



The ethical shortlisting problem

Güneş Erdoğan *

School of Management, University of Bath, BA1 7AY, UK

Department of Business Administration, MEF University, Ayazaga Cad. No:4 Maslak, 34396, Sariyer, Istanbul, Turkey

ARTICLE INFO

Keywords:

Computational ethics
Hiring
Shortlisting
Linear programming

ABSTRACT

Hiring is a fundamental, frequent activity for all organizations. Hiring decisions have been reported to be subject to conscious and unconscious biases in the literature. The field of Computational Ethics aims to quantify and maximize the ethicality of decisions. This paper attempts to apply Computational Ethics to the shortlisting process in hiring through the use of Linear Programming. Given a set of applicants for a job with numerical qualification values, the author aims to determine weights for each qualification type to compute scores and resulting rankings for each applicant. To this end, Abstract Moral Theories of Utilitarianism, Maximin/Leximin, Egalitarianism, and Prioritarianism are utilized and applied to a set of randomly generated applicant data. Computational experiments demonstrate that the models are scalable and return interpretable results. The necessity of a quota-based shortlisting system to alleviate disadvantaged candidates is highlighted. The author recommends the use of the Maximin model and iteratively eliminating the applicant with the lowest score.

1. Introduction

Hiring is a recurring task for every organization, may it be a commercial company, governmental body, or a charitable foundation. In the UK, for example, [Office for National Statistics \(2019\)](#) shows an average turnover rate of 28% between January 2017 and December 2018. Consequently, employment opportunities are regularly being advertised, applications submitted, longlisting and then shortlisting completed, shortlisted applicants interviewed, and decisions are made. Clearly, there are other variations, including the use of so-called “head-hunters”, aptitude tests, recruitment centers, etc. but the core processes of longlisting and shortlisting are common to nearly all recruitment.

However methodically undertaken, the hiring process is still subject to human cognition and consequently, its hidden and not-so-hidden biases. For example, [Ruffle and Shtudiner \(2015\)](#) analyzed the effect of having an attractive-looking profile picture on a CV and found significant differences in terms of callbacks from employers. [Fernandez-Mateo and Fernandez \(2016\)](#) studied the reasons for gender inequality in top management jobs and stated that gender differences affect the process mainly while the candidate pool for interviews is being formed, which we interpret to be the shortlisting process.

One solution to the challenges of bias can be found in the burgeoning field of Human Resources (HR) Analytics (see [Marler and Boudreau \(2017\)](#) and [Margherita \(2021\)](#) for recent surveys). [Jimenez et al. \(2018\)](#), for example, studied the shortlisting process, converting the qualification of each applicant to a ranking and then using fuzzy

logic to compute an overall objective and a fair ranking. The book by [Edwards and Edwards \(2019\)](#) features a chapter on the use of predictive analytics for hiring and selection, where the authors state that male applicants are 3.3 times more likely to be shortlisted with respect to the females. [Pessach et al. \(2020\)](#) provide machine learning algorithms and mathematical programming formulations to optimize the hiring process based on the objectives of maximally satisfying workforce demand, match between role requirements and applicant profiles, and diversity of the workforce. Yet biases can produce discriminatory outcomes into the algorithmic approaches and their outcomes, as pointed out by [Lambrech and Tucker \(2019\)](#). Consequently, in this paper, we address this issue by combining HR Analytics with insights from the relatively new field of Computational Ethics.

Ethics is an inherently qualitative concept, since there is no exact way of measuring how ethical an assessment, decision, or action is. Computational Ethics is a relatively new field of research that aims to answer these questions. In their groundbreaking work, [Anderson et al. \(2006\)](#) stated the importance of embedding ethics into Artificial Intelligence, and presented their integration of *Hedonic Act Utilitarianism* into a knowledge-based decision support system that provides guidance on biomedical ethical dilemmas. In the same issue, [Allen et al. \(2006\)](#) underlined the importance of explicitly involving ethics into the design of new technologies under development. We refer the interested reader to the book of [Anderson and Anderson \(2011\)](#), the survey of [Cervantes et al. \(2020\)](#), and to the recent paper of [Segun \(2021\)](#) for an in-depth discussion of the topic.

* Correspondence to: School of Management, University of Bath, BA1 7AY, UK.
E-mail address: g.erdogan@bath.ac.uk.

In this paper, we aim to apply existing Abstract Moral Theories to a critical part of hiring process, shortlisting. The mathematical tool we use is Linear Programming (LP), which aligns surprisingly well with our aim. The ethical theories we employ are based on the excellent book by Tännsjö (2019), in which the author presented four main Abstract Moral Theories: *Utilitarianism*, *Maximin/Leximin*, *Egalitarianism*, and *Prioritarianism* as well as their implications for the practice of healthcare. The interested reader is referred to the aforementioned work and the references therein for the details of the literature on Abstract Moral Theories and their applications.

It is worth noting that our approach differs significantly from the Social Choice Theoretical perspective that aims to combine the choices and preferences of a set of individuals. In our problem setting there are no voters, but a single, automated decision maker that aims to make a quantitatively ethical decision. A related but distinctly different method is Data Envelopment Analysis (DEA), through which the applicants may be assessed based on their maximal efficiency. For a recent example of DEA and its application, we refer the interested reader to Zhu et al. (2021). However, the resources used by an applicant are rarely included in an application, and judging every applicant by a different set of criteria is against our understanding of ethics.

Our primary contribution is the explicit incorporation of Abstract Moral Theories into optimization models, which has not been achieved before, to the best of our knowledge. Secondly, we demonstrate the necessity of a quota-based shortlisting system to alleviate disadvantaged candidates. Finally, we provide an iterative algorithm that eliminates the candidate with the lowest score, which we believe to be the maximally ethical method of shortlisting.

The rest of this paper is organized as follows. In Section 2, we provide a motivating example and the details of what we aim to achieve. In Section 3, we present the models that result from applying the Abstract Moral Theories to the shortlisting process. In Section 4, we demonstrate the use of the models on a set of data and derive insights. In Section 5, we provide the results of an iterative algorithm that eliminates the applicant with the lowest score, which we observe to return consistent results for some models. Finally, in Section 6, we provide our conclusions.

2. Motivation

Consider the case when an academic department aims to recruit at the Associate Professor/Senior Lecturer level. Table 1 contains data about five longlisted applicants, including their number of papers, h-indices, grant incomes, number of co-authors, and average teaching evaluations in the past three years. The data has been generated randomly and then modified to ensure non-dominance among the applicants, i.e. each applicant is strictly better than the others in one qualification type. Note that having a higher value is better for each qualification type listed. We realize that a job application may involve a much larger amount of applicants and types of qualifications, and the norms are different for each academic field of research, but we believe that Table 1 is sufficient for the purpose of demonstrating our motivation.

We aim to shortlist two of the applicants based on this quantitative data, so as to avoid hidden biases and an unethical assessment. We thereby need a way of converting the qualification data of each applicant to a score. We postulate that an ethical scoring method should satisfy the following necessary conditions.

1. If two applicants have the exact same qualifications, they should have the same score.
2. If an applicant has qualifications that are greater than or equal to those of another applicant for each type of qualification, then the score of the former applicant must be greater than or equal to that of the latter.

Mathematically, the first condition would translate to the scoring method being a function, whereas the second would point to the function being monotonic (i.e. nondecreasing in multiple variables). Since a linear function with nonnegative coefficients satisfies both conditions, we propose to score the applicants by multiplying each of their quantitative qualifications by an appropriate weight. To equalize the qualification types that have different scales, we find it appropriate to normalize all qualification values to the interval $[0, 1]$ by dividing them with the maximum value of the qualification among the applicants.

Finally, the recruiting department may want to provide lower and upper bounds on the weight to be assigned to each qualification type, as per the requirements of the role as well as to avoid an unbalanced scoring scheme. For example, a research role may require a greater weight to be placed on the first three qualification types. On the other hand, a networking role may require more emphasis on the number of co-authors, whereas a teaching role more may require a greater weight on the average teaching evaluations. These lower and upper bounds may be viewed as an ethical commitment to find an applicant that fulfills the requirements of the job.

The final piece of the model is the objective function of maximizing ethicality, which will be provided by the Abstract Moral Theories. In the next section, we provide LP models that can address these requirements.

3. Models

We define I as the set of applicants and J as the set of qualification types. We denote the level of each applicant $i \in I$ for qualification type $j \in J$ as q_{ij} . We normalize the qualifications to the range $[0, 1]$ as $\hat{q}_{ij} = q_{ij} / \max_{i' \in I} \{q_{i'j}\} \forall i \in I, j \in J$.

Our aim is to assign weights to the types of qualifications to maximize the ethical objective adopted. Based on the recruiting organization's objectives and the specification of the role, each type of qualification has a lower bound $0 \leq l_j \leq 1$ on the weight that it can be assigned. Note that the condition $\sum_{j \in J} l_j \leq 1$ is required to avoid infeasibility. The recruiting organization also imposes an upper bound $l_j \leq u_j \leq 1$ on the weights to ensure a balanced scoring scheme, with the feasibility condition $\sum_{j \in J} u_j \geq 1$.

Let x_j be the decision variable that corresponds to the weight of qualification type j . Furthermore, let y and z denote the minimum and the maximum scores attained by the applicants, respectively. In the rest of this section, we provide models for our problem based on the perspectives of Utilitarianism, Maximin/Leximin, Egalitarianism, and Prioritarianism.

3.1. Utilitarian model

Tännsjö (2019) defines Utilitarianism as "the idea that we ought to maximize the sum total of happiness". In the context of our problem, we correlate the happiness of each applicant with the score they are

Table 1
Illustrative example — qualifications for a set of applicants.

Applicant	Number of papers	h-index	Grant income (£)	Number of co-authors	Average teaching evaluation
1	12	9	10872	2	3.86
2	10	6	85548	7	3.59
3	16	3	66968	3	4.92
4	18	3	1867	7	4.10
5	10	8	40790	9	2.80

assigned. Hence we set the objective of our first model as maximizing the sum of all scores assigned to the applicants. The model is then:

(U)

$$\text{maximize} \quad \sum_{i \in I} \sum_{j \in J} \hat{q}_{ij} x_j \quad (1)$$

$$\text{subject to} \quad \sum_{j \in J} x_j = 1 \quad (2)$$

$$l_j \leq x_j \leq u_j \quad \forall j \in J. \quad (3)$$

The objective function (1) corresponds to the objective stated above. Constraint (2) sets the sum of the weight assigned to the types of qualifications equal to 1. Finally, constraint set (3) states the upper and lower bounds of each weight. These two constraints will be repeated throughout the rest of the models we will present.

Being an LP model with a single equality constraint, (U) can be solved to optimality without the need of a sophisticated algorithm. Algorithm 1 we provide below finds an optimal solution in $O(|J| \log(|J|))$.

Algorithm 1 Solution algorithm for model (U)

- 1: **Input** Objective function coefficients $\sum_{i \in I} \hat{q}_{ij} \quad \forall j \in J$
 - 2: Initialize $x_j := l_j \quad \forall j \in J$
 - 3: Determine the remaining weight as $r := 1 - \sum_{j \in J} l_j$
 - 4: Sort qualification types in nonincreasing order of their objective function coefficients, with $\pi(k)$ denoting the k^{th} qualification type in the order.
 - 5: $k = 1$
 - 6: **While** $k \leq |J|$ **and** $r > 0$
 - 7: $x_{\pi(k)} := x_{\pi(k)} + \max\{r, u_{\pi(k)} - l_{\pi(k)}\}$
 - 8: $r := r - \max\{r, u_{\pi(k)} - l_{\pi(k)}\}$
 - 9: $k := k + 1$
 - 10: **End While**
 - 11: **Return** $x_j \quad \forall j \in J$
-

It should be noted that each applicant contributes to the objective function as much as their qualification values, and the stronger applicants with higher qualification values contribute more, effectively skewing the gradient of the objective function in their favor. From another perspective, the objective function represents the average of the applicants, and penalizes applicants whose qualification profiles differ significantly from the average. In other words, Utilitarianism favors the strong and penalizes the outliers. Maximin idea that we explore in the next subsection attempts to remedy this.

3.2. Maximin model

Maximin is considered to be a complementary perspective to the Utilitarian idea. In the words of Tännsjö (2019), Maximin idea can be stated as “absolute priority should be given to the person who is worst off”. In the context of the shortlisting process, we interpret this as the maximization of the lowest score among all applicants, and model it by using the variable y we have previously defined. The resulting model is:

(M)

$$\text{maximize} \quad y \quad (4)$$

$$\text{subject to} \quad y \leq \sum_{j \in J} \hat{q}_{ij} x_j \quad \forall i \in I \quad (5)$$

$$y \geq 0 \quad (6)$$

(2), (3).

In model (M), the objective (4) is to maximize y . Constraint set (5) ensures that y is less than or equal to the minimum score among the applicants. Constraint (6) states that y is nonnegative. Constraints (2) and (3) identical to model (U).

Unlike (U), model (M) does not admit closed a form optimal solution in general. However, the number of variables ($|J|$) and constraints ($|I| + 1$) it involves would require less than 0.1 s for any realistic sized instance with any LP solver.

3.3. Leximin model

Leximin is an improvement upon the Maximin idea. To quote (Tännsjö, 2019): “Once the needs of those who are worst off have been catered to, we ought to tend to the needs of those who come next in line”. This definition conveniently matches a lexicographic maximization objective, with y being maximized first, being followed by the sum of all applicant scores. We denote the optimal solution value of (M) as y^* , to become a lower bound for all applicant scores. The formulation for Leximin is then:

(L)

$$\text{maximize} \quad \sum_{i \in I} \sum_{j \in J} \hat{q}_{ij} x_j \quad (7)$$

$$\sum_{j \in J} \hat{q}_{ij} x_j \geq y^* \quad \forall i \in I \quad (8)$$

(2), (3).

Model (L) differs from (M) in terms of its objective function (7) that maximizes the sum of scores, as well as constraint set (8) that enforces the lower bound found by (M) upon the scores. The computational complexity of (L) is equivalent to that of (M).

We provide a short verbal proof of the fact that models (M) and (L) can have different optimal solutions. Consider an instance with a particular applicant k , who has scaled qualification values that are all equal to each other and less than or equal to those of the rest of the candidates, i.e. $q_{kl} = \min_{(i,j) \in (I \setminus \{k\}) \times J} \{q_{ij}\} \quad \forall l \in J$. This candidate determines the objective function value for (M). Furthermore, all feasible solutions for (M) have the same objective function value and are optimal for this instance. However, if the objective function coefficient vector of (L) is not equal to the qualifications of applicant k multiplied by a scalar, then the optimal solution(s) for (L) would be a subset of the feasible solutions of (M).

3.4. Egalitarian model

Tännsjö (2019) defines Egalitarianism as: “...a family of theories all resting on the idea that inequalities are of negative value; an unequal distributive pattern is, in one respect at least, better if it is even rather than if it exhibits different levels of happiness among the recipients”. The inequality, in our case, happens in terms of the scores assigned to the applicants. So, we formulate our model to find weights that will minimize the difference between the maximum and the minimum scores of the applicants. Utilizing the variable z we have previously defined to represent the maximum score among the applicants, the corresponding formulation is:

(E)

$$\text{minimize} \quad z - y \quad (9)$$

$$\text{subject to} \quad z \geq \sum_{j \in J} \hat{q}_{ij} x_j \quad \forall i \in I \quad (10)$$

$$z \geq 0 \quad (11)$$

(2), (3), (5), (6).

The objective function (9) attempts to equalize all scores. Constraint set (10) dictates that z is greater than or equal to the maximum score among all applicants, and constraint (11) ensures that z is nonnegative. The rest of the constraints are identical to model (M). Model (E) has a polynomial level of computational complexity as models (M) and (L).

3.5. Prioritarian models

Prioritarianism can be described as giving priority to those who are disadvantaged. Tännsjö (2019) explains Prioritarianism as: “The rationale behind prioritarianism is the idea that suffering has a special moral importance. This means that a person who momentarily suffers has a special moral claim for improvement of her hedonic situation”. The author also states that Prioritarianism results in a family of theories rather than a single theory, based on how important each type of disadvantage is.

We emphasize that Prioritarianism is in conflict the first necessary condition we have postulated in Section 2, which states that two applicants should be scored equally if they have equal qualifications. The Prioritarian view would be that if one of these two applicants have a disadvantage, they should be assigned a higher score. This is not mathematically possible with a linear weighting function unless the disadvantage is viewed as a qualification, which we believe to be inappropriate. On an individual level, an intuitive method would be to use the information about disadvantaged applicants as a tiebreaker when the scores are being converted to rankings. In what follows, we aim to prioritize the subset of disadvantaged applicants at an aggregate level.

It is hard to define what “disadvantage” is. In the absence of *a priori* knowledge that a subset of applicants are disadvantaged, we can only attempt to infer disadvantaged applicants based on their qualification values. However, the results of such an analysis can also be controversial. For example, if an applicant is significantly worse than others in terms of teaching scores, does that point to a disadvantage or just to a lack of commitment to teaching? Furthermore, let us assume we know *a priori* that a subset of applicants are disadvantaged. How would we compare a permanent physical disability to a temporary disadvantage such as a two-year leave of absence due to caring duties, or even to another permanent disadvantage such as being from a deprived minority?

For the purposes of this paper, we assume that all disadvantages are equal. Let us define $D \subseteq I$ to be the subset of applicants that we know *a priori* to be disadvantaged. In what follows, we model Prioritarianism in terms of its reflection onto models (U), (M), (L), and (E). Let y' be the minimum score among the disadvantaged applicants, and z' be the maximum score among the non-disadvantaged applicants.

3.5.1. Prioritarian/Utilitarian

Based on model (U), our first model for Prioritarianism is:

(P/U)

$$\begin{aligned} \text{maximize} \quad & \sum_{i \in D} \sum_{j \in J} \hat{q}_{ij} x_j & (12) \\ \text{subject to} \quad & (2), (3). \end{aligned}$$

The objective function (12) aims to maximize the sum of the scores for the disadvantaged applicants. Due to the identical constraint structure with model (U), model (P/U) can also be solved using Algorithm 1. Note that this model is equivalent to model (U) if $D = I$.

3.5.2. Prioritarian/Maximin

We model Prioritarianism based on model (M) as:

(P/M)

$$\text{maximize} \quad y' \tag{13}$$

$$\text{subject to} \quad y' \leq \sum_{j \in J} \hat{q}_{ij} x_j \quad \forall i \in D \tag{14}$$

$$y' \geq 0 \tag{15}$$

(2), (3).

For (P/M), the objective is to maximize the minimum score among all disadvantaged applicants. The constraint structure is similar to model (M). We underline that this model is equivalent to (M) if $D = I$.

3.5.3. Prioritarian/Leximin

Let us denote the optimal solution value of model (P/M) as $(y')^*$. The model for Prioritarian/Leximin is then:

(P/L)

$$\text{maximize} \quad \sum_{i \in I} \sum_{j \in J} \hat{q}_{ij} x_j \tag{16}$$

$$\begin{aligned} \text{subject to} \quad & \sum_{j \in J} \hat{q}_{ij} x_j \geq (y')^* \quad \forall i \in D & (17) \\ & (2), (3). \end{aligned}$$

Model (P/L) aims maximize the sum of the scores for all the applicants, subject to the constraint set (17) that sets a lower bound on the minimum score of the disadvantaged applicants. This model is equivalent to (L) if $D = I$.

3.5.4. Prioritarian/Equalitarian

Finally, we incorporate Prioritarianism into model (E) as:

(P/E)

$$\text{minimize} \quad z' - y' \tag{18}$$

$$\text{subject to} \quad z' \geq \sum_{j \in J} \hat{q}_{ij} x_j \quad \forall i \in I \setminus D \tag{19}$$

$$z' \geq 0 \tag{20}$$

(2), (3), (14), (15).

This modified version of (E) aims to maximize the minimum score among the disadvantaged applicants, while simultaneously attempting to minimize the maximum score among the non-disadvantaged applicants, as per (18). The constraint structure of the model is identical to (E).

We conclude this section with the statement that there may be many other ways of modeling Prioritarianism, depending on how to compare different types of disadvantages. This is in line with the viewpoint of Tännsjö (2019) about the family of theories that Prioritarianism gives birth to.

4. Computational experiments

One of the advantages of Mathematics over Abstract Moral Theory is the possibility of conducting numerical experiments instead of thought experiments. In this section, we present our results for the data provided in Section 2. We first examine the case of no applicants with disadvantages, then move over to the case with disadvantaged applicants. We also analyze an instance with a larger number of applicants. All models and instances we have tested have required less than 0.1 CPU second to be solved, with no discernible difference in performance. Our aim is to generate insights rather than test the computational reach, since the LP models we have presented can be easily scaled up to thousands of applicants and tens of qualification types.

The choice of the lower and upper bounds for the weights depends on the preference of the decision maker. As a rule of thumb, we recommend using the formula $l_j = 1/(\sqrt{\rho} \times |J|)$ $\forall j \in J$ and $u_j = \sqrt{\rho}/|J|$ $\forall j \in J$, where ρ is the upper bound of the ratio of the maximum weight to the minimum weight. In our computational experiments we have used $\rho = 4$, meaning that no qualification can have a weight that is more than four times any other, resulting in $l_j = 0.1$ and $u_j = 0.4$ $\forall j \in J$. Clearly, more specific lower and upper bounds may be chosen based on the requirements of the job.

Table 2
Optimal weights for the models without disadvantaged applicants.

Model	Weights				
	Number of papers	h-index	Grant income (£)	Number of co-authors	Average teaching evaluation
(U)	0.3000	0.1000	0.1000	0.1000	0.4000
(M)	0.3247	0.2466	0.1000	0.1000	0.2287
(L)	0.3247	0.2466	0.1000	0.1000	0.2287
(E)	0.3935	0.2740	0.1000	0.1000	0.1326

Table 3
Scores and rankings for the models without disadvantaged applicants.

Applicant	(U)		(M)		(L)		(E)	
	Score	Rank	Score	Rank	Score	Rank	Score	Rank
1	0.6488	4	0.6774	3	0.6774	3	0.6752	4
2	0.7030	3	0.6894	2	0.6894	2	0.6757	3
3	0.8116	1	0.7111	1	0.7111	1	0.6852	1
4	0.7466	2	0.6774	3	0.6774	3	0.6752	4
5	0.6309	5	0.6774	3	0.6774	3	0.6852	1

4.1. Without disadvantaged applicants

We first present the optimal weights for the models (U), (M), (L), and (E). All models have been solved using a Simplex algorithm based solver. The results are provided in Table 2. There is a significant difference between the results for (U) and the rest of the models, due to the lack of concern for the minimum score or the difference between the minimum and maximum scores. Models (M) and (L) return identical results, due to their similar structure. The difference between the results of (M), (L), and (E) seem to be insignificant at first sight. However, the scores of the applicants result in rankings that are significantly different.

Table 3 displays the scores and rankings for the models without disadvantaged applicants. The scores have been rounded to the fourth significant digit, and the rankings are determined accordingly. All models agree that Applicant 3 is ranked first. While models (M) and (L) completely agree due to their close structure, they disagree with the other two models about most of the rankings. For example, for the second place in ranking, (U) suggests Applicant 4, (M) and (L) suggest Applicant 2, and (E) suggests Applicant 5 as in a tie for the first place with Applicant 3. This is reminiscent of individuals with different world views that struggle to agree.

4.2. With disadvantaged applicants

We assume that we know Applicant 1 and Applicant 5 to be disadvantaged, i.e. $D = \{1, 5\}$, and run models (P/U), (P/M), (P/L), and (P/E) on the same set of data. The optimal weights are given in Table 4, whereas the resulting scores and rankings are provided in Table 5. These models favor h-index as the most important qualification, for which the disadvantaged applicants have the highest qualification values.

Analysis of Table 5 reveals a significant overall change in scores and rankings. The differences in scores are more pronounced for these models, due to the prioritization of the disadvantaged applicants. All Prioritarian models successfully lift both disadvantaged applicants to the first two places in the ranking. On the other hand, Applicant 3, who was unanimously ranked first by the previous set of models is now unanimously ranked fourth. This is a stark warning about the effect of prioritization on the non-disadvantaged applicants.

We find a surprising result when we run a second experiment with $D = \{1, 4\}$, which is presented in Tables 6 and 7. In this case, Prioritarian models cannot find weights that can provide an advantage to both disadvantaged applicants. Model (P/U) lifts Applicant 4 to the second rank at the cost of leaving applicant 1 at fourth place. The other three models, with their emphasis on equality, cannot provide either of the disadvantaged applicants with a first or second rank.

Based on these experiments, we observe that Prioritarian models run the risk of providing too much advantage to the disadvantaged applicants when their qualification profiles match, and too little when their qualification profiles do not match. We conjecture that with a larger number of disadvantaged applicants, it is more likely to have a mismatch. We test this conjecture in the next subsection.

4.3. Larger number of applicants

We have generated random qualification data for 15 additional applicants, ensuring that the first five applicants are still not dominated. We have kept applicants 1 and 5 as disadvantaged, to retain their closely aligned profiles. Three of the additional applicants are disadvantaged, resulting in set $D = \{1, 5, 12, 14, 16\}$. The resulting data is presented in Table 8. For this larger instance, we aim to shortlist five applicants.

Table 9 demonstrates that the Prioritarian models indeed change the weights provided by the non-Prioritarian models. However, to observe their true effect we need the resulting rankings, provided in Table 10 with the disadvantaged applicants are shown in boldface.

We can observe from the results in Table 10 that all 8 models agree on applicant 20 being shortlisted. With respect to model (U), model (P/U) successfully lifts Applicant 1 by 11 ranks and Applicant 5 by 10 ranks, and shortlists 2 disadvantaged applicants. Models (P/M) and (P/L) return identical results, but differ from (M) and (L) by removing disadvantaged Applicant 5 from the shortlist range. Model (P/E), in an effort to minimize the score differences, ranks all disadvantaged applicants out of the shortlist range. We observe our conjecture to be true for this larger instance, regarding models (P/M), (P/L), and (P/E). We attribute this to the conflict between the Prioritarian perspective and first necessary condition we have postulated. As a remedy, we propose that a quota for disadvantaged applicants to be established, and a non-Prioritarian model of choice to be used separately on the disadvantaged and non-disadvantaged applicants. The top ranked applicants in both rankings may then be selected to fill the quota and satisfy the total number of applicants to be shortlisted.

5. Eliminating the applicant with the lowest score

It is worth noting that models (M) and (L) are prone to return low quality rankings in the presence of dominated applicants, due to their focus on the minimum score. Consider adding a sixth applicant to the original set of 5 applicants, with half of the qualification values of Applicant 1. This new applicant, who will clearly have the lowest score, sets the gradient for the objective function. Due to the alignment of their qualification profiles with the weakest applicant, Applicant 1

Table 4
Optimal weights for the models with disadvantaged applicants, $D = \{1,5\}$.

Model	Weights				
	Number of papers	h-index	Grant income (£)	Number of co-authors	Average teaching evaluation
(P/U)	0.1000	0.4000	0.1000	0.1000	0.3000
(P/M)	0.1000	0.4000	0.1000	0.1075	0.2925
(P/L)	0.1000	0.4000	0.1000	0.1075	0.2925
(P/E)	0.3000	0.4000	0.1000	0.1000	0.1000

Table 5
Scores and rankings for the models with disadvantaged applicants, $D = \{1,5\}$.

Applicant	(P/U)		(P/M)		(P/L)		(P/E)	
	Score	Rank	Score	Rank	Score	Rank	Score	Rank
1	0.7370	1	0.7328	1	0.7328	1	0.7134	2
2	0.7189	3	0.7193	3	0.7193	3	0.6841	3
3	0.6338	4	0.6288	4	0.6288	4	0.6116	4
4	0.5633	5	0.5629	5	0.5629	5	0.5966	5
5	0.7295	2	0.7328	1	0.7328	1	0.7268	1

Table 6
Optimal weights for the models with disadvantaged applicants, $D = \{1,4\}$.

Model	Weights				
	Number of papers	h-index	Grant income (£)	Number of co-authors	Average teaching evaluation
(P/U)	0.4000	0.1000	0.1000	0.1000	0.3000
(P/M)	0.2021	0.1979	0.1000	0.1000	0.4000
(P/L)	0.2021	0.1979	0.1000	0.1000	0.4000
(P/E)	0.3935	0.2740	0.1000	0.1000	0.1326

Table 7
Scores and rankings for the models with disadvantaged applicants, $D = \{1,4\}$.

Applicant	(P/U)		(P/M)		(P/L)		(P/E)	
	Score	Rank	Score	Rank	Score	Rank	Score	Rank
1	0.6370	4	0.6814	3	0.6814	3	0.6752	4
2	0.6856	3	0.7139	2	0.7139	2	0.6757	3
3	0.8005	1	0.7572	1	0.7572	1	0.6852	1
4	0.7633	2	0.6814	3	0.6814	3	0.6752	4
5	0.6295	5	0.6635	5	0.6635	5	0.6852	1

Table 8
Qualification values for the additional applicants.

Applicant	Number of papers	h-index	Grant income (£)	Number of co-authors	Average teaching evaluation	Disadvantaged
6	9	5	81627	8	3.49	No
7	9	7	7230	5	3.13	No
8	14	5	81776	5	2.71	No
9	14	4	62508	5	4.00	No
10	15	5	74263	6	3.13	No
11	15	4	79399	4	3.96	No
12	8	4	68172	5	3.64	Yes
13	6	2	49790	8	3.64	No
14	11	4	63000	6	3.84	Yes
15	17	5	16617	6	3.53	No
16	15	7	24347	7	4.79	Yes
17	9	4	54205	1	4.76	No
18	8	3	16806	8	2.98	No
19	16	6	75450	1	3.41	No
20	10	5	74888	8	4.65	No

becomes ranked as first for models (M) and (L), despite the fact that none of the models ranked this applicant above third rank before.

The remedy for this shortcoming stems from the ability of (M) to identify the weakest applicant. The algorithm we propose is to solve model (M) repeatedly, eliminating the applicant (or applicants) with the minimum score every time. The same algorithm may be applied to the results of (U) and (E), although these models are more robust

due to their focus on the maximum score. The results of the proposed algorithm are presented in [Table 11](#).

As expected, models (U) and (E) returned similar results in a single run and with the elimination algorithm. However, the use of the elimination algorithm resulted in models (M) and (L) (that still return identical results) to yield results that agree with either (U) or (E) for each shortlisted candidate. As a final remark, we note that the

Table 9
Optimal weights for the larger instance.

Model	Weights				
	Number of papers	h-index	Grant income (E)	Number of co-authors	Average teaching evaluation
(U)	0.3000	0.1000	0.1000	0.1000	0.4000
(P/U)	0.1000	0.3000	0.1000	0.1000	0.4000
(M)	0.1000	0.1809	0.1000	0.2574	0.3616
(P/M)	0.1000	0.2131	0.1869	0.1000	0.4000
(L)	0.1000	0.1809	0.1000	0.2574	0.3616
(P/L)	0.1000	0.2131	0.1869	0.1000	0.4000
(E)	0.3252	0.1509	0.1771	0.2086	0.1382
(P/E)	0.2895	0.2017	0.2258	0.1000	0.1830

Table 10
Ranking of the applicants for the larger instance.

Applicant	(U)	(P/U)	(M)	(P/M)	(L)	(P/L)	(E)	(P/E)
1	15	4	16	14	16	14	16	16
2	10	3	3	3	3	3	4	3
3	2	5	6	4	6	4	9	7
4	5	16	10	17	10	17	13	14
5	16	6	5	11	5	11	8	11
6	12	7	4	5	4	5	6	9
7	19	18	18	19	19	19	18	19
8	13	15	14	13	14	13	7	8
9	6	10	9	7	9	7	10	10
10	7	11	7	8	7	8	1	1
11	4	8	11	6	11	6	5	4
12	17	17	15	14	15	14	15	15
13	18	19	13	18	13	18	17	18
14	11	12	8	10	8	10	14	12
15	8	14	12	16	12	16	12	13
16	1	1	2	2	2	2	3	6
17	14	13	18	12	20	12	19	17
18	20	20	18	20	18	20	19	20
19	9	9	17	9	17	9	11	5
20	3	2	1	1	1	1	1	2

final shortlisting decisions for all models are almost identical, with differences in rankings. We observe that the iterative elimination algorithm frequently returns ties for the lowest scored applicant for models (M) and (L). In such a case, all such candidates should be eliminated from the shortlist or retained in it. This may result in the number of shortlisted candidates to be less than the intended number. We recommend the use of model (U) as a tiebreaker in such cases.

6. Conclusions

In this short paper, we have applied four Abstract Moral Theories to the shortlisting process in hiring, with the objective of maximizing ethicality. We have presented LP models, applied them to a small and a larger instance, and presented our findings on the performance of the models. We have concluded that to ensure appropriate representation of disadvantaged applicants, a quota should be established, and disadvantaged individuals should be evaluated separately. We have also found that iteratively eliminating the applicant with the lowest score results in better results, particularly for models with a focus on the lowest score.

The choice of model and solution algorithm have a minor effect on the computational effort required. Hence, the choice depends on the ethical theory that the decision maker believes to be appropriate. Based on our analysis, we recommend the use of the iterative elimination algorithm with the Maximin model (M). This combination ensures that each eliminated applicant could not possibly be scored higher, which we believe to be a valid reason for elimination.

There is a number of limitations to the models we present. They require the decision maker to carefully consider the most relevant qualification types for the applicants, followed by meticulous preprocessing to extract the data from the application documentation. This is

Table 11
Comparison of the results of the single run and eliminating the applicant with the lowest score.

Applicant	Single run			Eliminate lowest score		
	(U)	(M), (L)	(E)	(U)	(M), (L)	(E)
1						
2						
3		3	4		4	1
4	2			2	1	
5				5		
6		5				
7		4				
8						
9						
10						
11			1		4	1
12	4		5	4	3	3
13						
14						
15						
16	1	2	3	1	4	5
17						
18						
19						
20	3	1	1	3	2	4

a manual task and can be time consuming. The models also assume that the applicants are all applying to the same level of job. For a multi-level job opening (e.g., Assistant / Associate / Full Professor), applicants of significantly different qualification levels will apply, and the models will eliminate all junior candidates in this scenario. Finally, the models assume that “more is better” for each qualification type, which may not be true in certain cases.

A possible topic for future work is the application of the models to real-world data and observe the discrepancies between human and model-based ranking decisions. Another venue of research would be the implementation of the existing Social Choice Theory algorithms onto the results of the models we have presented, i.e. treating the models as individuals voting for the ranking of applicants and using a Concordet method to determine the outcome. Finally, we believe that it would be a valuable contribution to model solve the case with a multi-level job opening.

The models we present are not based on past data, they are scalable to larger instances, and their results are interpretable. We hope that these models will pave the way to the development of more sophisticated models and a more ethical hiring process, as well as new application fields for Computational Ethics.

CRediT authorship contribution statement

Güneş Erdoğan: Conceptualization, Data curation, Formal analysis, Methodology, Software, Writing – original draft, Writing – review & editing.

Acknowledgments

Thanks are due to Danish Mishra and Ralph Behnke for a fruitful discussion on the recruitment process as part of the groundwork for the paper. We are grateful to Prof Michael Lewis for his feedback on an earlier version of the paper and to the two anonymous reviewers for their constructive feedback that helped improve the paper.

References

Allen, C., Wallach, W., Smit, I., 2006. Why machine ethics? *IEEE Intell. Syst.* 21 (4), 12–17.

- Anderson, M., Anderson, S.L., 2011. *Machine Ethics*. Cambridge University Press.
- Anderson, M., Anderson, S.L., Armen, C., 2006. An approach to computing ethics. *IEEE Intell. Syst.* 21 (4), 56–63.
- Cervantes, J.-A., López, S., Rodríguez, L.-F., Cervantes, S., Cervantes, F., Ramos, F., 2020. Artificial moral agents: A survey of the current status. *Sci. Eng. Ethics* 26, 501–532.
- Edwards, M.R., Edwards, K., 2019. *Predictive HR Analytics: Mastering the HR Metric*. Kogan Page Publishers.
- Fernandez-Mateo, I., Fernandez, R.M., 2016. Bending the pipeline? Executive search and gender inequality in hiring for top management jobs. *Manage. Sci.* 62 (12), 3636–3655.
- Jimenez, C.C.E., Matsuzaki, K., Gustilo, R.C., 2018. Fuzzy-based intelligent shortlisting process for human resource job recruitment procedures. *Int. J. Eng. Technol. (UAE)* 7, 229–233.
- Lambrecht, A., Tucker, C., 2019. Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Manage. Sci.* 65 (7), 2966–2981.
- Margherita, A., 2021. Human resources analytics: A systematization of research topics and directions for future research. *Hum. Resour. Manag. Rev.* <http://dx.doi.org/10.1016/j.hrmr.2020.100795>, in press.
- Marler, J.H., Boudreau, J.W., 2017. An evidence-based review of HR analytics. *Int. J. Hum. Resour. Manag.* 28, 3–26.
- Office for National Statistics, 2019. Employee turnover levels and rates by industry section, UK, January 2017 to December 2018. Online; Accessed 07 April 2021, <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/adhocs/10685employeeturnoverlevelsandratesbyindustrysectionukjanuary2017todecember2018>.
- Pessach, D., Singer, G., Avrahami, D., Ben-Gal, H.C., Shmueli, E., Ben-Gal, I., 2020. Employees recruitment: A prescriptive analytics approach via machine learning and mathematical programming. *Decis. Support Syst.* 134, 113290.
- Ruffle, B.J., Shtudiner, Z., 2015. Are good-looking people more employable? *Manage. Sci.* 61 (8), 1760–1776.
- Segun, S.T., 2021. From machine ethics to computational ethics. *AI Soc.* 36, 263–276.
- Tännsjö, T., 2019. *Setting Health-Care Priorities: What Ethical Theories Tell Us*. Oxford University Press, USA.
- Zhu, L., Liu, J., Xie, J., Yu, Y., Gao, L., Li, S., Duan, H., 2021. Can efficiency evaluation be applied to power plant operation improvement? A combined method with modified weighted russell directional distance model and pattern matching. *Comput. Oper. Res.* 134, 105406.