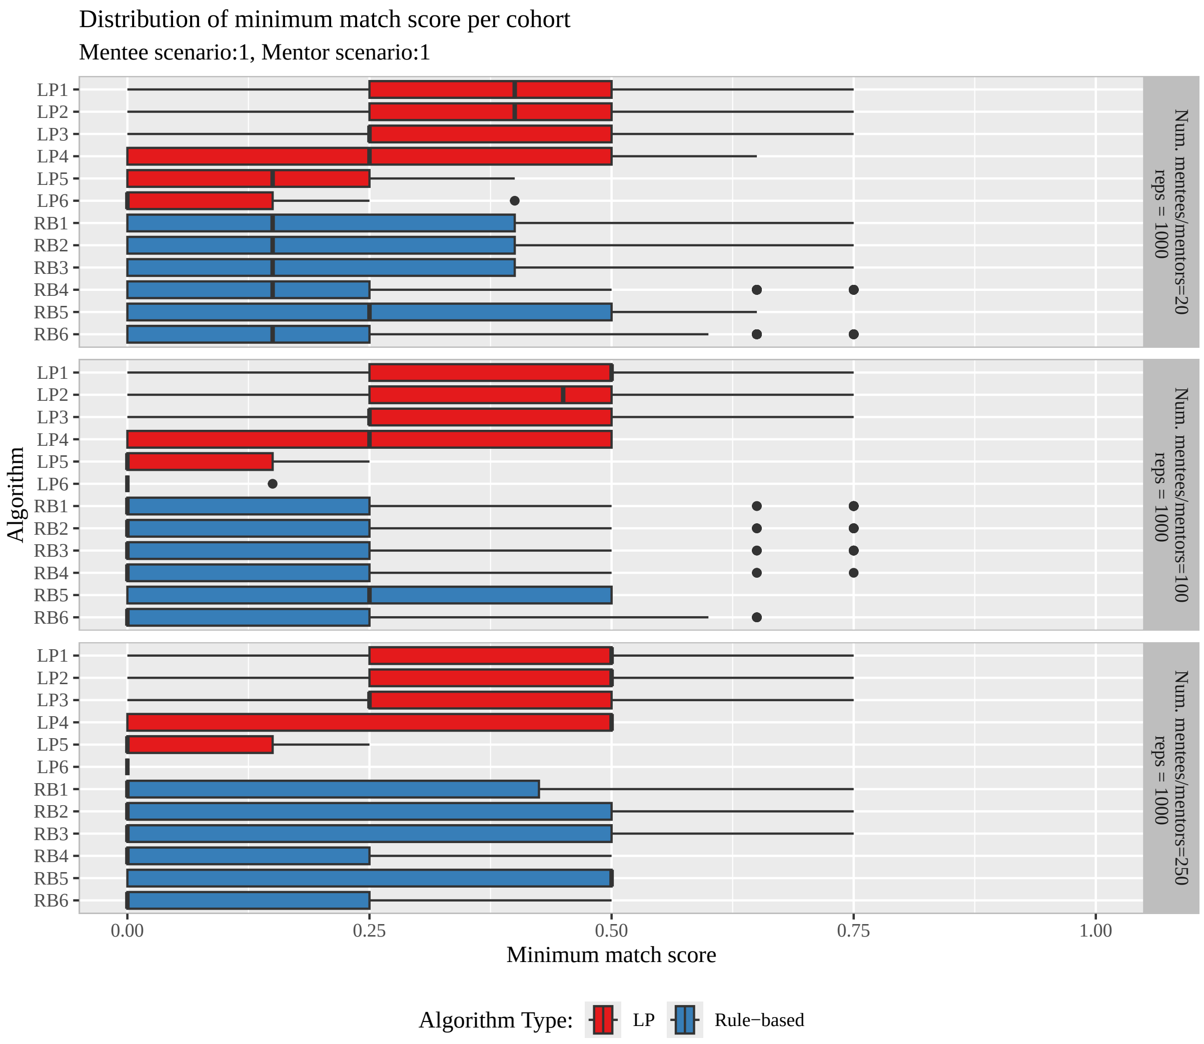

Figure 3: Distribution of minimum match scores $R^{\mathbf{w}}_{\text{min}}$ for algorithms across 1000 cohorts of mentees (randomly generated from mentee scenario 1) and mentors (randomly-generated from mentor scenario 1) and in which there are equal numbers of mentees and mentors

Table 16: Numerical experiments: summary statistics of minimum match scores $R^{\mathbf{w}}_{\min}$ for all algorithms for 1000 cohorts of mentees (randomly generated from mentee scenario 1) and mentors (randomly generated from mentor scenario 1)

		$\|S\| = \|M\| = 20$			$\|S\| = \|M\| = 100$			$\|S\| = \|M\| = 250$		
Alg.	n	mean	median	s.d.	mean	median	s.d.	mean	median	s.d.
LP1	1000	0.414	0.400	0.167	0.440	0.500	0.165	0.433	0.500	0.159
LP2	1000	0.377	0.400	0.149	0.404	0.450	0.146	0.417	0.500	0.152
LP3	1000	0.326	0.250	0.170	0.323	0.250	0.168	0.339	0.250	0.171
LP4	1000	0.248	0.250	0.214	0.257	0.250	0.229	0.282	0.500	0.236
LP5	1000	0.124	0.150	0.114	0.071	0.000	0.101	0.054	0.000	0.092
LP6	1000	0.046	0.000	0.072	0.000	0.000	0.005	0.000	0.000	0.000
RB1	1000	0.207	0.150	0.204	0.155	0.000	0.199	0.172	0.000	0.216
RB2	1000	0.212	0.150	0.208	0.164	0.000	0.201	0.183	0.000	0.219
RB3	1000	0.206	0.150	0.211	0.159	0.000	0.201	0.182	0.000	0.219
RB4	1000	0.166	0.150	0.187	0.098	0.000	0.155	0.099	0.000	0.173
RB5	1000	0.248	0.250	0.217	0.252	0.250	0.233	0.278	0.500	0.240
RB6	1000	0.184	0.150	0.194	0.132	0.000	0.176	0.130	0.000	0.191

7.4.2. Attribute match percentages

Figure 4 presents the attribute match percentages, that is, the percentage of mentees allocated a mentor who matched based on gender, subject, programme and type. The figure shows the results for 1000 cohorts of mentees (randomly generated from mentee scenario 1) and mentors (randomly generated from mentor scenario 1) for small-, medium- and large-sized mentoring programmes with equal numbers of mentees and mentors. As expected, the results are directly related to the criteria used in each algorithm. For example, LP5 and LP6 do not include gender as a criterion, therefore, on average only 60% of mentees received a mentor of their preferred gender under this algorithm. Under LP6, which uses programme as the only criterion, there was a high percentage (approximately 75%) of mentees that were assigned a mentor with a relevant programme. Overall, across all algorithms, apart from LP5 and LP6, approximately 80 − 90% of mentees received a mentor of their preferred gender, and apart from LP4, LP6 and RB4, approximately 80 − 90% of mentees received a mentor of the required type. We also found that the percentage of mentees receiving a mentor with a matching programme and subject is highly dependent on the algorithm.

Across the three programme sizes, similar patterns are observed. However, as was the case when considering the distribution of the mean match score, the percentages of mentees receiving a mentor with the required attributes (such as those shown in Figure 4) have less variability as the size of the mentoring programme increases.

The trends observed in Figure 4 for mentee scenario 1 and mentor scenario 1 are similar to those for other scenarios, with one notable exception when mentor scenario 1 is paired with mentee scenario 2 and scenario 3. In these cases, the proportion of male and female mentors is on average equal (50%), but the proportion of mentees requesting a female mentor is high (80%). In these scenarios, the proportion of mentees receiving a mentor of their preferred gender is approximately 60 − 70%. This highlights an important point; the performance of the algorithm will always be affected by the characteristics of the mentees and mentors in a particular cohort of a mentoring programme.

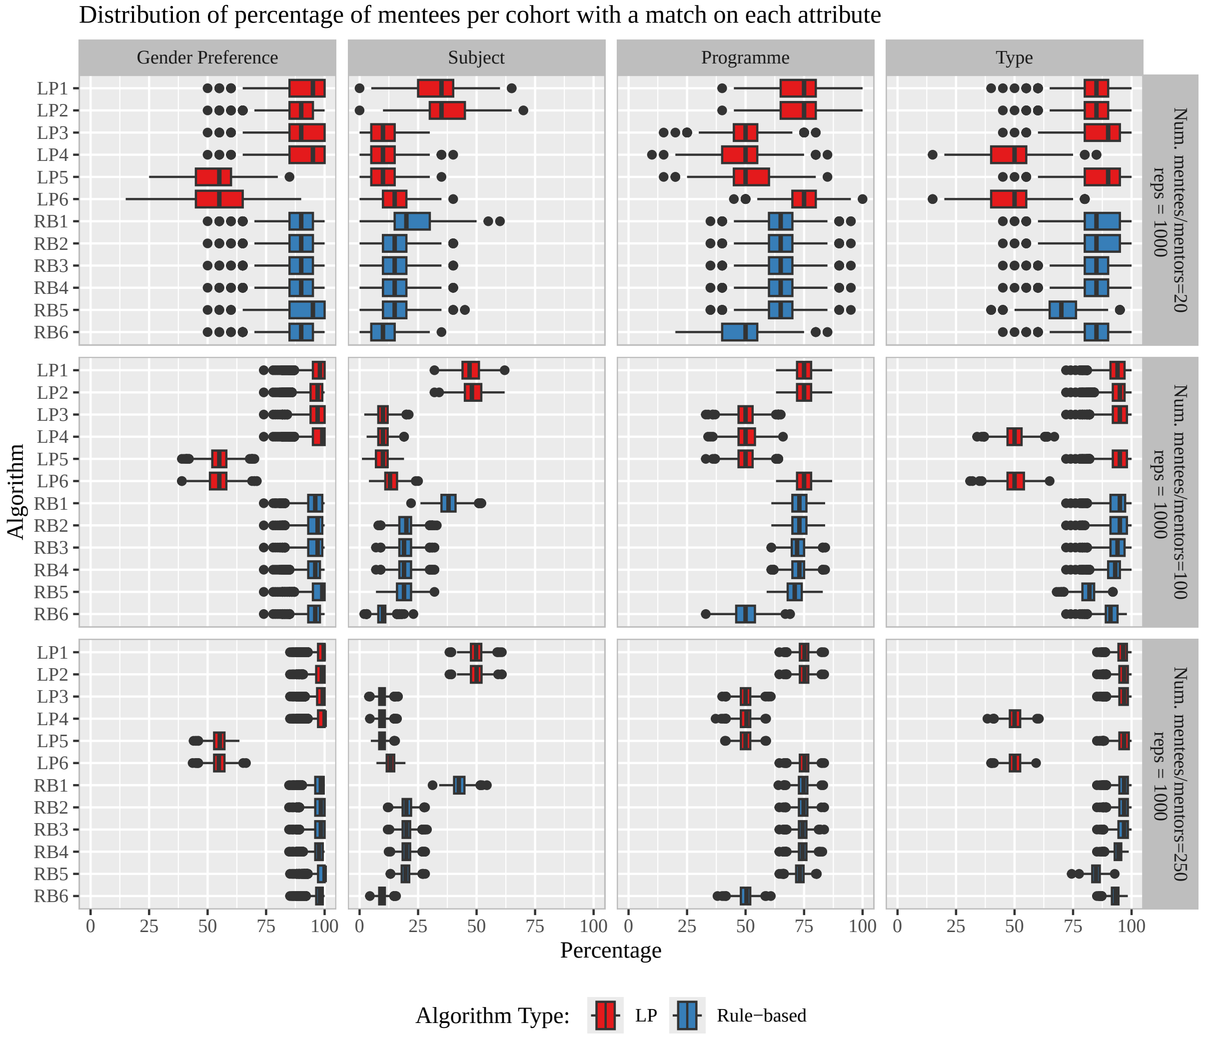

Figure 4: Distribution of percentage of mentees per cohort with a match on each attribute for 1000 cohorts of mentees (randomly generated from mentee scenario 1) and mentors (randomly generated from mentor scenario 1).

7.5. Computation time benchmarking

To investigate the computational efficiency associated with the algorithms, the computation time was recorded for each algorithm for the small-, medium- and large-sized mentoring programmes. In addition to providing a way to compare the LP models and the rule-based heuristics, it also provides a useful comparison with the time required for the current manual matching process.

Mentee scenario 1 and mentor scenario 1 were used to randomly generate 1000 cohorts with equal numbers of mentees and mentors. The computation time required to run the algorithms on a Nectar (`nectar.org.au`) virtual machine with 16 GB RAM and 8 VCPUs is presented in Figure 5. Note that the computation time is presented using a log scale. For a small mentoring programme with 20 mentees and mentors, the computation time of the LP and rule-based heuristics are similar (less than 1 second), with the rule-based heuristics with more iterations having a slightly longer computation time (i.e. time for RB1 > RB6). As the size of the mentoring programme increases, the computation time for the rule-based heuristics does not change much, with most replications completed in under 1 second. However, the LP models take significantly longer as the size of the mentoring programme increases, with cohorts from a medium-sized mentoring programme taking approximately 2–5 seconds and cohorts from a large-mentoring programme taking approximately 30–120 seconds. This is not surprising, given that the LP methods involve scoring every mentee with every mentor, and the numbers of LP variables and constraints increase with the number of mentees/mentors, whereas the rule-based heuristics involve grouping and sorting.

In reality, if a mentorship programme were to run for multiple weeks or months, waiting 2 minutes to obtain an optimal solution seems acceptable. However, if the algorithm were to be run more frequently, the computation time may need to be considered. Computation time is also a consideration if the methods are to be applied to larger programmes. It is expected that the time required for the LP methods will continue to increase significantly with an increase in the size of the programme. With the programme sizes tested here, both the approaches offer significant time savings over the current manual matching process.

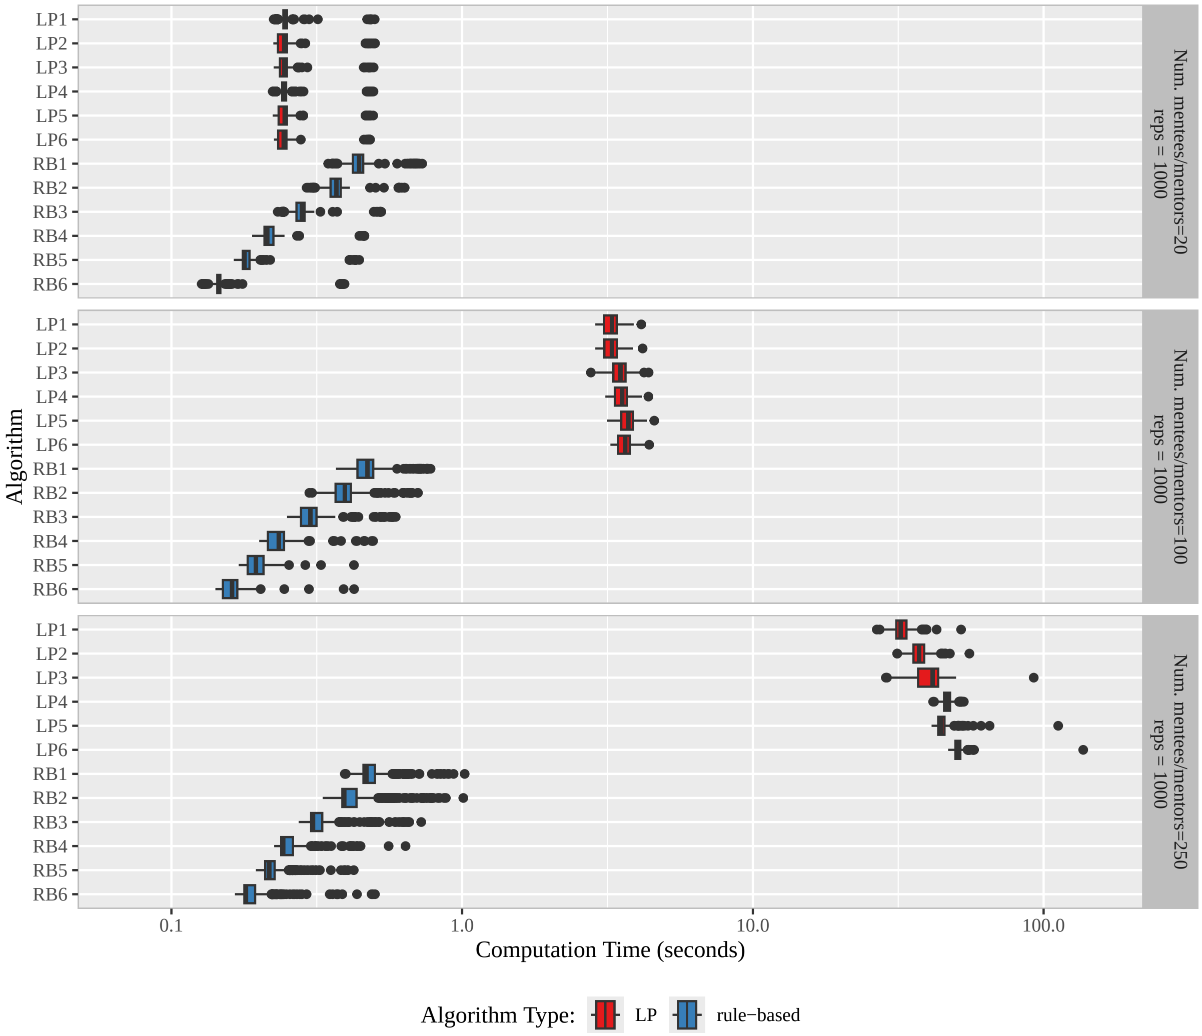

Figure 5: Distribution of computation time to run algorithms for 1000 cohorts of mentees (randomly generated from mentee scenario 1) and mentors (randomly generated from mentor scenario 1). Times in seconds are presented on a log scale.

8. Case study: AUT Women in Tech mentorship programme

The algorithms described above were applied to data from the AUT Women in Tech mentorship programme from semester 2, 2021. In this mentoring programme there were 31 mentees and 32 mentors. To demonstrate the composition of the participants of the mentoring programme, the proportions of mentees and mentors with each attribute are presented in Figure 6. As noted above, the performance of the algorithms will be affected by the composition of the cohort. In this real-life example, slightly over 50% of the mentees prefer a female mentor, with the remainder having no preference. Approximately 75% of the mentors are female. Approximately one quarter of the mentees are first-year students and approximately half of the mentors are academics. Most mentees and mentors are in the field of computer and information sciences, followed by engineering. The mentees and mentors are spread across the 13 subject areas, with some not specifying a subject (these individuals are presented as NA in Figure 6).

Table 17 presents the results from the LP models and rule-based heuristics demonstrating that six of the approaches (LP1, LP2, RB1, RB2, RB3 and RB4) outperformed the original manual matching, in the mean score. Five of the approaches (LP1, LP2, RB2, RB3 and RB4) outperformed the original manual matching on all metrics. Approaches LP1 and LP2 outperformed the original manual matching on all the performances measures, and RB1 and RB2 outperformed or equalled the original manual matching on all except one of the performance measure. Even RB4, which included only four iterations (compared with the 11 iterations of RB1), outperformed or equalled the original manual matching on all performance measures. This suggests that the quality of the mentee/mentor matching could be improved by using the algorithms introduced in this paper.

It is important to note that in addition to the attributes used in the algorithms described in this paper, there may be other factors that the organisers of the mentoring team considered when assigning matches between mentees and mentors. For example, some mentors may have good reputations or important relationships with the university and are therefore given priority, if there are more mentors than mentees. Other information such as the similarities of subjects in which the mentee and mentor may engage may have been considered; for example, although a mentor did not report having expertise in the mentee's subject, the organisers may deem that the mentor's subject area is sufficiently close for a match.

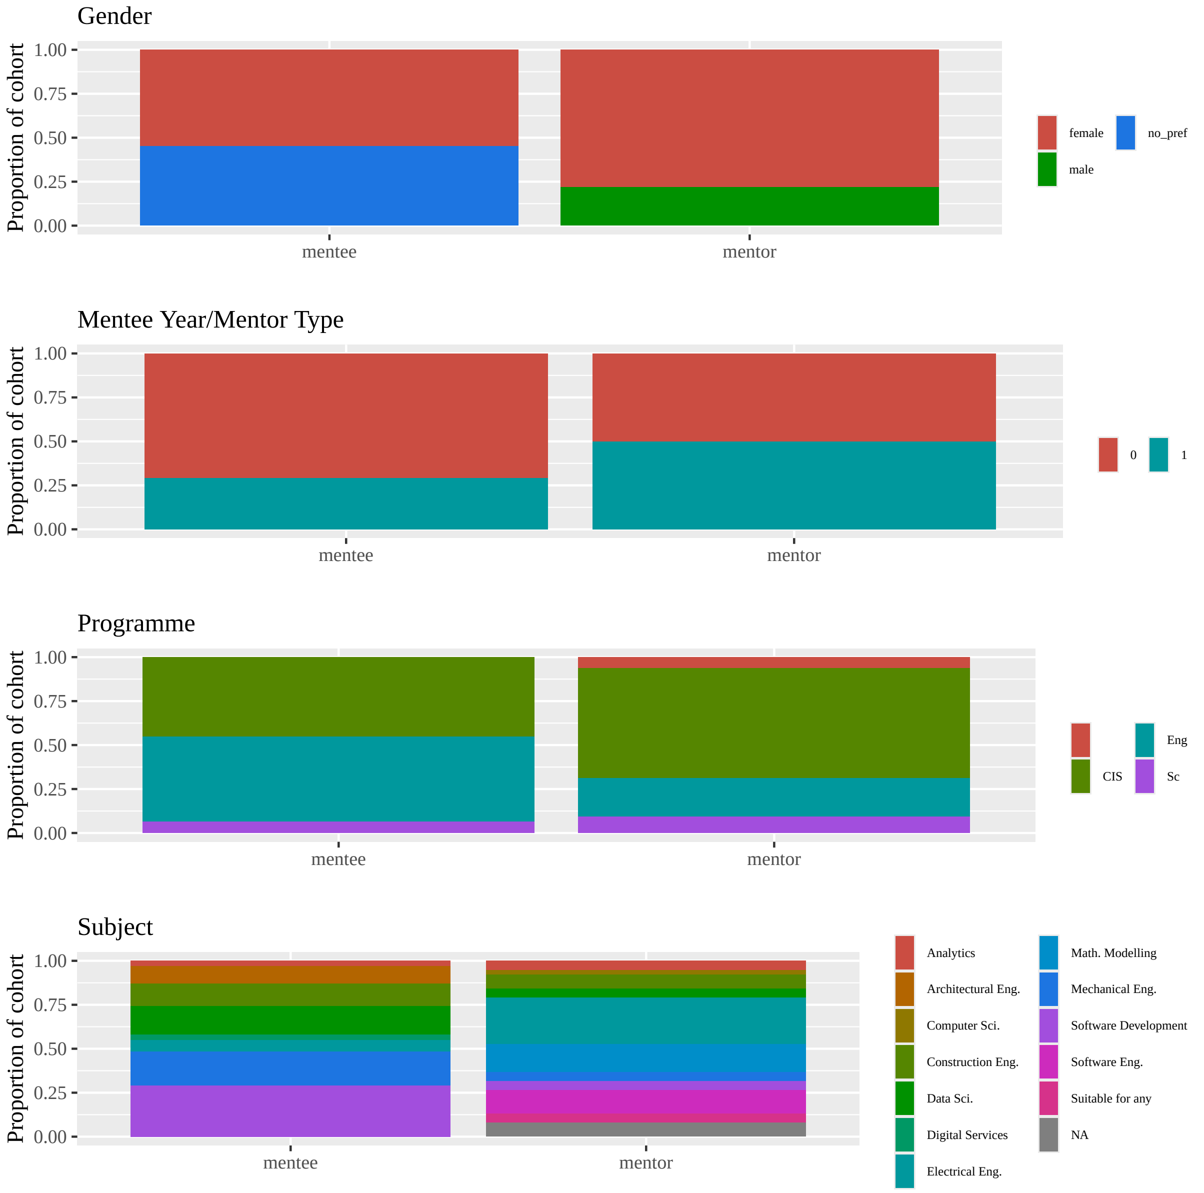

Figure 6: Proportions of each attribute value for mentees and mentors for the semester 2, 2021, AUT Women in Tech mentorship programme.

Table 17: Summary results for the semester 2, 2021 AUT Women in Tech mentorship programme.

					% of mentees with matching mentor			
Algorithm	Total score $R^{\mathbf{w}}$	Obj val Z	Mean score $\bar{R}^{\mathbf{w}}$	Min score $R^{\mathbf{w}}_{\text{min}}$	Gender pref.	Type	Prog.	Subj.
LP1 **	27.45 *	27.45	0.89 *	0.65 *	100.00 *	80.65 *	96.77 *	38.71 *
LP2 **	27.45 *	24.50	0.89 *	0.65 *	100.00 *	80.65 *	96.77 *	38.71 *
LP3	25.40	28.00	0.82	0.50 *	100.00 *	80.65 *	74.19	6.45
LP4	22.25	31.00	0.72	0.50 *	100.00 *	41.94	70.97	6.45
LP5	23.40	25.00	0.75	0.00	87.10	80.65 *	74.19	6.45
LP6	22.15	30.00	0.71	0.15	83.87	54.84	96.77 *	12.90 *
RB1	26.75 *	-	0.86 *	0.50 *	100.00 *	80.65 *	90.32	25.81 *
RB2 **	26.70 *	-	0.86 *	0.65 *	100.00 *	80.65 *	93.55 *	19.35 *
RB3 **	26.70 *	-	0.86 *	0.65 *	100.00 *	80.65 *	93.55 *	19.35 *
RB4 **	26.70 *	-	0.86 *	0.65 *	100.00 *	80.65 *	93.55 *	19.35 *
RB5	24.20	-	0.78	0.50 *	100.00 *	54.84	87.10	12.90 *
RB6	25.10	-	0.81	0.50 *	100.00 *	74.19 *	74.19	12.90 *
Original	25.50	-	0.82	0.50	100.00	67.74	93.55	12.90

* indicates algorithm outperformed or equalled the original manual matching on the given performance measure.

** indicates algorithm outperformed or equalled the original manual matching on all performance measures.

Obj val = Objective function value; Gender pref. = Gender preference; Prog. = Programme; Subj. = Subject

9. Conclusion

This paper has presented two approaches (an LP model and a rule-based heuristic) that can be used for matching mentees with the most appropriate mentors. The approaches have been described with reference to a mentoring programme for female students studying STEM subjects at university level; however, the approaches could easily be extended or altered to suit a different application.

Numerical experiments were used to investigate multiple versions of each approach. The two best-performing LP approaches (LP1 and LP2) outperformed all other approaches in the numerical experiments and the application to real data. The two best-performing rule-based heuristics (RB1 and RB2) performed only slightly worse than the best LP approaches. RB4, which is a simple four-iteration version of the rule-based heuristic, still performed well in the numerical experiments and in the case study. However, the LP methods had a clear advantage in minimising the level of badness of the worst matches. As discussed, this is not surprising because the rule-based heuristics have a greedy approach, assigning the best matches first, and then assigning whoever is left at the end. In contrast, the LP models take a more holistic view, in that they include the match scores from all pairs in the objective function.

The computation time required to perform the rule-based heuristics was quite consistent as the programme sizes increased; however, the computation time increased dramatically when using the LP approach. While the computation time was under 2 minutes even for the largest problem size when using the LP approach (which is reasonable if the programme for which matching is required is running for multiple weeks or months), if the programme size and thus computation time increased, it could be a consideration against using the LP approach. In such circumstances, specialised algorithms, such as the Hungarian method (Kuhn, 1955), could be considered. Anecdotal evidence from organisers of the AUT Women in Tech mentorship programme was that the matching of mentees and mentors took several days to perform manually, so vast time savings could be achieved by applying the methods proposed in this paper

The key contribution of this paper is the development of a rule-based heuristic to match mentees and mentors. The rule-based heuristic provides a simple and easily understandable method for assigning mentees to mentors. As demonstrated through the numerical experiments, the rule-based heuristics can perform nearly as well as the LP models, and for larger programme sizes, require less computation time.

A limitation of the approaches tested in this study is that they included a limited number of attributes: gender preference/gender; type of mentee/mentor; programme and subject. With appropriate data collected in the expression of interest form, this could easily be extended to other attributes provided they can be 'coded' to fit the scoring functions described above. Incorporating factors that do not fit the current scoring functions (e.g. prioritising mentors with important relationships with the university) is an area for future work. Currently, mentees and mentors are asked to write several sentences about their reasons for participating in the programme. Future work could extract relevant information from these text fields and incorporate this within the algorithms.

Another avenue for future work is to examine the factors that do lead to a good match. In this study, the attributes were chosen in consultation with the organisers of the AUT Women in Tech mentorship programme; however, it would be valuable for future research to determine whether these attributes do indeed lead to the best matches. For example, factors such as the mentees' and mentors' personality types could potentially be included.

Motivated by current practices in the AUT Women in Tech mentorship programme, this study prioritised gender preference as the main criterion for matching mentors and mentees. However, it is possible that some mentees may value a mentor's professional background or specific area of expertise more than their gender. Future work could explore how important various attributes are in a mentoring relationship, both in general and to individual mentees. The expression of interest form could be expanded to provide a more comprehensive preference elicitation process, allowing mentor matching to better reflect each mentee's unique career goals and personal aspirations, rather than focusing primarily on gender or year of study.

Another limitation of this study is that the scoring focused on the benefits to mentees. There are also well-documented benefits to mentors of participating in a mentoring scheme. Future work could modify the scoring method to incorporate the benefit to mentors as well.

This study has demonstrated the effectiveness of a rule-based heuristic and provides a useful baseline for future work in the area.

Acknowledgments

The authors wish to acknowledge the Centre for eResearch at the University of Auckland for their help in facilitating this research http://www.eresearch.auckland.ac.nz.

Funding

This work was supported by an AUT ECMS Women in Tech Research Grant.

Ethics

This research was approved by the Auckland University of Technology Ethics Committee, 21/364.

References

Ágoston, K.C., Biró, P., Kováts, E., Jankó, Z., 2022. College admissions with ties and common quotas: Integer programming approach. European Journal of Operational Research 299, 722–734. doi:10.1016/j.ejor.2021.08.033.

Ahmed, T., Johnson, J., Latif, Z., Kennedy, N., Javier, D., Stinson, K., Vishwanatha, J.K., 2021. MyNRMN: A national mentoring and networking platform to enhance connectivity and diversity in the biomedical sciences. FASEB bioAdvances 3, 497–509. doi:10.1096/fba.2020-00102.

Amelink, C.T., 2009. Overview: Mentoring and Women in Engineering, in: Bogue, B., Cady, E. (Eds.), Apply Research to Practice (ARP) Resources. Society of Women Engineers Assessing Women & Men in Engineering (SWE AWE), p. 15.

Avecilla, P.A., Capiña, X.E.R., Javier, A.Y., 2023. Teachers' Teaching Style as Perceived by Students and its Influence on Students' Level of Self-Regulation and Motivation in Learning Psychology. Technium Social Sciences Journal 43, 213–240. doi:10.47577/tssj.v43i1.8693.

Barrett, J.L., Mazerolle, S.M., Nottingham, S.L., 2017. Attributes of Effective Mentoring Relationships for Novice Faculty Members: Perspectives of Mentors and Mentees. Athletic Training Education Journal 12, 152–162. doi:10.4085/1202152.

Berkelaar, M., et al., 2023. lpSolve: Interface to 'Lp_solve' v. 5.5 to Solve Linear/Integer Programs.

Bielczyk, N., Veldsman, M., Ando, A., Caldinelli, C., Makary, M.M., Nikolaidis, A., Scelsi, M.A., Stefan, M., OHBM Student and Postdoc Special Interest Group, Badhwar, A., 2019. Establishing online mentorship for early career researchers: Lessons from the Organization for Human Brain Mapping International Mentoring Programme. European Journal of Neuroscience 49, 1069–1076. doi:10.1111/ejn.14320.

Biró, P., Gyetvai, M., 2022. Online voluntary mentoring: Optimising the assignment of students and mentors. European Journal of Operational Research doi:10.1016/j.ejor.2022.08.008.

Bjursell, C., Florin Sädbom, R., 2018. Mentorship programs in the manufacturing industry. European Journal of Training and Development 42, 455–469. doi:10.1108/EJTD-05-2018-0044.

Bobbio, F., Carvalho, M., Lodi, A., Rios, I., Torrico, A., 2023. Capacity Planning in Stable Matching: An Application to School Choice, in: Proceedings of the 24th ACM Conference on Economics and Computation, Association for Computing Machinery, New York, NY, USA. p. 295. doi:10.1145/3580507.3597771.

Bokova, I., 2017. Cracking the code: Girls' and women's education in science, technology, engineering and mathematics (STEM) - UNESCO Digital Library. https://unesdoc.unesco.org/ark:/48223/pf0000253479.

Burkard, R.E., Dell'Amico, M., Martello, S., 2009. Assignment Problems. Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor 6, Philadelphia, PA 19104).

Cao, N.V., Fragnière, E., Gauthier, J.A., Sapin, M., Widmer, E.D., 2010. Optimizing the marriage market: An application of the linear assignment model. European Journal of Operational Research 202, 547–553. doi:10.1016/j.ejor.2009.06.009.

Cross, M., Lee, S., Bridgman, H., Thapa, D.K., Cleary, M., Kornhaber, R., 2019. Benefits, barriers and enablers of mentoring female health academics: An integrative review. PloS one 14, e0215319.

Deng, C., Gulseren, D.B., Turner, N., 2022. How to match mentors and protégés for successful mentorship programs: A review of the evidence and recommendations for practitioners. Leadership & Organization Development Journal 43, 386–403. doi:10.1108/LODJ-01-2021-0032.

Dutta, I., Shah, A., Safa, M.M., Jayavel, K., 2022. Building a recommendation system for career advice for students from professionals, in: 2022 International Conference on Computer Communication and Informatics (ICCCI), IEEE. pp. 1–10.

Fenoaltea, E.M., Baybusinov, I.B., Zhao, J., Zhou, L., Zhang, Y.C., 2021. The Stable Marriage Problem: An interdisciplinary review from the physicist's perspective. Physics Reports 917, 1–79. doi:10.1016/j.physrep.2021.03.001.

Gale, D., Shapley, L.S., 1962. College Admissions and the Stability of Marriage. The American Mathematical Monthly 69, 9–15. doi:10.2307/2312726, arXiv:2312726.

Gale, D., Shapley, L.S., 2013. College Admissions and the Stability of Marriage. The American Mathematical Monthly 120, 386–391. doi:10.4169/amer.math.monthly.120.05.386.

Garg, N., Kavitha, T., Kumar, A., Mehlhorn, K., Mestre, J., 2010. Assigning Papers to Referees. Algorithmica 58, 119–136. doi:10.1007/s00453-009-9386-0.

Haas, C., Hall, M., 2019. Two-Sided Matching for mentor-mentee allocations—Algorithms and manipulation strategies. PloS one 14, e0213323.

Hagler, M.A., Rhodes, J.E., 2018. The Long-Term Impact of Natural Mentoring Relationships: A Counterfactual Analysis. American Journal of Community Psychology 62, 175–188. doi:10.1002/ajcp.12265.

Hee, J., Toh, Y.L., Yap, H.W., Toh, Y.P., Kanesvaran, R., Mason, S., Krishna, L.K.R., 2020. The Development and Design of a Framework to Match Mentees and Mentors Through a Systematic Review andThematic Analysis of Mentoring Programs Between 2000 and 2015. Mentoring & Tutoring: Partnership in Learning 28, 340–364. doi:10.1080/13611267.2020.1778836.

Hoenen, S., Kolympiris, C., 2020. The Value of Insiders as Mentors: Evidence from the Effects of NSF Rotators on Early-Career Scientists. The Review of Economics and Statistics 102, 852–866. doi:10.1162/rest_a_00859.

Kricorian, K., Seu, M., Lopez, D., Ureta, E., Equils, O., 2020. Factors influencing participation of underrepresented students in STEM fields: Matched mentors and mindsets. International Journal of STEM Education 7, 1–9.

Kuhn, H.W., 1955. The Hungarian method for the assignment problem. Naval Research Logistics Quarterly 2, 83–97. doi:10.1002/nav.3800020109.

LEMON Graph Library, 2018. Welcome to COIN-OR::LEMON. https://lemon.cs.elte.hu/trac/lemon.

Lovász, L., Plummer, M.D., 1986. Matching Theory. North-Holland Mathematics Studies, 121, North-Holland, Amsterdam.

Martinez, A., Christnacht, C., 2021. Women are nearly half of US workforce but only 27% of STEM workers. https://www.census.gov/library/stories/2021/01/women-making-gains-in-stem-occupations-but-still-underrepresented.html.

McClain, C.M., Kelner, W.C., Elledge, L.C., 2021. Youth mentoring relationships and college social and academic functioning: The role of mentoring relationship quality, duration, and type. American Journal of Community Psychology 68, 340–357.

Moghaddam, M., 2016. Best Matching Theory & Applications. Automation, Collaboration, & E-Services ; v. 3, Springer, Cham.

Onuma, F.J., Berhane, B., Fries-Britt, S.L., 2022. "I've always been in private school": The role of familial norms and supports in Black immigrant students' preparation for STEM majors. Journal of Diversity in Higher Education 15, 241–253. doi:10.1037/dhe0000285.

O'Rourke, B., 2021. Growing gap in STEM supply and demand. https://news.harvard.edu/gazette/story/2021/11/increasing-access-and-opportunity-in-stem-crucial-say-experts/.

Patterson, L., Varadarajan, D.S., Salim, B.S., 2021. Women in STEM/SET: Gender gap research review of the United Arab Emirates (UAE) – a meta-analysis. Gender in Management 36, 881–911. doi:10.1108/GM-11-2019-0201.

Pettersson, W., Delorme, M., García, S., Gondzio, J., Kalcsics, J., Manlove, D., 2021. Improving solution times for stable matching problems through preprocessing. Computers & Operations Research 128, 105128. doi:10.1016/j.cor.2020.105128.

Poulsen, K.M., 2013. Mentoring programmes: Learning opportunities for mentees, for mentors, for organisations and for society. Industrial and Commercial Training 45, 255–263. doi:10.1108/ICT-03-2013-0016.

R Core Team, 2023. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing.

Rehman, R., Khan, F., Kayani, N., Ali, T.S., 2022 Jul-Aug. Reflection of mentors and mentees at initiation of Faculty Mentorship Program at Aga Khan University: A perspective. Pakistan Journal of Medical Sciences 38, 1691. doi:10.12669/pjms.38.6.5454.

Sassler, S., Glass, J., Levitte, Y., Michelmore, K.M., 2017. The missing women in STEM? Assessing gender differentials in the factors associated with transition to first jobs. Social Science Research 63, 192–208. doi:10.1016/j.ssresearch.2016.09.014.

Weathers, D., 2011. Forbes Insights Study Identifies Strong Link between Diverse Talent and Innovation | Business Wire. https://www.businesswire.com/news/home/20110713006204/en/Forbes-Insights-Study-Identifies-Strong-Link-Diverse.

Zavala, A., Esquivel, A., Gutierrez, M.J., Plascencia, G.L., García, O., García-Arellano, G., 2022. Women Empowering Women: A mentoring program, in: 2022 IEEE Global Engineering Education Conference (EDUCON), pp. 480–484. doi:10.1109/EDUCON52537.2022.9766626.