



US011030583B1

(12) **United States Patent**
Garg et al.

(10) **Patent No.:** **US 11,030,583 B1**
(45) **Date of Patent:** **Jun. 8, 2021**

(54) **SYSTEM, METHOD, AND COMPUTER PROGRAM FOR AUTOMATICALLY REMOVING DATA FROM CANDIDATE PROFILES THAT MAY INFLUENCE BIAS**

(71) Applicant: **Eightfold AI Inc.**, Mountain View, CA (US)

(72) Inventors: **Ashutosh Garg**, Sunnyvale, CA (US);
Varun Kacholia, Sunnyvale, CA (US)

(73) Assignee: **Eightfold AI Inc.**, Mountain View, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 300 days.

(21) Appl. No.: **16/209,834**

(22) Filed: **Dec. 4, 2018**

(51) **Int. Cl.**

G06Q 10/00 (2012.01)

G06Q 10/10 (2012.01)

G06N 20/00 (2019.01)

G06Q 10/06 (2012.01)

G06F 40/295 (2020.01)

(52) **U.S. Cl.**

CPC **G06Q 10/1053** (2013.01); **G06F 40/295** (2020.01); **G06N 20/00** (2019.01); **G06Q 10/063112** (2013.01)

(58) **Field of Classification Search**

CPC **G06Q 10/1053**; **G06Q 10/063112**; **G06F 40/295**; **G06N 20/00**

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,472,097 B1 12/2008 Scarborough et al.
9,665,641 B1 5/2017 Zhang

10,185,712 B2 1/2019 Gidney
2005/0004905 A1 1/2005 Dresden
2005/0086186 A1 4/2005 Sullivan et al.
2005/0216295 A1 9/2005 Abrahamsohn
2005/0261956 A1 11/2005 Kato
2006/0235884 A1 10/2006 Pfenniger et al.
2006/0271421 A1 11/2006 Stenecker et al.
2007/0047802 A1 3/2007 Puri
2007/0112585 A1 5/2007 Breiter et al.
2009/0144075 A1 6/2009 Flinn et al.

(Continued)

OTHER PUBLICATIONS

Time for a Diversity 'Reboot' DePompa, Barbara, SC Magazine 29.4: 26-29, Haymarket Media, Inc. (Jul./Aug. 2018).*

(Continued)

Primary Examiner — Gabrielle A McCormick

(74) Attorney, Agent, or Firm — Zhong Law, LLC

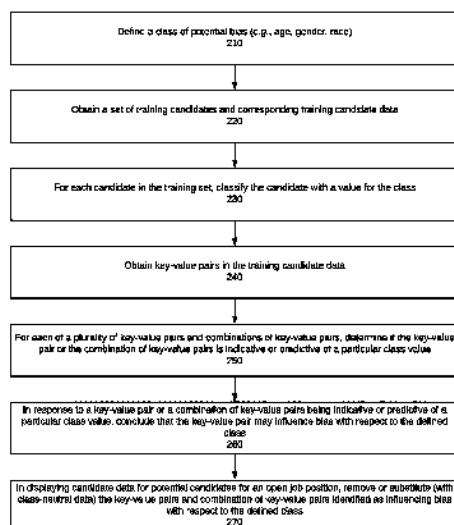
(57)

ABSTRACT

The present disclosure describes a system, method, and computer program for removing or replacing information in candidate profiles that may influence bias. The system identifies data in candidate profiles that may influence bias with respect to one or more defined classes of bias, such as gender, race, or age. The identified data is removed or replaced with an abstracted/neutral form of the data. The system may also remove information that personally identifies a candidate. In certain embodiments, the system determines what data is relevant for the job for which a candidate is applying and, for data that is relevant, at what level it is relevant (i.e., at the base value or an abstracted level). In such embodiments, the system removes data from the candidate's profile that is not relevant for the job role and, where applicable, abstracts some of the remaining data to levels that are relevant for the job role.

29 Claims, 10 Drawing Sheets

REMOVING CANDIDATE DATA THAT MAY INFLUENCE BIAS WITH RESPECT TO A DEFINED CLASS OF BIAS



(56)

References Cited

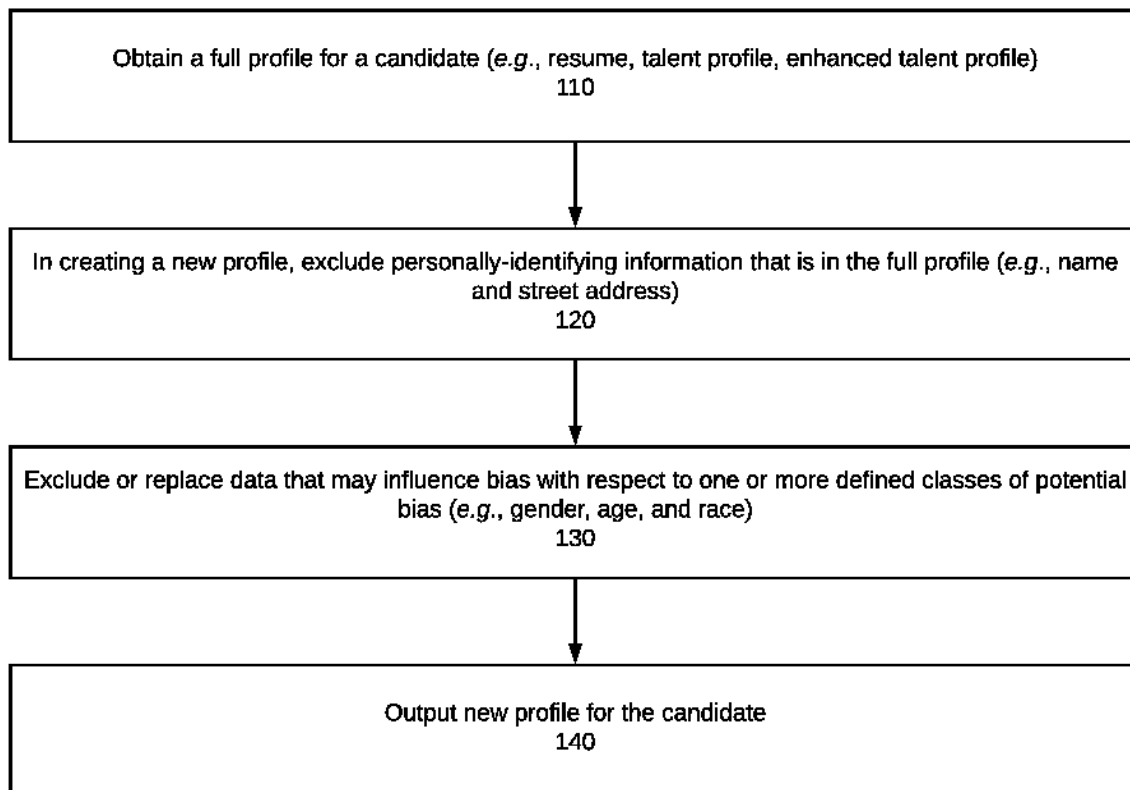
U.S. PATENT DOCUMENTS

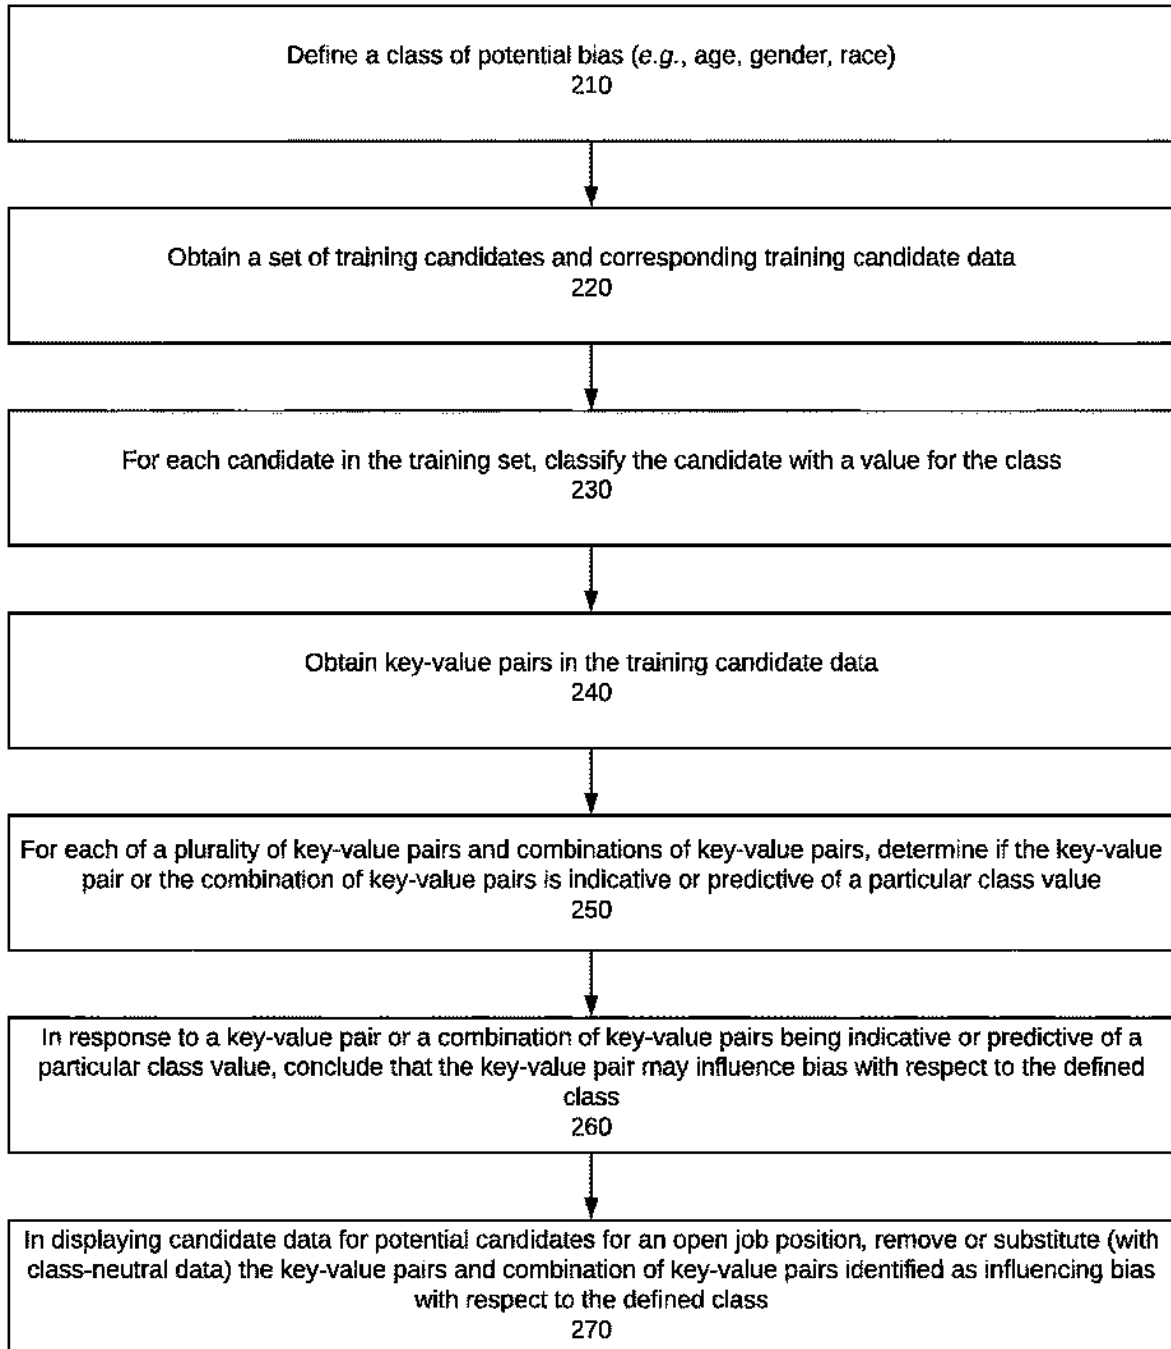
- 2010/0153149 A1 6/2010 Prigge et al.
 2010/0153150 A1 6/2010 Prigge et al.
 2011/0055098 A1 3/2011 Stewart
 2011/0276505 A1 11/2011 Schmitt
 2013/0096991 A1 4/2013 Gardner et al.
 2013/0290208 A1 10/2013 Bonmassar et al.
 2014/0039991 A1 2/2014 Gates et al.
 2014/0122355 A1 5/2014 Hardtke et al.
 2014/0282586 A1 9/2014 Shear et al.
 2014/0330734 A1 11/2014 Sung et al.
 2015/0142711 A1 5/2015 Pinckney et al.
 2015/0161567 A1 6/2015 Mondal et al.
 2015/0178682 A1 6/2015 Matthews et al.
 2015/0244850 A1 8/2015 Rodriguez et al.
 2015/0309986 A1 10/2015 Brav et al.
 2015/0317610 A1 11/2015 Rao et al.
 2016/0012395 A1 1/2016 Omar
 2016/0034853 A1 2/2016 Wang et al.
 2016/0055457 A1* 2/2016 Mather G06Q 10/1053
 705/321
 2016/0098686 A1 4/2016 Younger
 2017/0061081 A1 3/2017 Jagannathan et al.
 2017/0243162 A1 8/2017 Gavrielides et al.
 2017/0344555 A1 11/2017 Yan et al.
 2017/0357945 A1 12/2017 Ashkenazi et al.
 2018/0039946 A1 2/2018 Bolte et al.
 2018/0150484 A1 5/2018 Dupey et al.
 2018/0218330 A1 8/2018 Choudhary et al.
 2018/0232751 A1* 8/2018 Terhark G06Q 10/063118
 2018/0308061 A1 10/2018 Jadda et al.
 2018/0357557 A1* 12/2018 Williams G06Q 30/0201
 2018/0373691 A1 12/2018 Alba et al.
 2019/0066056 A1* 2/2019 Gomez G06Q 10/06393
 2019/0114593 A1 4/2019 Champaneria
 2019/0197487 A1 6/2019 Jersin et al.
 2019/0205838 A1 7/2019 Fang et al.
 2020/0007336 A1* 1/2020 Wengel G06F 16/1824
 2020/0065769 A1 2/2020 Gupta et al.
 2020/0117582 A1* 4/2020 Srivastava G06F 11/3688
 2020/0160050 A1 5/2020 Bhotika et al.

OTHER PUBLICATIONS

What resume? Hiring is in the midst of a tech revolution; Cutting costs, turnover; Eliminating bias. Lejalde-Ruiz, Alexia. South Florida Sun—Sentinel [Fort Lauderdale, Fla] Jul. 22, 2018: G.1.*
 4 Ways technology has changed recruitment for better (and worse) White, Sarah K. CIO CXO Media, Inc. (Mar. 15, 2017).*
 Feds try to blank out bias. Dobson, Sarah. Canadian HR Reporter 30.9: 1.9. IIAB Press Limited. (May 15, 2017).*
 How Congress Could Reduce Job Discrimination By Promoting Anonymous Hiring. Hausman, David. Stanford Law Review 64.5: 1343-1369. Stanford University, Stanford Law School. (May 2012).*
 International Application No. PCT/US2020/012317, International Search Report and Written Opinion dated Apr. 9, 2020, 8 pages.
 CUSTOMERGLUE “Hire the best candidate for your Company using Artificial Intelligence” (2016), medium.com/SCustomerGlu, 2016, 5 pages.
 Pedreschi et al., “Discrimination-Aware Data Mining,” Aug. 24-27, 2008, KDD 08, Las Vegas, Nevada, 9 pages.
 Dixon et al., “Measuring and Mitigating Unintended Bias in Text Classification,” Proceeding of the 2018 AAAI/ACM Conf. on AI, Ethics, and Society. Feb. 2-3, 2018, 7 pages.
 Liu et al., “Delayed Impact of Fair Machine Learning,” arXiv:1803.04383v2, Apr. 7, 2018, 37 pages.
 Hardt et al., “Equality of Opportunity in Supervised Learning,” arXiv:1610.02413v1, Oct. 7, 2016, 22 pages.

* cited by examiner

Creating "New Profile" for Candidate**FIG. 1**

REMOVING CANDIDATE DATA THAT MAY INFLUENCE BIAS WITH RESPECT TO A DEFINED CLASS OF BIAS**FIG. 2**

**EXAMPLE METHOD FOR DETERMINING WHETHER A KEY-VALUE PAIR IS INDICATIVE OR
PREDICTIVE OF A PARTICULAR CLASS VALUE**

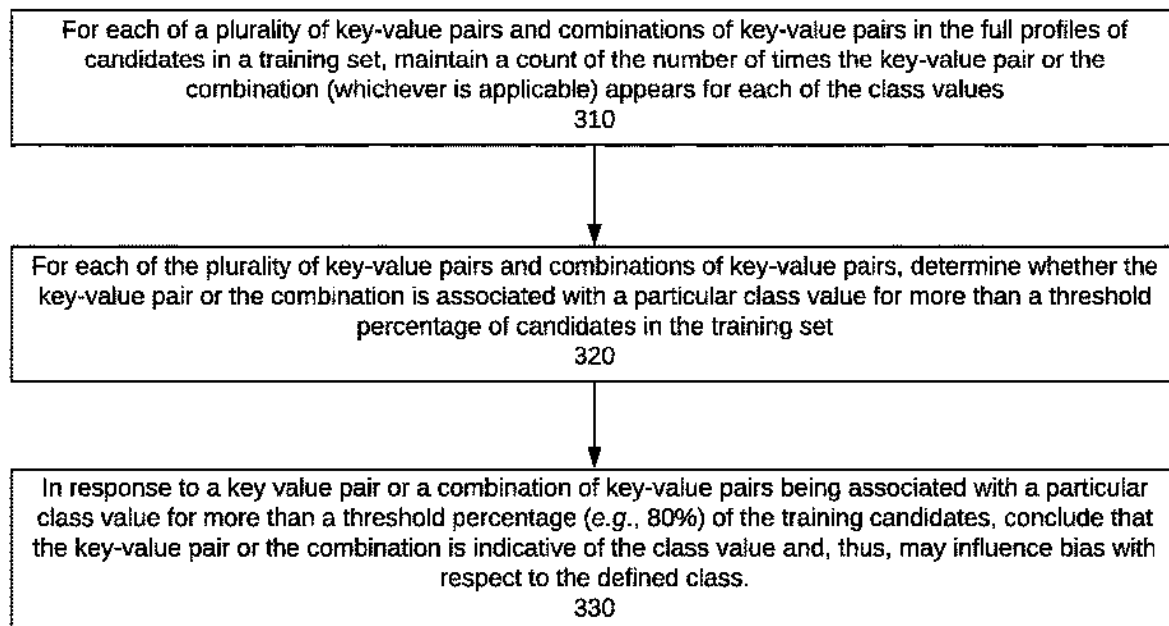


FIG. 3

**EXAMPLE MACHINE-LEARNING METHOD FOR DETERMINING WHETHER A KEY-VALUE PAIR IS
INDICATIVE OR PREDICTIVE OF A PARTICULAR CLASS VALUE**

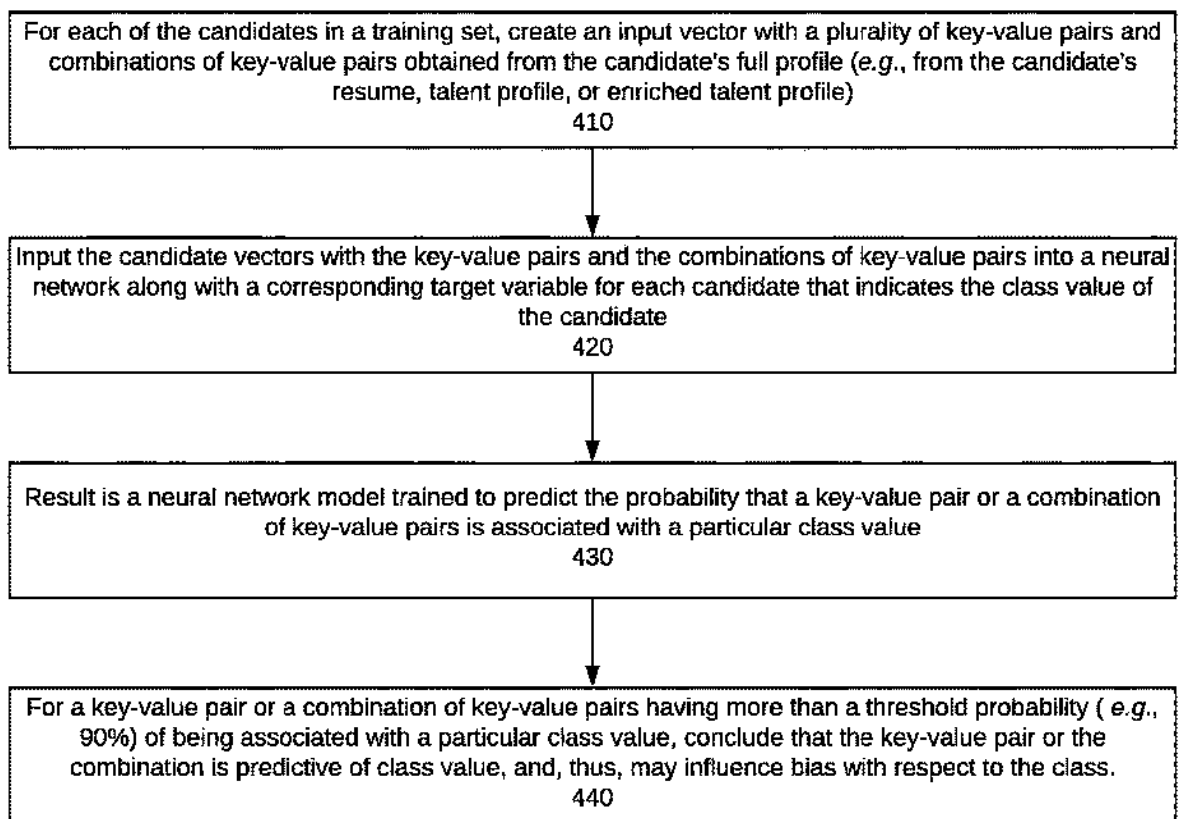


FIG. 4

**EXAMPLE OF METHOD FOR IDENTIFYING KEY-VALUE PAIRS THAT MAY INFLUENCE GENDER BIAS
(Using Threshold Method)**

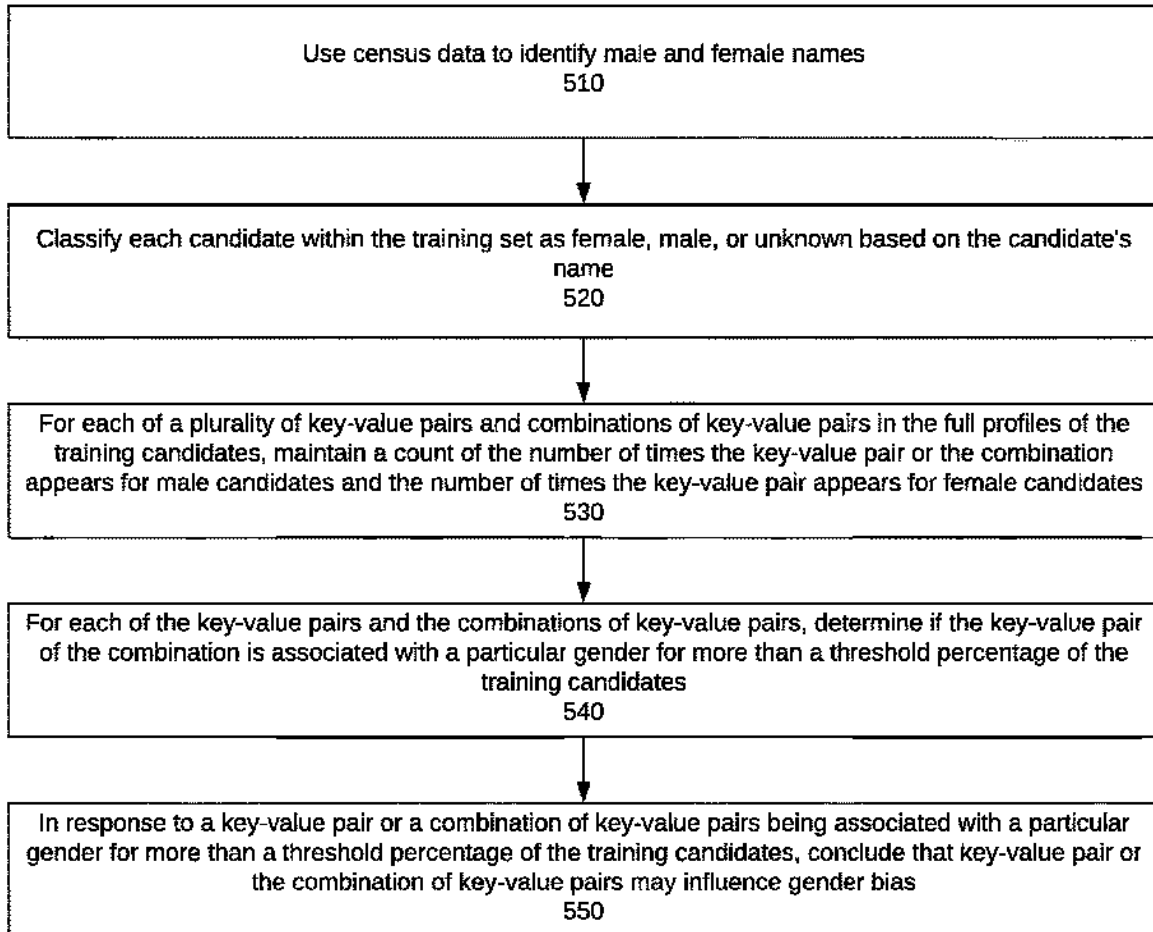


FIG. 5

**EXAMPLE OF METHOD FOR IDENTIFYING KEY-VALUE PAIRS THAT MAY INFLUENCE RACE BIAS
(Using Neural Network)**

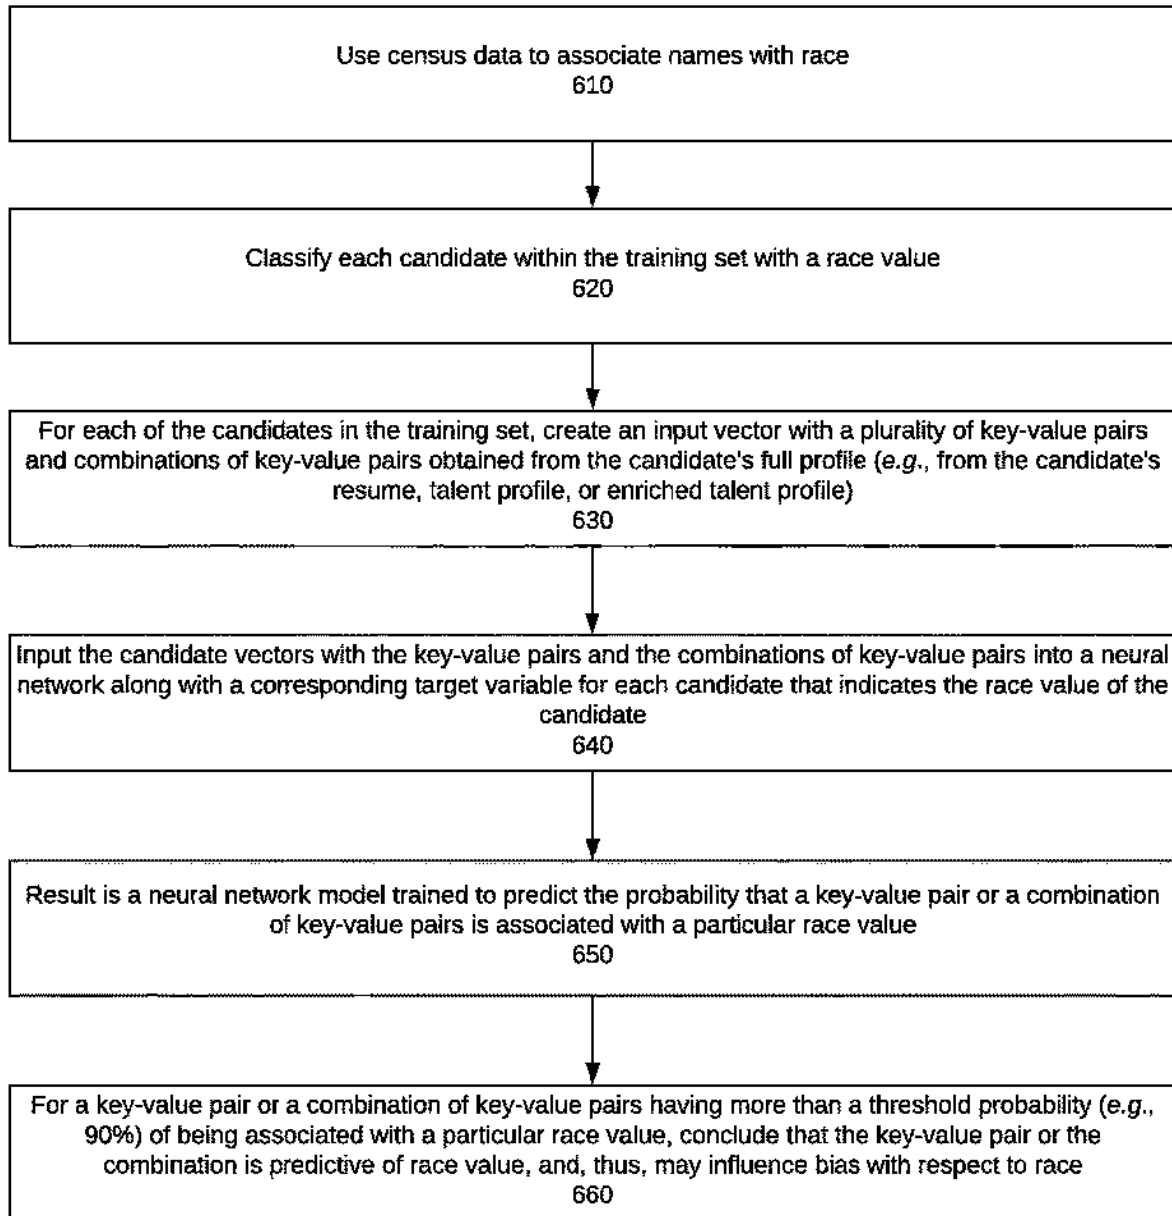


FIG. 6

**EXAMPLE OF METHOD FOR IDENTIFYING KEY-VALUE PAIRS THAT MAY INFLUENCE AGE BIAS
(Using Threshold Method)**

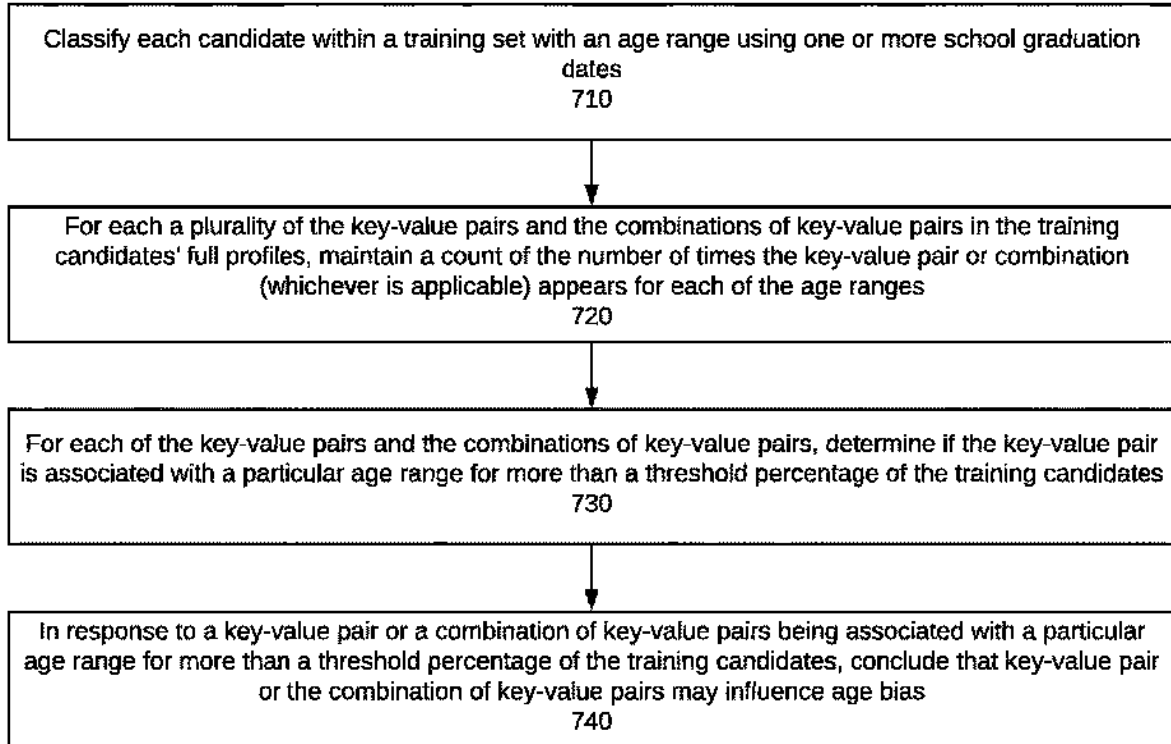
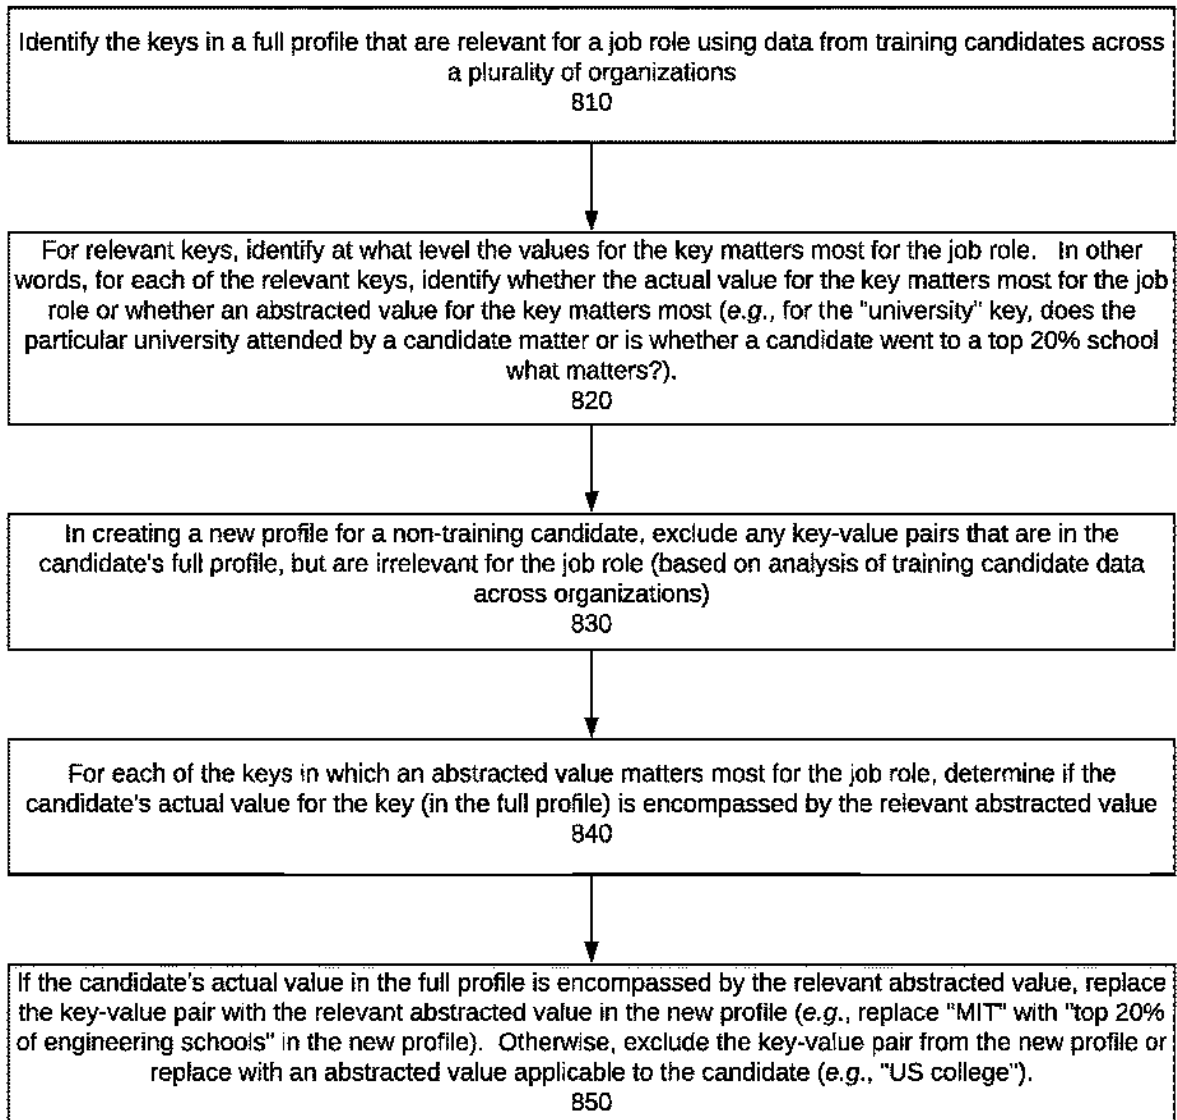
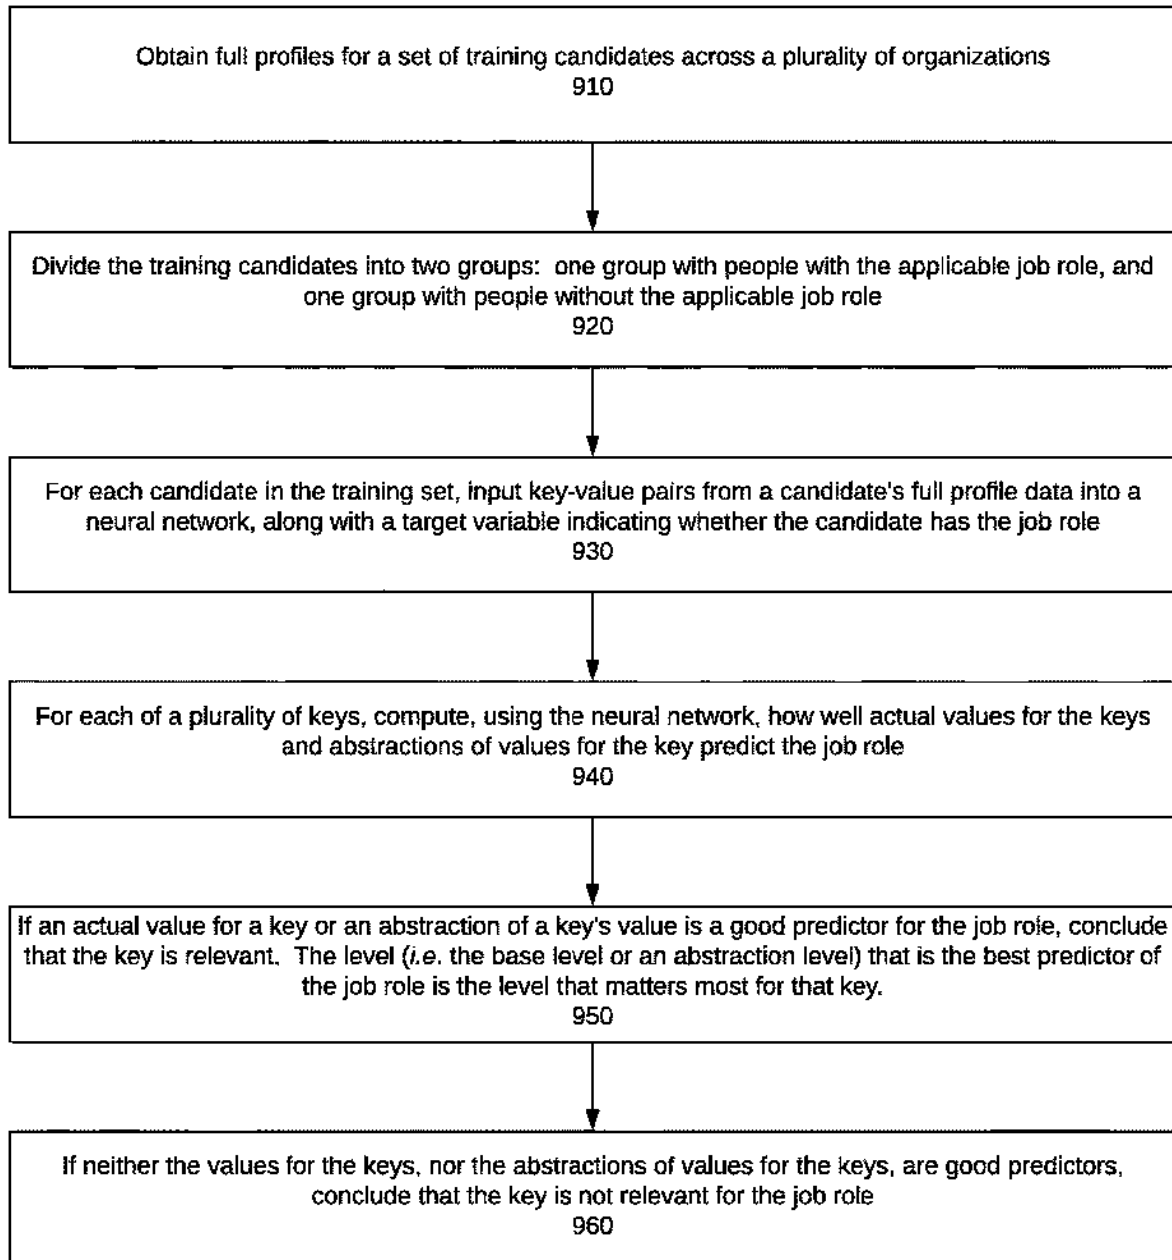
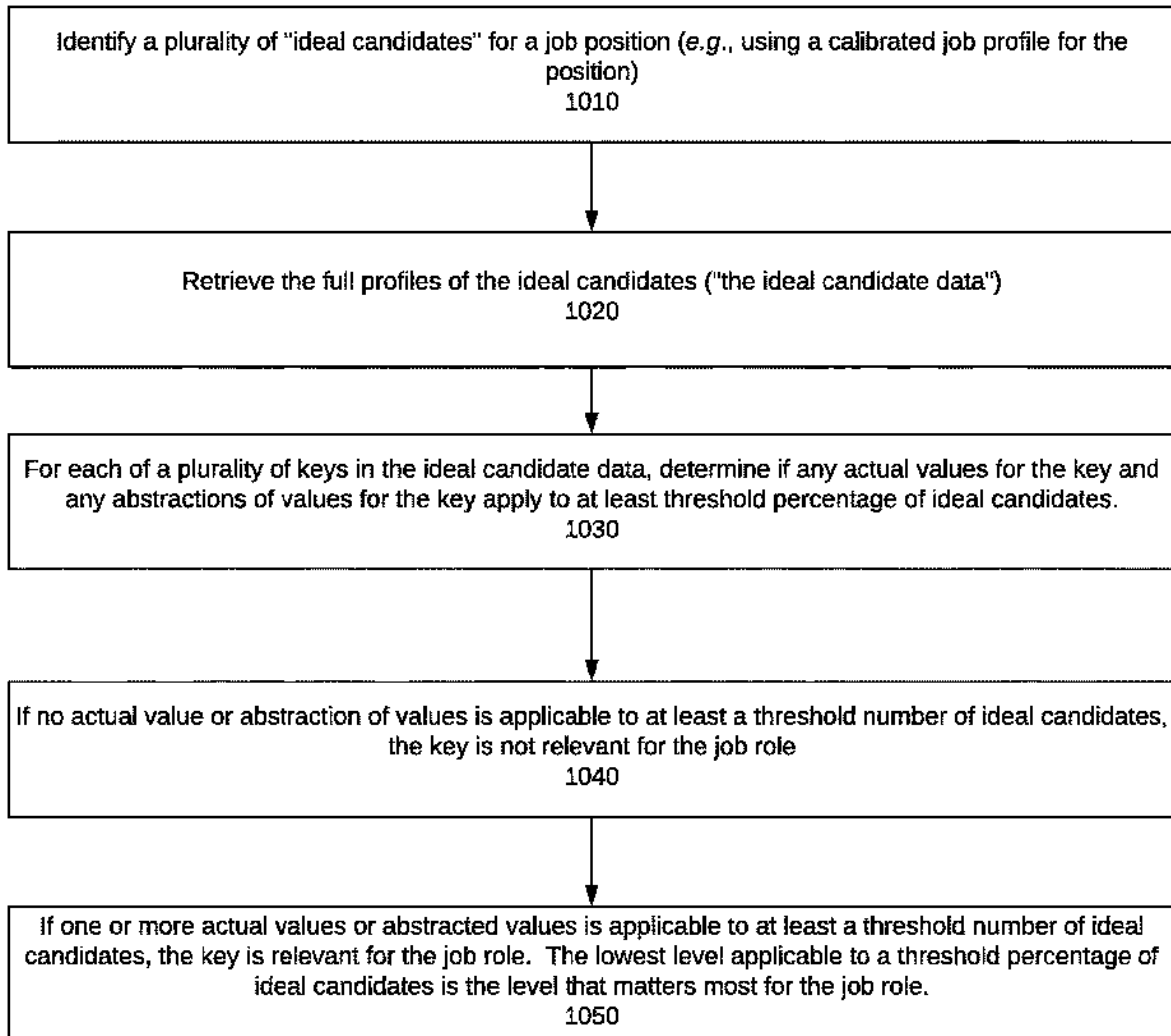


FIG. 7

ABSTRACTING CANDIDATE DATA TO RELEVANT LEVELS**FIG. 8**

IDENTIFYING RELEVANT KEYS AND ABSTRACTION LEVELS FOR KEY-VALUE PAIRS USING A NEURAL NETWORK**FIG. 9**

IDENTIFYING RELEVANT KEYS AND ABSTRACTION LEVELS FOR KEY-VALUE PAIRS USING A CALIBRATED JOB PROFILE**FIG. 10**

1

SYSTEM, METHOD, AND COMPUTER PROGRAM FOR AUTOMATICALLY REMOVING DATA FROM CANDIDATE PROFILES THAT MAY INFLUENCE BIAS

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates generally to human resource applications, and more specifically to removing data from candidate profiles that may influence bias.

2. Description of the Background Art

Many companies and organizations use various HR applications to identify candidates for open job positions. An example of a system that identifies candidates for job openings is described in U.S. patent application Ser. No. 16/121,401, filed on Sep. 4, 2018, and titled "System, Method, and Computer Program for Automatically Predicting the Job Candidates Most Likely to be Hired and Successful in a Job," the contents of which are incorporated by reference herein. Such systems typically display a list of potential candidates and enable a user (e.g., an HR manager) at the organization to view profiles for the potential candidates. The profiles may be resumes, talent profiles, or enhanced talent profiles as described in the incorporated U.S. patent application Ser. No. 16/121,401.

In reviewing candidate profiles, conscious or unconscious biases by reviewers at the organization may cause them to overlook candidates that are well qualified for a job position. A bias may be a type of bias that is generally known to organizations, such as those related to gender, race, or age. Reviewers may also have biases that are unknown to the organization. For example, a reviewer may make assumptions about a candidate based on a hobby listed on the resume. For example, a reviewer looking for a salesperson with an outgoing personality and a "people person" may assume that someone who plays chess is an introvert, which may or may not be true. In such case, if the candidate otherwise meets the qualifications for the job, it would be better to interview the candidate instead of dismissing the candidate based on his/her profile.

To reduce the likelihood of bias in the initial screening process, it would be helpful to remove data from a candidate's profile that may influence bias. Therefore, there is demand for an HR application that not only identifies potential candidates for open job positions, but creates new profiles for the candidates that exclude data that may influence bias.

SUMMARY OF THE DISCLOSURE

The present disclosure describes a system, method, and computer program for removing or replacing information in candidate profiles that may influence bias. The method is performed by a computer system that identifies potential candidates for open job positions and displays profiles for the identified candidates ("the system").

In one embodiment, the system first creates or obtains a "full profile" for a candidate and then creates a "new profile" for the candidate that excludes or substitutes data in the full profile that may influence reviewer bias. The "new profiles" may be the profiles initially displayed to a user screening candidates for an open job position.

2

In one embodiment, the system excludes or substitutes the following types of information in creating the new profile: (1) data that personally identifies the candidate, such as the candidate's name, and (2) data that is indicative of gender, race, and/or age (or another defined class of bias). In one embodiment, a method for creating a profile for a job candidate that excludes data that may influence reviewer bias comprises the following steps:

- defining a class of potential bias;
- identifying key-value pairs in candidate data that may influence bias with respect to the defined class by performing the following:
 - obtaining a set of training candidates;
 - for each candidate in the set, classifying the candidate with a value for the class;
 - obtaining key-value pairs in the candidate data for the training set (the "training candidate data")
 - for each of a plurality of key-value pairs in the training candidate data, determining if the key-value pair is indicative of a particular class value; and
 - in response to a key-value pair being indicative of a particular class value, concluding that the key-value pair may influence bias with respect to the defined class;
- obtaining a first profile for a non-training candidate, wherein the first profile includes the non-training candidate's name and one or more key-value pairs identified as influencing bias in the training candidate data; and
- creating a second profile for the non-training candidate from the first profile by removing or substituting the non-training candidate's name and the key-value pair(s) identified as influencing bias.

In a further embodiment, the system may also exclude any information in the new profile that is not relevant to the job role for which the candidate is applying. In such embodiments, the system determines what data is relevant for the job for which the candidate is applying and, for data that is relevant, at what level it is relevant (i.e., at the base value or an abstracted level). The system then removes data from the profile that is not relevant for the job role and, where applicable, abstracts some of the remaining data to levels that are relevant for the job role.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a flowchart that illustrates a method for creating a new profile for a candidate that excludes data that may influence bias according to one embodiment.

FIG. 2 is a flowchart that illustrates a method, according to one embodiment, for identifying candidate data that may influence bias with respect to a defined class of bias.

FIG. 3 is a flowchart that illustrates a method, according to one embodiment, for determining whether a key-value pair is indicative of a particular class value.

FIG. 4 is a flowchart that illustrates a method, according to an alternate embodiment, for determining whether a key-value pair is indicative of a particular class value.

FIG. 5 is a flowchart that illustrates an example method for identifying key-value pairs that may influence gender bias.

FIG. 6 is a flowchart that illustrates an example method for identifying key-value pairs that may influence race bias.

FIG. 7 is a flowchart that illustrates an example method for identifying key-value pairs that may influence age bias.

3

FIG. 8 is a flowchart that illustrates a method, according to one embodiment, for abstracting candidate data to relevant levels for a job role.

FIG. 9 is a flowchart that illustrates a method, according to one embodiment, for determining what keys in candidate data are relevant for a job role, and, for keys that are relevant, at what level the key matters most.

FIG. 10 is a flowchart that illustrates a method, according to an alternate embodiment, for determining what keys in candidate data are relevant for a job role, and, for keys that are relevant, at what level the key matters most.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The present disclosure describes a system, method, and computer program for removing or replacing information in candidate profiles that may influence bias. The method is performed by a computer system that identifies potential candidates for open job positions and displays profiles for the identified candidates ("the system"). An example of the system is described in U.S. patent application Ser. No. 16/121,401 (incorporated herein above).

For each open job position in an organization, the system displays a list (typically ranked) of identified candidates for the job position. A user (e.g., an HR manager) can select any of the candidate and see a profile associated with the candidate. In one embodiment, the initial profile viewed by the user is a "new profile" created by the system from a "full profile" for the candidate. The "full profile" may be the candidate's resume, talent profile, or enhanced talent profile (e.g., the "enhanced talent profile" described in the U.S. patent application Ser. No. 16/121,401, incorporated herein above). The new profile is based on the candidate's full profile, but excludes data that may influence bias. An organization may configure the system to enable the user to see the full profile at a later point in the interviewing/screening process. A method for creating the new profile is described below.

1. Creating a New Candidate Profile from a Candidate's Full Profile

FIG. 1 illustrates a method for creating a new profile for a candidate that excludes data that may influence bias. The system obtains a full profile for the candidate (step 110). For avoidance of doubt, the term "full profile" or "full candidate profile" herein means any profile (e.g., resume, talent profile, enhanced talent profile, etc.) for the candidate obtained or created by the system before data in the profile is removed or replaced in accordance with the methods described herein.

In creating the new profile, the system excludes personally-identifiable information that is in the full profile, such as name and/or street address information (step 120). The system also excludes or replaces candidate data that may influence bias with respect to one or more defined classes of bias (e.g., gender, race, or age) (step 130). This step is described in more detail with respect to FIGS. 2-7. The system then outputs the new profile, which is displayed in a user interface for the system (step 140).

2. Identifying Data that May Influence Bias with Respect to a Defined Class of Bias

FIG. 2 illustrates a method for identifying data in full candidate profiles that may influence bias with respect to a

4

defined class of bias. The system receives as an input a defined class of bias, such as race, gender, and age (step 210). The defined class may be inputted by a system administrator or a user of the system via a user interface.

The system obtains candidate data for set of training candidates, preferably from across a plurality of organizations and a variety of professions (step 220). The training candidate data includes their full candidate profiles (e.g., a resume, talent profile, or enhanced talent profile) and data that enables each of the candidates to be classified with a class value (e.g., name or school graduation date). The system may obtain the training candidate data from a talent repository managed by the system (or an organization/company) or from public data sources that store job-candidate/employment profiles. The system classifies each of the training candidates with a class value (e.g., male or female) (step 230).

The system obtains key-value pairs from the full candidate profiles of the training candidates (step 240), and for each of a plurality of key-values pairs and combinations of key-value pairs, the system determines if the key-value pair or combination of key-value pairs is indicative of a particular class value (step 250). In response to a key-value pair or a combination of key-value pairs being indicative of a particular class value, the system concludes that the key-value pair or combination of key-value pairs may influence bias with respect to the defined class (step 260). In creating a new profile for a candidate, the system removes or substitutes (with class-neutral data) the key-value pairs and combination of key-value pairs identified as influencing bias with respect to the defined class (step 270). "Neutral" data that serves as a substitute for key-value pairs may be an abstracted form of the key-value pair. For example, a particular US college may be replaced with an abstracted value of the college, such as "4-year US college."

2.1 Determining Whether a Key-Value Pair is Indicative of a Particular Class Value

FIG. 3 illustrates a method for determining whether a key-value pair is indicative of a particular class value. This method is referred to herein as the "threshold method" in order to distinguish it from the method of FIG. 3, which is referred to as the "machine learning method." For each of a plurality of key-value pairs and combinations of key-value pairs in the full profiles of the training candidates, the system maintains a count of the number of times the key-value pair or the combination of key-value pairs appears for each of the class values (step 310), and determines whether the key-value pair or the combination (whichever is applicable) is associated with a class value for more than a threshold percentage (e.g., 80%) of candidates in the training set (step 320). If a key-value pair or a combination of key-value pairs is associated with a particular class value for more than the threshold percentage of candidates, the system concludes the key-value pair or the combination of key value-pairs (whichever is applicable) is indicative of the class value and, thus, may influence bias with respect to the defined class (step 330).

FIG. 4 illustrates an alternate method for determining whether a key-value pair is indicative of a particular class value. This method uses a neural network, such as a deep neural network, and, as stated above, is referred to herein as the "machine-learning method." For each of the candidates in the training set, the system creates an input vector for the candidate with a plurality of key-value pairs and combination of key-value pairs obtained from the training candidate's full profiles (step 410). The input vectors may include all the key-value pairs in the full profiles or just a subset of

5

key-value pairs. To train a neural network, the system inputs the vectors for each of the training candidates into the neural network, along with a corresponding target variable for each of the candidates that indicates the class value of the candidate (step 420). The result is a neural network that is trained to predict the probability that a key-value pair or combination of key-value pairs is associated with a particular class value (step 430). For a key-value pair or a combination of key-value pairs having more than a threshold probability (e.g., 90%) of being associated with a particular class value, the system concludes that the key-value pair or the combination of key-value pairs (whichever is applicable) is indicative of class value, and, thus, may influence bias with respect to the class (step 440).

The methods of FIGS. 3 and 4 could be applied to just individual key-value pairs (and not combinations of key-value pairs).

2.2 Example Method for Identifying Key-Value Pairs that May Influence Gender Bias.

FIG. 5 illustrates an example method for identifying key-value pairs that may influence gender bias. This example uses the threshold method described above with respect to FIG. 3, but the machine learning method could be used instead. Census data is inputted or imported into the system, which analyzes the data to identify male and female names (step 510). Specifically, names used for a threshold percentage (e.g., 90%) of one gender are classified as associated with that gender. The system uses this information to classify each candidate in the training set as female, male, or unknown based on the candidate's name (step 520).

For each of a plurality of key-value pairs and combinations of key-value pairs in the full profile of the training candidates, the system maintains a count of the number of times the key-value pair (or the combination) appears for male candidates and the number of times the key-value pair (or the combination) appears for female candidates (step 530), and determines whether the key-value pair or the combination (whichever is applicable) is associated with a particular gender for more than a threshold percentage (e.g., 80%) of candidates in the training set (step 540). If a key-value pair or a combination of key-value pairs is associated with a particular gender for more than the threshold percentage of candidates, the system concludes the key-value pair or the combination of key-value pairs (whichever is applicable) is indicative of the class value and, thus, may influence gender bias (step 550).

2.3 Example Method for Identifying Key-Value Pairs that May Influence Race Bias

FIG. 6 illustrates an example method for identifying key-value pairs that may influence race bias. This example uses the machine-learning method described above with respect to FIG. 4, but the threshold method could be used instead. Census data is inputted or imported into the system, which analyzes the data to associate names with races (step 610). Specifically, names used for a threshold percentage (e.g., 90%) of one race are classified as associated with that race value. The system uses this information to classify each candidate in the training set with a race value (which may include "unknown") based on the candidate's name (step 620).

For each of the training candidates, the system creates an input vector for the training candidate with a plurality of key-value pairs and combination of key-value pairs obtained from the training candidate's full profile (step 630). To train a neural network, the system inputs the vectors for each of the training candidates into the neural network, along with the candidate's race value (step 640). The result is a neural

6

network that is trained to predict the probability that a key-value pair or a combination of key-value pairs is associated with a particular race value (step 650). For a key-value pair or a combination of key-value pairs having more than a threshold probability (e.g., 90%) of being associated with a particular race value, the system concludes that the key-value pair or the combination of key-value pairs (whichever is applicable) is indicative of race value, and, thus, may influence racial bias (step 660).

2.4 Example Method for Identifying Key-Value Pairs that May Influence Age Bias.

FIG. 7 illustrates an example method for identifying key-value pairs that may influence age bias. This example uses the threshold method described above with respect to FIG. 4, but the machine learning method could be used instead. The system analyzes high school or college graduation dates, and uses these dates to classify each of the candidates in a training set with an age range value (e.g., 20-30 years old, 30-40 years old, etc.) (step 710). For example, in the year 2020, the system may classify someone that graduated from college in the year 2018, with a "20-30-year-old" age range.

For each of a plurality of key-value pairs and combinations of key-value pairs in the full profiles of the training candidates, the system maintains a count of the number of times the key-value pair (or the combination) appears for each of the age ranges (step 720), and determines whether the key-value pair or the combination (whichever is applicable) is associated with a particular age range for more than a threshold percentage (e.g., 80%) of candidates in the training set (step 730). If a key-value pair or a combination of key-value pairs is associated with a particular age range for more than the threshold percentage of candidates, the system concludes the key-value pair or the combination of key-value pairs (whichever is applicable) is indicative of age and, thus, may influence age bias (step 740).

3. Further Embodiments—Abstracting Candidate Data to Relevant Levels

In certain embodiments, creating a new profile for candidate also includes removing any data that is not relevant to the job role for which the candidate is applying. The methods describe above with respect to FIGS. 2-7 reduce the likelihood that a reviewer is influenced by a known type of bias, such as gender, race, or age, whereas abstracting candidate profiles to only what is relevant for a job also reduces the likelihood of a reviewer being influenced by reviewer-specific biases that may be unknown to the organization. For example, a particular reviewer may make assumptions about a candidate based on a hobby listed by the candidate, and the candidate's hobbies may not be relevant for obtaining or succeeding at the job for which the candidate is applying.

FIG. 8 illustrates a method for abstracting candidate data to relevant levels for a job role. The system identifies the keys (i.e., the fields) in a full candidate profile that are relevant for a job role using training candidate data, preferably from across a plurality of organizations (step 810). The training candidate data may be obtained from a talent repository maintained by the system or an organization, or from public data sources with employment profile data. Methods for performing this step are described with respect to FIGS. 9-10. If key-value pairs have already been eliminated or substituted in accordance with the method of FIG. 2, then this is performed for the remaining keys in the candidate's full profile.

For each of the relevant keys, the system identifies at what level the key matters most for the job role (step 820). In other words, for each of the relevant keys, the system identifies whether the actual value for the key matters most or whether an abstracted value for a key matters most. For example, for the "university" key, does the particular university attended by a candidate matter or is whether a candidate went to a top 20% school what matters?

In creating the new profile for the candidate, the system excludes any key-value pairs that are in the full candidate profile but are irrelevant for the job role (step 830). For each of the relevant keys in which an abstracted value matters most for the job role, the system determines whether the candidate's actual value for the key is encompassed by the relevant abstracted value (step 840). For example, if what matters most for the "university" key is whether a candidate went to a top 20% school, then system determines whether the university attended by the candidate is a top 20% school. They system may use published or inputted university rankings to make this determination.

If the candidate's actual value in his/her full profile is encompassed by the relevant abstracted value, the system replaces the key-value pair in the candidate's full profile with the relevant abstracted value in the new profile (step 850). For example, the system may replace "Massachusetts Institute of Technology" in a full profile with "top 20% of engineering schools" in the new profile. Otherwise, the system either excludes the key-value pair from the new profile or replaces the key-value pair with an abstracted value relevant to the candidate in the new profile, depending on how the system is configured (also step 850). For example, if the candidate attended a 4-year college that is not ranked in the top 20% of schools (according to the ranking(s) used by the system), then the system may not specify college information for the candidate or the system may replace the candidate's specific college with something like "US university." Key-value pairs that are not abstracted or removed in accordance with step 850 remain in the new profile.

If a candidate applies for multiple job positions at an organization, then the system may create a new profile for the candidate for each of the job roles, as what is relevant for one job role may not be relevant for another job role.

3.1 Identifying Relevant Keys and Abstraction Levels for Key-Value Pairs

FIGS. 9 and 10 illustrate two different methods for determining which keys in the candidate data are relevant for a particular job role and, for keys that are relevant, determining at what level the key matters most (i.e., at the base level or at an abstracted level). FIG. 9 uses a neural network, and FIG. 10 uses data from "ideal candidates" for a job position.

Turning to FIG. 9, the system obtains full profiles for a set of training candidates, preferably comprising people from across a plurality of organizations (step 910). The training set includes both people who have never held the applicable job role and people who currently have or have previously had the applicable job role. The system divides the candidates in the training set into two groups: One group with people who have or have had the applicable job role, and one group with people who have not held the applicable job role (step 920). For each candidate in the set, the system inputs key-value pairs from the candidate's full profile (in the form of an input vector) into a neural network, along with a target variable indicating whether the candidate has (or has had) the job role. In one embodiment, the neural network is a recurring neural network (RNN).

For each of a plurality of keys in the candidate data, the system computes, using the neural network, how well actual values for the keys and abstractions of values for the keys predict the job role (step 940). The "abstracted values" may be preconfigured by a system administrator or may be determined automatically by clustering values for keys into groups, where a group encompasses multiple values. The system may test multiple abstraction levels for a key.

If an actual value for a key or an abstraction of a key's value is a good predictor for the job role, the system concludes that the key is relevant (step 950). The level (i.e., the base level or an abstracted level) that is the best predictor of the job role is the level that matters most for the key. For example, for the "undergraduate college" key, if "top 20% school" is a better predictor than a particular university value (e.g., "Stanford"), then "top 20% school," which is an abstracted value, is the level that matters most for that key. Conversely, if for the "skills" key, the value "java" is a better predictor for a job role than an abstracted value that encompasses a wider range of skills, then the base level (i.e., the actual value) for the key is the level that matters most for that key. If neither the actual values for a key, nor abstractions of values for a key, are good predictors, the system concludes that the key is not relevant for the job role (step 960).

FIG. 10 illustrates an alternate method for identifying the relevant keys and abstraction levels for a job role. The system identifies a plurality of "ideal candidates," across organizations, for a job role, and retrieves full profiles for the ideal candidates ("the ideal candidate data") (steps 1010, 1020). In one embodiment, the system obtains ideal candidate data using a calibrated job profile for the job role. Calibrated job profiles are described in U.S. patent application Ser. No. 16/121,401 which is incorporated by reference herein above.

For each of a plurality of keys in the ideal candidate data, the system determines if any actual values for the key and any abstractions of values for the key apply to a least a threshold percentage (e.g., 80%) of ideal candidates (step 1030). If no actual value or abstraction of values is applicable to at least a threshold number of ideal candidates, the system concludes that the key is not relevant for the job role (step 1040). If one or more actual values or abstracted values is applicable to at least a threshold number of ideal candidates, the system concludes that the key is relevant for the job role (step 1050). The lowest level applicable to a threshold percentage of ideal candidates is the level that matters most for the job role. For example, if both an abstracted value for a key and a particular actual value for a key apply to at least a threshold percentage of ideal candidates, the base level (actual value) for the key is the level that matters most for that key.

4. General

The methods described herein are embodied in software and performed by a computer system (comprising one or more computing devices) executing the software. A person skilled in the art would understand that a computer system has one or more memory units, disks, or other physical, computer-readable storage media for storing software instructions, as well as one or more processors for executing the software instructions.

As stated an example of a computer system for performing the methods described herein is set forth in U.S. patent application Ser. No. 16/121,401 (incorporated herein above). In addition to the software modules described in

U.S. patent application Ser. No. 16/121,401, the system may have a software module for performing the methods described herein.

As will be understood by those familiar with the art, the invention may be embodied in other specific forms without departing from the spirit or essential characteristics thereof. Accordingly, the above disclosure is intended to be illustrative, but not limiting, of the scope of the invention, which is set forth in the following claims.

The invention claimed is:

1. A method for creating a profile for a job candidate that excludes data that influence reviewer bias, wherein the method is performed by a computer system that identifies potential candidates for an open job position, the method comprising:

defining a class of potential bias;

identifying key-value pairs in candidate data that influence bias with respect to the defined class by performing the following:

obtaining training candidate data representing candidate data for a set of training candidates;

for each training candidate in the set, classifying the training candidate with a value for the class;

obtaining a plurality of key-value pairs in the training candidate data;

for each of the plurality of key-value pairs in the training candidate data, determining if the key-value pair is indicative of a particular class value; and

in response to the key-value pair being indicative of the particular class value, concluding that the key-value pair influences bias with respect to the defined class;

obtaining a first profile for a non-training candidate, wherein the first profile includes the non-training candidate's name and one or more key-value pairs identified as influencing bias in the training candidate data; and

creating a second profile for the non-training candidate from the first profile by removing or substituting the non-training candidate's name and the one or more key-value pairs identified as influencing bias.

2. The method of claim 1, wherein the determining and concluding steps further comprises:

for each of the plurality of key-value pairs, determining if the key-value pair is associated with a particular class value for more than a threshold percentage of the candidates in the training set; and

in response to a key-value pair being associated with a particular class value for more than a threshold percentage of the candidates in the training set, concluding that the key-value pair is indicative of the particular class value and influences bias with respect to the defined class.

3. The method of claim 1, wherein the training candidate data is used to train a neural network to predict a probability that a key-value pair is associated with a particular class value, and the computer system concludes that a key-value pair is indicative of a particular class value in response to the key-value pair having at least a threshold probability of being associated with the particular class value.

4. The method of claim 1, wherein the defined class of bias is one of the following: gender, race, or age.

5. The method of claim 4, wherein, in response to the defined class being gender, the classifying step comprises: using census data to identify male and female names; and classifying each training candidate in the set as female, male, or unknown based on the candidate's name.

6. The method of claim 2, wherein in response to the defined class being race, the classifying step comprises: using census data to associate names with race; and classifying each training candidate within the set with a race value based on the candidate's name.

7. The method of claim 2, wherein in response to the defined class being age, the classifying step comprises classifying each training candidate within the set with an age range using one or more school graduation dates.

8. The method of claim 1, wherein first profile is at least one of a resume, talent profile, or enhanced talent profile.

9. The method of claim 8, wherein the training candidate data has same keys as the first profile.

10. The method of claim 1, wherein the training candidates are obtained across a plurality of organizations.

11. The method of claim 1, wherein a substitute for a key-value pair is an abstraction of the key-value pair.

12. The method of claim 1, wherein defining a class of bias comprises receiving input from a user identifying the class of bias.

13. The method of claim 1, wherein the non-training candidate's address is also excluded from the second profile.

14. The method of claim 1, wherein creating the second profile further comprises: abstracting data in the first profile to relevant levels for a job role for which the non-training candidate is applying.

15. The method of claim 14, wherein abstracting data to relevant levels for the job role comprises:

identifying keys in the first profile that are relevant for the job position using candidate data across a plurality of organizations;

for each of a plurality of relevant keys, determining whether an actual value for the key relating to the job position or whether an abstracted value for the key relating to the job position; and

in creating the second profile, performing the following:

excluding key-value pairs from the second profile that correspond to keys identified as irrelevant for the job position; and

for key-value pairs corresponding to keys in which an abstracted value for the key for relating to the job position, determining whether the non-training candidate's actual value for the key is encompassed by the abstracted value that relates to and performing the following:

in response to the abstracted value encompassing the non-training candidate's actual value for the key, replacing the key-value pair with the abstracted value in the second profile; and

in response to the abstracted value not encompassing the non-training candidate's actual value for the key, replacing the key-value pair with a relevant abstracted value for the non-training candidate or excluding the key-value pair from the second profile.

16. A non-transitory computer-readable medium comprising a computer program, that, when executed by a computer system that identifies potential candidates for an open job position, enables the computer system to perform the following method for creating a profile for a job candidate that excludes data that may influence reviewer bias, the method comprising:

defining a class of potential bias;

identifying key-value pairs in candidate data that may influence bias with respect to the defined class by performing the following:

11

obtaining training candidate data representing candidate data for a set of training candidates (the "training candidate data");

for each training candidate in the set, classifying the training candidate with a value for the class;

obtaining key-value pairs in the training candidate data; for each of a plurality of key-value pairs in the training candidate data, determining if the key-value pair is indicative of a particular class value; and

in response to a key-value pair being indicative of a particular class value, concluding that concluding that the key-value pair may influence bias with respect to the defined class;

obtaining a first profile for a non-training candidate, wherein the first profile includes the non-training candidate's name and one or more key-value pairs identified as influencing bias in the training candidate data; and

creating a second profile for the non-training candidate from the first profile by removing or substituting the non-training candidate's name and the one or more key-value pairs identified as influencing bias.

17. The non-transitory computer-readable medium of claim 16, wherein the determining and concluding steps further comprises:

for each of the plurality of key-value pairs, determining if the key-value pair is associated with a particular class value for more than a threshold percentage of the candidates in the training set; and

in response to a key-value pair being associated with a particular class value for more than a threshold percentage of the candidates in the training set, concluding that the key-value pair is indicative of the particular class value and influences bias with respect to the defined class.

18. The non-transitory computer-readable medium of claim 16, wherein the training candidate data is used to train a neural network to predict a probability that a key-value pair is associated with a particular class value, and the computer system concludes that a key-value pair is indicative of a particular class value in response to the key-value pair having at least a threshold probability of being associated with the particular class value.

19. The non-transitory computer-readable medium of claim 16, wherein the defined class of bias is one of the following: gender, race, or age.

20. The non-transitory computer-readable medium of claim 19, wherein, in response to the defined class being gender, the classifying steps comprises:

using census data to identify male and female names; and classifying each training candidate in the set as female, male, or unknown based on the candidate's name.

21. The non-transitory computer-readable medium of claim 19, wherein in response to the defined class being race, the classifying step comprises:

using census data to associate names with race; and classifying each training candidate within the set with a race value based on the candidate's name.

22. The non-transitory computer-readable medium of claim 19, wherein in response to the defined class being age, the classifying step comprises classifying each training candidate within the set with an age range using one or more school graduation dates.

23. A computer system for creating a profile for a job candidate that excludes data that influence reviewer bias, the system comprising:

one or more processors;

12

one or more memory units coupled to the one or more processors, wherein the one or more memory units store instructions that, when executed by the one or more processors, cause the one or more processors to perform:

defining a class of potential bias;

identifying key-value pairs in candidate data that influence bias with respect to the defined class by performing the following:

obtaining training candidate data representing candidate data for a set of training candidates;

for each training candidate in the set, classifying the training candidate with a value for the class;

obtaining key-value pairs in the training candidate data; for each of a plurality of key-value pairs in the training candidate data, determining if the key-value pair is indicative of a particular class value; and

in response to a key-value pair being indicative of a particular class value, concluding that concluding that the key-value pair may influence bias with respect to the defined class;

obtaining a first profile for a non-training candidate, wherein the first profile includes the non-training candidate's name and one or more key-value pairs identified as influencing bias in the training candidate data; and

creating a second profile for the non-training candidate from the first profile by removing or substituting the non-training candidate's name and the one or more key-value pairs identified as influencing bias.

24. The computer system of claim 23, wherein the determining and concluding steps further comprises:

for each of the plurality of key-value pairs, determining if the key-value pair is associated with a particular class value for more than a threshold percentage of the candidates in the training set; and

in response to a key-value pair being associated with a particular class value for more than a threshold percentage of the candidates in the training set, concluding that the key-value pair is indicative of the particular class value and influences bias with respect to the defined class.

25. The computer system of claim 23, wherein the training candidate data is used to train a neural network to predict the probability that a key-value pair is associated with a particular class value, and the computer system concludes that a key-value pair is indicative of a particular class value in response to the key-value pair having at least a threshold probability of being associated with the particular class value.

26. The computer system of claim 23, wherein the defined class of bias is one of the following: gender, race, or age.

27. The computer system of claim 26, wherein, in response to the defined class being gender, the classifying steps comprises:

using census data to identify male and female names; and classifying each training candidate in the set as female, male, or unknown based on the candidate's name.

28. The computer system of claim 26, wherein in response to the defined class being race, the classifying step comprises:

using census data to associate names with race; and classifying each training candidate within the set with a race value based on the candidate's name.

29. The computer system of claim 26, wherein in response to the defined class being age, the classifying step comprises

13

classifying each training candidate within the set with an age range using one or more school graduation dates.

* * * * *

14