



# Human-Centric Multimodal Machine Learning: Recent Advances and Testbed on AI-Based Recruitment

Alejandro Peña<sup>1</sup> · Ignacio Serna<sup>1</sup> · Aythami Morales<sup>1</sup> · Julian Fierrez<sup>1</sup> · Alfonso Ortega<sup>1</sup> · Ainhoa Herrarte<sup>1</sup> · Manuel Alcantara<sup>1</sup> · Javier Ortega-Garcia<sup>1</sup>

Received: 10 October 2022 / Accepted: 9 February 2023  
© The Author(s) 2023

## Abstract

The presence of decision-making algorithms in society is rapidly increasing nowadays, while concerns about their transparency and the possibility of these algorithms becoming new sources of discrimination are arising. There is a certain consensus about the need to develop AI applications with a Human-Centric approach. Human-Centric Machine Learning needs to be developed based on four main requirements: (i) utility and social good; (ii) privacy and data ownership; (iii) transparency and accountability; and (iv) fairness in AI-driven decision-making processes. All these four Human-Centric requirements are closely related to each other. With the aim of studying how current multimodal algorithms based on heterogeneous sources of information are affected by sensitive elements and inner biases in the data, we propose a fictitious case study focused on automated recruitment: FairCVtest. We train automatic recruitment algorithms using a set of multimodal synthetic profiles including image, text, and structured data, which are consciously scored with gender and racial biases. FairCVtest shows the capacity of the Artificial Intelligence (AI) behind automatic recruitment tools built this way (a common practice in many other application scenarios beyond recruitment) to extract sensitive information from unstructured data and exploit it in combination to data biases in undesirable (unfair) ways. We present an overview of recent works developing techniques capable of removing sensitive information and biases from the decision-making process of deep learning architectures, as well as commonly used databases for fairness research in AI. We demonstrate how learning approaches developed to guarantee privacy in latent spaces can lead to unbiased and fair automatic decision-making process. Our methodology and results show how to generate fairer AI-based tools in general, and in particular fairer automated recruitment systems.

**Keywords** Automated recruitment · Bias · Biometrics · Computer vision · Deep learning · FairCV · Fairness · Multimodal · Natural language processing

## Introduction

Artificial Intelligence plays a key role in people's lives nowadays, with automatic systems being deployed in a large variety of fields, such as healthcare, education, or jurisprudence. The data science community's breakthroughs of the

last decades along with the large amounts of data currently available have made possible such deployment, allowing us to train deep models that achieve a performance never seen before. The emergence of deep learning technologies has generated a paradigm shift, with handcrafted algorithms being replaced by data-driven approaches. However, the application of machine learning algorithms built using training data collected from society can lead to adverse effects, as these data may reflect current socio-cultural and historical biases [1]. In this scenario, automated decision-making models have the capacity to replicate human biases present in the data, or even amplify them [2–6] if appropriate measures are not taken.

There are relevant models based on machine learning that have been shown to make decisions largely influenced by demographic attributes in various fields. For example,

---

Paper based on the keynote by Prof. Julian Fierrez at ICPRAM 2021.

---

This article is part of the topical collection "Pattern Recognition Applications and Methods" guest edited by Ana Fred, Maria De Marsico and Gabriella Sanniti di Baja.

---

✉ Alejandro Peña  
alejandro.penna@uam.es

Extended author information available on the last page of the article

Google's [7] and Facebook's [8] ad delivery systems generated undesirable discrimination with disparate performance across population groups. In 2016, ProPublica researchers [9] analyzed several Broward County defendants' criminal records 2 years after being assessed with the recidivism system COMPAS, finding that the algorithm was biased towards black defendants. New York's insurance regulator probed UnitedHealth Group over its use of an algorithm that researchers found to be racially biased, the algorithm prioritized healthier white patients over sicker black ones [10]. Apple Credit service granted higher credit limits to men than women even though it was programmed to be blind to that variable [11]. Face analysis technologies have also shown a gap in performance between some demographic groups [2, 12–14] as a major consequence of an undiverse representation of society in the training data. Moreover, as Balakrishnan et al. pointed out [15], the problem of data bias goes beyond the training set, as we need a bias-free evaluation set in order to correctly assess algorithmic fairness.

The usage of AI technologies is also growing in the labor market [16], where automatic decision-making systems are commonly used in different stages within the hiring pipeline [17]. However, automatic tools in this area have also exhibited worrying biased behaviors, such as Amazon's recruiting tool preferring male candidates over female ones [18]. Ensuring that all social groups have equal opportunities in the labor market is crucial to overcome differences with minority groups, which have been historically penalized [19]. Some works are starting to address fairness in hiring [20–22], but the lack of transparency (i.e., both the models and their training data are usually private for legal or corporate reasons [20]) hinders the technical evaluation of these systems.

In response to the deployment of automatic systems, along with the concerns about their fairness, the governments are adopting regulations in this matter, placing special emphasis on personal data processing and preventing algorithmic discrimination. Among these regulations, the European Union's General Data Protection Regulation (GDPR)<sup>1</sup> is specially relevant for its impact on the use of machine learning algorithms [23]. The GDPR aims to protect EU citizens' rights concerning data protection and privacy by regulating how to collect, store, and process personal data (e.g., Articles 17 and 44), and requires measures to prevent discriminatory effects while processing sensitive data (according to Article 9, sensitive data includes “personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs”). Thus, research on transparency, fairness, or explicability in machine learning is not only an ethical matter, but a legal concern and the basis for

the development of responsible and helpful AI systems that can be trusted [24].

On the other hand, one of the most active areas in Machine Learning (ML) is around the development of new multimodal models capable of understanding and processing information from multiple heterogeneous sources of information [25]. Among such sources of information we can include structured data (e.g., tabular data), and unstructured data from images, audio, and text. The implementation of these models in society must be accompanied by effective measures to prevent algorithms from becoming a source of discrimination. In this scenario, where multiple sources of both structured and unstructured data play a key role in algorithms' decisions, the task of detecting and preventing biases becomes even more relevant and difficult.

In this environment of desirable fair and trustworthy AI, the main contributions of this work are:

- We review the latest advances in Human-Centric ML research with special focus on the public available databases proposed by the community.
- We present a new public experimental framework around automated recruitment, aimed to study how multimodal machine learning is influenced by demographic biases present in the training datasets: FairCVtest.<sup>2</sup>
- We have evaluated the capacity of both pre-trained models and data-driven technologies to extract demographic information and learn biased target functions from multimodal sources of information, including images, texts, and structured data from resumes.
- We evaluated a discrimination-aware learning method based on the elimination of sensitive information such as gender or ethnicity from the learning process of multimodal approaches, and apply it to our automatic recruitment testbed for improving fairness among demographic groups.

Our results demonstrate the high capacity of commonly used learning methods to expose sensitive information (e.g., gender and ethnicity) from different data domains, and the necessity to implement appropriate techniques to guarantee discrimination-free decision-making processes.

A preliminary version of this article was published in [26]. This article significantly improves [26] in the following aspects:

- We extend FairCVdb to incorporate a name and a short biography to each profile. To the best of our knowledge, this upgrade makes FairCVdb the first fairness research database including image, text and structured data.

<sup>1</sup> <https://gdpr.eu/>.

<sup>2</sup> <https://github.com/BiDALab/FairCVtest>.

- We provide more extensive experiments within FairCVtest, where we analyze the impact of data bias on an automatic recruitment tool under different scenarios. In these experiments, we use commonly used fairness criteria to quantify this impact. We also measure the sensitive information exploited in the decision-making process, whereas [26] limited the experiments to a more qualitative analysis. Furthermore, by including text data to our dataset, we extend FairCVtest with Natural Language Processing techniques.
- We provide a survey on fairness research in AI, in which we review some of the methods proposed in recent years to prevent algorithmic discrimination, and the most commonly used databases in the field.

The rest of the paper is structured as follows: “[Human-Centric Research in Machine Learning](#)” presents an overview on explainability in ML models, discrimination-aware ML approaches, and Human-Centric ML databases. “[FairCVdb: Dataset for Multimodal Bias Research](#)” describes the considered automatic hiring pipeline, examines the information available in a resume highlighting the sensitive data associated to it, and describes the dataset created in this work: FairCVdb. “[General Learning Framework](#)” presents the general framework for our work including problem formulation. “[Experiments and Results](#)” reports the experiments in our testbed FairCVtest after describing the experimental methodology and the different scenarios evaluated. Finally, “[Conclusions](#)” summarizes the main conclusions.

## Human-Centric Research in Machine Learning

The recent advances in AI and the large amounts of data available have made possible the deployment of automatic decision-making algorithms in our society. Due to their great impact in people’s lives, especially in high stake settings, is essential that these systems are responsible and trustworthy. However, there are many models that have been shown to make decisions based on attributes considered as private (e.g., gender<sup>3</sup> and ethnicity), or exhibiting systematically discrimination against individuals belonging to disadvantaged groups. We can find examples of such unfair treatment in various fields, such as healthcare [10, 29], ad delivery

<sup>3</sup> We are aware of the studies that move away from the traditional view of gender as a binary variable [27], and the difference between gender identity and biological sex. Despite the limitations of such model [28], in this paper we use “gender” to refer to the external perception of biological sex, in line with the work historically developed in gender classification into male and female individuals.

systems [7, 8, 30], hiring [16, 18], and both facial analysis [5, 12, 13] and NLP technologies [31, 32].

In the following sections, we will present recent advances in Human-Centric ML research related with: (1) explainability and interpretability of ML models; (2) discrimination-aware ML approaches; and (3) databases for Human-Centric ML research.

## Interpretable and Explainable ML

One of the long-term goals in deep learning is to learn abstract representations, which are generally invariant to local changes in the input [33]. It has been observed that many learned representations correspond to human-interpretable concepts. But it is not quite clear what function they serve and whether it has a causal role that reveals how the network models its higher-level notions [34]. Research is showing that not all representations in the convolutional layers of a DNN correspond to natural parts, raising the possibility of a different decomposition of the world than humans might expect, calling for further study into the exact nature of the learned representations [35, 36].

There is significant work on understanding neural networks. Most methods typically focus on what a network looks at when making a decision [37, 38]; other approaches seek to train explanatory models [39] or networks [40] that generate human-readable text.

We can distinguish between two types of approaches for generating a better understanding of an AI model: interpretable and explainable. As defined in [41], an interpretation is the mapping of an abstract concept (e.g., a predicted class) into a domain that the human can make sense of, e.g., images or text; and an explanation is the collection of features of the interpretable domain that have contributed for a given example to produce a decision.

On the interpretation side, we have Activation Maximization, which consists of looking for an input pattern that produces a maximum response of the model. It was introduced in [42], but such visualization technique has a limitation: as complexity increases, it becomes more difficult to find a simple representation of a higher layer unit, because the optimization does not converge to a single global minimum. Simonyan et al. came up with the suggestion to perform the optimization with respect to the input image, obtaining an artificial image representative of the class of interest [43].

One way of improving activation maximization to enable enhanced visualizations of learned features is with the so-called expert. That is, in the function to be maximized, the L2-norm regularizer (a term that penalizes inputs that are farther away from the origin) is replaced by a more sophisticated one, called expert [35, 44, 45]. Another way is via deep generative models, incorporating such model in the activation maximization framework [46, 47].

On the explanation side, we have Sensitivity Analysis: how much do changes in each pixel affect the prediction. Initially intended for pruning neural networks and reducing the dimensionality of their input vector, was particularly useful for understanding the sensitivity of performance with respect to their structure, parameters, and input variables [48, 49]. More recently, it has been used for explaining the classification of images by deep neural networks. Simonyan et al. [43] applied partial derivatives to compute saliency maps. They show the sensitivity of each of the input image pixels, where the sensitivity of a pixel measures to what extent small changes in its value make the image to belong more or less to the class (local explanation).

Alternatives for explaining deep neural network predictions are backward propagation techniques. Some are: deconvolution, layer-wise relevance propagation (LRP) and guided backprop.

Zeiler and Fergus [50] proposed deconvolution to compute a heatmap showing which input pattern originally caused a certain activation in the feature maps. The idea behind the deconvolution approach is to map the activations from the network back to pixel space using a backpropagation rule. The quantity being propagated can be filtered to retain only what passes through certain neurons or feature maps.

The LRP method [37] applies a propagation rule that distributes back (without gradients) the classification output  $f(x)$  decomposed into pixel relevance onto the input variables. This algorithm can be used to visualize the contribution of pixels both for and against a given class.

Guided backprop is the extension of the deconvolution approach for visualizing features learned by CNNs. Proposed in [51], it combines backpropagation and deconvolution by masking out the values for which at least one of the entries of the top gradient (deconvnet) or bottom data (backpropagation) is negative.

Another very well-known backpropagation-based method combining gradients, network weights, and activations is Grad-CAM [38]. Gradient-weighted Class Activation Mapping (Grad-CAM) uses the gradients of the class score with respect to the input image to produce a coarse localization map highlighting the important regions in the image for predicting the concept. It can be combined with guided backpropagation for fine-grained visualizations of class-discriminative features.

These methods selectively illustrate one of the multiple patterns a filter represents, explanatory graphs provide a workaround. [52] proposed a method disentangling part patterns from each filter to represent the semantic hierarchy hidden inside a CNN.

Some other methods have gone beyond visualization of CNNs and diagnosed CNN representations to gain a deep understanding of the features encoded in a CNN. Others

report the inconsistency of some widely deployed saliency methods, as they are not independent of both the data on which the model was trained and the model parameters [53].

Szegedy et al. [54] reported the existence of blind spots and counter-intuitive properties of neural networks. They found that it is possible to change the network's prediction by applying an imperceptible optimized perturbation to the input image, which they called an adversarial example. Paving the way for a series of works that sought to produce images with which to fool the models [55–57].

Other studies aiming to understand deep neural networks are neuron ablation techniques. These seek a complete functional understanding of the model, trying to elucidate its inner workings or shed light on its internal representations. Bau et al. found evidence for the emergence of disentangled, human-interpretable units (of objects, materials and colors) during training [34].

## Discrimination-Aware Learning

In order to prevent automated systems from making decisions based on protected attributes or reproduce biased behaviors against disadvantaged groups, the research community has devised various ways to improve fairness in AI systems. These approaches are usually divided in the literature between pre-processing, in-processing, and post-processing techniques [24].

The pre-processing techniques aim to transform the input domain to prevent discrimination and remove sensitive information. The authors of [58] propose to remove sensitive information while improving model interpretability by learning a data-to-data transformation in the input domain, where the new representation achieves certain fairness criterion. This transformation is based in both neural style transfer and kernel Hilbert spaces. A similar approach is proposed in [59], which seeks to generate a new dataset similar to a given one, but fairer with respect to a certain protected attribute. For this purpose, a fairness criterion is added to the loss function of an auxiliary GAN [60]. In [61], the authors address the pre-processing transformation as an optimization problem which trades off discrimination and utility at probabilistic level, while controlling sample distortion on an individual level. More recently, Ramaswamy et al. proposed [62] a method for augmenting real datasets with GAN-generated synthetic images by modifying vectors in the GAN latent space to de-correlate sensitive and target attributes.

In-processing approaches focus on the learning process as the key point to prevent biased models, by changing the optimization objective or imposing fairness constraints. In [63], the authors propose an adaptation of Domain Adaptation Neural Networks [64] to generate agnostic feature representations, unbiased related to certain protected attribute. Also based in domain adaptation, in [65], the authors reduce

racial biases in face recognition using mutual information and unsupervised domain adaptation, from a labeled domain (i.e., Caucasian individuals) to an unlabeled one (i.e., non-Caucasian individuals). A method to mitigate bias in occupation classification without having access to protected attributes is developed in [66], by reducing the correlation between the classifier's output for each individual and the word embeddings of their names. Wang and Deng studied in [13] the use of an adaptive margin in large margin face recognition loss functions [67] to reduce the gap in performance between different ethnicity groups. They proposed to use deep Q-learning to adaptively find the margin for each demographic group during training.

More recently, in-processing approaches based on adversarial learning frameworks [68] have been explored. A joint learning and unlearning method is proposed in [69] to simultaneously learn the main classification task while unlearning biases by applying confusion loss, based on computing the cross entropy between the output of the best bias classifier and a uniform distribution. The authors of [70] introduced a new regularization loss based on mutual information between feature embeddings and bias, training the networks using adversarial and gradient reversal [64] techniques. In [71] an extension of triplet loss [72] is applied to remove sensitive information in feature embeddings, without losing performance in the main task.

Finally, post-processing methods assume that the output of the model may be biased, so they apply a transformation on it to improve fairness between demographic groups. Some works in this line have proposed to prevent unfairness using discrimination-aware data-mining [73, 74]. In [75], the authors propose a framework that enables a human manager to select how to make the trade-off among fairness and utility. Then, the method selects a threshold for each demographic group to obtain an optimal classifier according to the manager's preferences. Post-processing techniques are also common among studies on fairness in ranking [76–78], which are close to our work here.

## Databases

The datasets used for learning or inference may be the most critical elements of the machine learning process where bias can appear. As these data are collected from society, they may reflect socio-cultural biases [1], or reflect an unbalanced representation of the different demographic groups composing it. A naive approach would be to remove all sensitive information from data, but this is almost infeasible in a general AI setup (e.g., [31] demonstrates how removing explicit gender indicators from personal biographies is not enough to remove the gender bias from an occupation classifier, as other words may serve as “proxy”). On the other hand, collecting large datasets that represent broad social diversity in

a balanced manner can be extremely costly, and not enough to avoid disparate treatment between groups [13].

The biases introduced in the dataset used to train machine learning models typically reflect human biases present in society, or are related to an inaccurate representation of groups [89, 90]. In view of this situation, the scientific community has put lots of effort into collecting databases that improve the representation of different demographic groups, which can be used to suppress the presence of bias. In this section, we discuss some of the most commonly used databases in AI fairness research, either because of the biases they present, or their absence (i.e., databases more balanced in terms of certain demographic attributes). Table 1 provides an overview of these databases, including the number of samples, data modality and the demographic attributes studied with each one. The **Adult Income** dataset [79] from the UCI repository is frequently used on gender and ethnicity bias mitigation. The main task of the database is predict whether a person will earn more or less than \$50K per year. The database includes 48,842 samples with 14 numerical/categorical attributes each, such as education level, capital-gain or occupation, and missing values.

The **German Credit** dataset [79] contains 1K entries with 20 different categorical/numerical attributes, where each entry represents a loan applicant by a bank. The applicants are classified as good or bad risk credit, showing age bias toward young people. Also related to age biases, the **Bank Marketing** database [80] contains marketing campaign data of a Portuguese bank institution. With more than 41K samples, the goal is to predict if the client will subscribe a term deposit, based on 20 categorical/numerical attributes including personal data and socioeconomic contextual information.

The **ProPublica Recidivism** dataset [9] provides more than 11K pretrial defendants records, assessed with the COMPAS algorithm to predict their likelihood of recidivism. After a 2-year study, the researchers find out that the algorithm was biased towards African-Americans, showing both higher false positive and lower false-negative rates than white defendants.

In the study of demographic bias in NLP technologies,<sup>4</sup> we can cite the **Common Crawl Bios** dataset [31], which contains nearly 400K short biographies collected from Common Crawl. The goal of the dataset is to predict the occupation from these bios, out of 28 possible occupations showing high gender imbalances. The dataset also provides a “gender blinded” version of each bio, where

<sup>4</sup> There are several works that study demographic biases in word embeddings [32, 91], working with representation spaces trained with large corpus of texts from Wikipedia, Common Crawl or Google News, among other sources.

**Table 1** Summary of the most common public databases for AI fairness and bias research

Database	#Samples	Image	Text	Cat./num.	Demographic	Access
UCI Adult Income [79]	48.8K	✗	✗	✓	Ethnicity, gender	archive.ics.uci.edu/ml/datasets/adult
German Credit [79]	1K	✗	✗	✓	Age	archive.ics.uci.edu/ml/datasets/statlog+(german+credit+data)
Bank Marketing [80]	41.1K	✗	✗	✓	Age	archive.ics.uci.edu/ml/datasets/Bank+Marketing
ProPublica Recidivism [9]	11.8K	✗	✗	✓	Ethnicity	github.com/propublica/compas-analysis
Common Crawl Bios [31]	397K	✗	✓	✗	Gender	github.com/microsoft/biosbias
WinoBias [81]	3.2K	✗	✓	✗	Gender	github.com/uclanlp/corefBias
CelebA [82]	202.6K	✓	✗	✓	Gender	mmlab.ie.cuhk.edu.hk/projects/CelebA.html
IMDB-WIKI [83]	523K	✓	✗	✗	Age, gender	data.vision.ee.ethz.ch/cvl/rrothe/imdb-wiki/
Cleaned IMDB [69]	140K	✓	✗	✗	Age, gender	robots.ox.ac.uk/vgg/data/laofiw/
MORPH [84]	55K	✓	✗	✗	Age, ethnicity, gender	ebill.uncw.edu/C20231_ustores/web/product_detail.jsp?PRODUCTID=8
PPB [12]	1.3K	✓	✗	✗	Ethnicity, gender	gendershades.org/index.html
LAOFIW [69]	14K	✓	✗	✗	Ethnicity, gender	robots.ox.ac.uk/vgg/data/laofiw/
FairFace [85]	108.5K	✓	✗	✗	Age, ethnicity, gender	github.com/jojojs/fairface
Diversity in Faces [86]	1M	✓	✗	✗	Age, ethnicity, gender	research.ibm.com/artificial-intelligence/trusted-ai/
DiveFace [71]	120K	✓	✗	✗	Ethnicity, gender	github.com/BiDALab/DiveFace
BFW [87]	20K	✓	✗	✗	Ethnicity, gender	github.com/visionjo/facerec-bias-bfw
DemogPairs [88]	10.8K	✓	✗	✗	Gender, ethnicity	download.hertasecurity.com/research/DemogPairs.zip
RFW [65]	40K	✓	✗	✗	Ethnicity	whdeng.cn/RFW/testing.html
BUPT-B [13]	1.3M	✓	✗	✗	Ethnicity	whdeng.cn/RFW/Trainingdataste.html
BUPT-G [13]	2M	✓	✗	✗	Ethnicity	whdeng.cn/RFW/Trainingdataste.html
FairCVdb (Ours)	24K	✓	✓	✓	Ethnicity, gender	github.com/BiDALab/FairCVtest

We specify the different modalities included in each dataset (i.e., images, texts, and categorical/numerical attributes), along with the demographic attributes typically studied with each one

explicit gender indicators have been removed (e.g., pronouns or names). On a closely related task, the **WinoBias** database [81] provides 3160 sentences, where the goal is to find all the expressions related to certain entity. Centered in people entities referred by their occupations, the dataset requires to link gender pronouns to male/female stereotypical occupations.

We now focus in face datasets, which are the basis for different face analysis task such as face recognition or gender classification. The **CelebA** database [82] contains nearly 202.6K images from more than 10K celebrities. Each image is annotated with 5 facial landmarks, along with 40 binary attributes including appearance features, demographic information, or attractiveness, which shows a strong gender bias.

The **IMDB-WIKI** dataset [83] provides 460.7K images from the IMDB profiles of 20,284 different celebrities, along with 62.3K images from Wikipedia. Images were labeled using the information available in the profiles (i.e., name, gender, and birth date), extracting an age label by comparing the timestamp of the images and the birth date. The dataset presents a gender bias in the age distributions, as we encounter younger females and older males. Due to the image acquisition process, some labels are noisy, so the

authors of [69] released the **cleaned IMDB** dataset, with 60K cleaned images for age prediction and 80K for gender classification obtained from the IMDB split.

Also related with age studies, the **MORPH** database [84] provides 55K images from 13K individuals, aimed to study the effect of age-progression on different facial tasks. The database is longitudinal with age, having pictures of the same user over time. The database is strongly unbalanced with respect to gender and ethnicity, with 65% images belonging to African-American males.

Some databases aim to mitigate biases in face analysis technologies by putting emphasis in demographic balance and diversity. **Pilot Parliaments Benchmark (PPB)** [12] is a dataset of 1270 parliamentarians images from 6 different countries in Europe and Africa. The images are balanced with respect to gender and skin color, which are available as labels (the skin color is codified using the six-point Fitzpatrick system). The **Labeled Ancestral Origin Faces in the Wild (LAOFIW)** dataset [69] provides 14K images manually divided into 4 ancestral origin groups. The database is balanced with respect to ancestral origin and gender, and a variety of pose and illumination. Also emphasizing ethnicity balance, the **FairFace** database [85] contains more

than 100K images equally distributed in 7 ethnicity groups (White, Black, Indian, East Asian, Southeast Asian, Middle East, and Latino), also providing gender and age labels. Aimed to study facial diversity, **Diversity in Faces** [86] provides 1 M images annotated with 10 different facial systems including gender, age, skin color, pose, and facial contrast labels, among others.

If we look at face recognition databases, **DiveFace** [71] contains face images equitably distributed among 6 demographic classes related to gender and 3 ethnic groups (Black, Asian, and Caucasian), including 24K different identities and a total of 120K images. The **DemogPairs** database [88] also proposes 6 balanced demographic groups related to gender and ethnicity, each one with 100 subjects and 1.8K images. On his part, the **Balanced Faces in the Wild (BFW)** database [87] presents 8 demographic groups related with gender and 4 ethnicity groups (Asian, Black, Indian and White), each one with 100 different users and 2.5K images. Finally, Wang and Deng proposed three different databases based on MS-Celeb-1 M [92], namely **Racial Faces in the Wild (RFW)** [65], **BUPT-B** [13] and **BUPT-G** [13]. While RFW is designed as a validation dataset, aimed to measure ethnicity biases, both BUPT-B and BUPT-G are proposed as ethnicity-aware training datasets. RFW defines 4 ethnic groups (Caucasian, Asian, Indian, and African), each one with 10K images and 3 different subjects. On the other hand, both BUPT-B and BUPT-G propose the same ethnic groups, the first one almost ethnicity-balanced with 1.3 M images and 28K subjects, while the latter contains 2 M images and 38K subjects, which are distributed approximating the world's population distribution.

## FairCVdb: Dataset for Multimodal Bias Research

### AI in Hiring Processes

The usage of predictive tools in recruitment processes is increasing nowadays. Employers have adopted these tools in an attempt to reduce the time and cost of hiring, or to maximize the quality of the hiring process, among other reasons [16]. Rather than a single-point decision, the hiring pipeline suppose a multi-stage process, which can be broadly divided in four stages [16]. In the *sourcing* stage, the employers attract potential candidates through advertisements or job posting. Then, during *screening*, the employers assess the applicants to choose a subset to *interview* individually. Finally, employers make a final decision (i.e., whether to hire or reject each applicant) in the *selection* stage. All of these stages can benefit from the use of

automatic algorithms,<sup>5</sup> as well as suffer from algorithmic discrimination if systems are not carefully designed. The labor market has a long history of unfair treatment of minority groups [19, 93], which makes bias prevention a crucial step in automatic hiring tools design. Although the study of fairness in algorithmic hiring has been limited [21], some works are starting to address this topic [20, 22, 94].

For the purpose of studying discrimination in Artificial Intelligence at large, and particularly in hiring processes, in this work, we propose a new experimental framework inspired in a fictitious automated recruiting system: FairCVtest. Our work can be framed within the screening stage of the hiring pipeline, where an automatic tool determines a score from a set of applicants resumes. We chose this application because it comprises personal information from different nature [95].

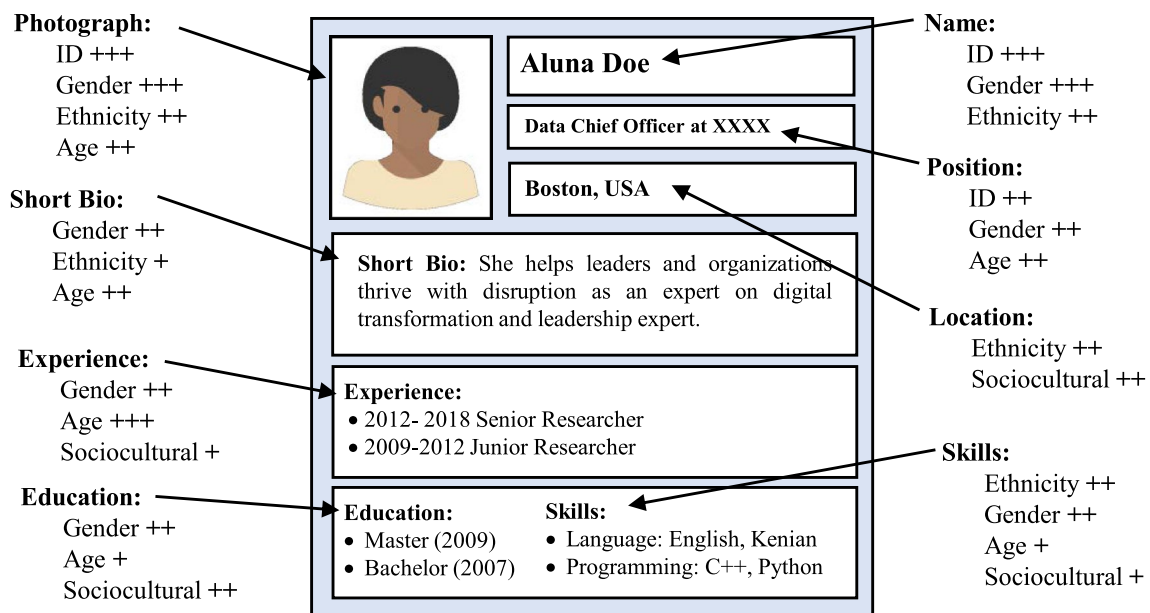
The resume is traditionally composed by structured data including name, position, age, gender, experience, or education, among others (see Fig. 1), and also includes unstructured data such as a face photo or a short biography. A face image is rich in unstructured information such as identity, gender, ethnicity, or age [96, 97]. That information can be recognized in the image, but it requires a cognitive or automatic process trained previously for that task. The text is also rich in unstructured information. The language and the way we use that language, determine attributes related to your nationality, age, or gender. Both, image and text, represent two of the domains that have attracted major interest from the AI research community during last years. The Computer Vision and the Natural Language Processing communities have boosted the algorithmic capabilities in image and text analysis through the usage of massive amounts of data, large computational capabilities (GPUs), and deep learning techniques.

The resumes used in the proposed FairCVtest framework include merits of the candidate (e.g., experience, education level, languages, etc.), two demographic attributes (gender and ethnicity), and a face photograph (see “[FairCVdb: Dataset Description](#)” for all the details).

### FairCVdb: Dataset Description

In this work, we present FairCVdb, a new dataset with 24,000 synthetic resume profiles for both fairness and multimodal research in AI. Each profile includes 2 demographic attributes (gender and ethnicity), an occupation, a face image, a name, 7 attributes obtained from 5 information blocks that are usually found in a standard resume, and a short biography. The profiles comprise data from different nature including structured and unstructured data:

<sup>5</sup> <https://www.hirevue.com/>.



**Fig. 1** Information blocks in a resume and personal attributes that can be derived from each one. The number of crosses represent the level of sensitive information (+++ = high, ++ = medium, + = low)

- Demographic attributes (*structured data*): Each profile has been generated according to two gender classes and three ethnicity classes. These demographic attributes determine the face image (gender and ethnicity related), name (gender related), and pronouns in the short biography (gender related).
- Face image (*unstructured data-image*): Each profile contains a real and unique face image assigned from the DiveFace database [71], which was introduced in “Databases”. DiveFace<sup>6</sup> contains face images from 24K different identities with their corresponding gender and ethnicity attributes.
- Short Biography (*unstructured data-text*): We use the Common Crawl Bios dataset [31] to associate a short biography, a name, and an occupation (from a pool of 10 different occupations) to each profile.
- Candidate competencies (*structured data*): The five information blocks are: (1) education attainment, (2) availability, (3) previous experience, (4) the existence of a recommendation letter, and (5) language proficiency in a set of three different and common languages. Each language is encoded with an individual feature (3 features in total) that represents the level of knowledge in that language. We will refer to these resume features as candidate competencies.

As we previously mentioned in “Databases”, the Common Crawl Bios dataset<sup>7</sup> [31] contains online biographies collected from Common Crawl relating 28 different occupations. Gender and occupation labels are available for each biography, as well as a “blinded” version of the bio, in which explicit gender indicators have been removed. For example, a biography labeled as [Attorney, Female] is presented as: *Andrea Jepsen is an attorney with the School Law Center, a law firm focusing on the rights of students and families in education and school law disputes. She has worked with people with disabilities since 1997 in a variety of roles, including as an early childhood special education service coordinator, and as a legal services provider working regularly in the courts and in administrative proceedings. Ms. Jepsen’s broad legal experience has involved representing clients in a variety of critical legal issues related to education, housing, elder law matters, public benefits, family law disputes, probate and other concerns.* Note that we underlined explicit gender indicators removed in the “blinded bio”, and that both name and occupation can be found in the first sentence of each biography, so this sentence was not included in the bios.

We select 24K biographies, and its corresponding blinded versions, from a subset of 10 different occupations. Each biography is associated according to gender to one FairCV profile, providing as well an occupation label and a name to the profiles, which we obtain by processing the first sentence

<sup>6</sup> <https://github.com/BiDALab/DiveFace>.

<sup>7</sup> <https://github.com/microsoft/biosbias>.



of each bio. We group the occupations in four professional sectors: (1) audiovisual communication and journalism, with *journalist*, *photographer*, and *filmmaker*; (2) administration and jurisdiction, with *attorney* and *accountant*; (3) healthcare, with *surgeon*, *nurse*, and *physician*; and (4) education, with *professor* and *teacher*. Each professional sector has the same number of samples (i.e., 6K bios), and is gender-balanced. Furthermore, we define a suitability attribute ( $S$ ), representing the affinity degree of each sector with the potential job to which the resumes apply. The association of this attribute with each sector has purely academic purposes, without seeking to state the usefulness or importance of each of them.

The score  $T^j$  for a profile  $j$  is generated by linear combination of the candidate competencies  $\mathbf{x}^j = [x_1^j, \dots, x_n^j]$  and the suitability attribute  $S^j$  as

$$T^j = \beta^j + \sum_{i=1}^n \alpha_i x_i^j + \alpha_s S^j \quad (1)$$

where  $n = 7$  is the number of features (competencies),  $\alpha_i$  are the weighting factors for each competency  $x_i^j$  (fixed manually based on consultation with a human recruitment expert), and  $\beta^j$  is a small Gaussian noise to introduce a small degree of variability (i.e., two profiles with the same competencies do not necessarily have to obtain the same result in all cases). Those scores  $T^j$  will serve as ground truth in our experiments.

Note that, by not taking into account gender or ethnicity information during the score generation in Eq. (1), these scores become agnostic to this information, and should be equally distributed among different demographic groups. Thus, we will refer to this target function as Unbiased scores  $T^U$ , from which we define two target functions that include two types of bias: Gender bias  $T^G$  and Ethnicity bias  $T^E$ . Biased scores are generated by applying a penalty factor  $T_\delta$  to certain individuals belonging to a particular demographic group. For the Gender-biased scores  $T^G$ , we apply a penalty factor on the female group, while in the Ethnicity-biased scores  $T^E$ , we apply the penalty factor to one ethnic group, and the inverse to another one (i.e., the individuals belonging to this group are overrated in  $T^E$ , showing a higher score than in  $T^U$ ). This leads to a set of scores where, with the same competencies, certain groups have lower scores than others, simulating the case where the process is influenced by certain cognitive biases introduced by humans, protocols, or automatic systems.

Table 2 summarizes the features that make up each profile, as well as their labels. We divided the FairCVdb in two splits, with 80% of the synthetic profiles (19,200 CVs) as training set, and the remaining 20% (4800 CVs) as validation set. Both sets are almost perfectly balanced among gender, ethnicity and professional sector. Figure 2 presents four visual examples of the resumes generated with FairCVdb.

**Table 2** Overview of the different attributes available in each FairCV profile

Name	Type	Data values
Education	$I$	$x_1 \in \{0.2, 0.4, 0.6, 0.8, 1\}$
Recommendation	$I$	$x_2 \in \{0, 1\}$
Availability	$I$	$x_3 \in \{0.2, 0.4, 0.6, 0.8, 1\}$
Previous experience	$I$	$x_4 \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$
Language proficiency	$I$	$x_i \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$ , $i \in \{5, 6, 7\}$
Face image	$I$	$I[m, n] / m, n \in [0, 119]$
Face embedding	$I$	$\mathbf{f} \in \mathbb{R}^{20} / \ \mathbf{f}\  = 1$
Agnostic face embedding	$I$	$\mathbf{f}_a \in \mathbb{R}^{20} / \ \mathbf{f}_a\  = 1$
Name	$I$	Text data
Biography	$I$	Text data
Agnostic biography	$I$	Text data
Gender	$IIT$	$G \in \{0, 1\}$
Ethnicity	$IIT$	$E \in \{0, 1, 2\}$
Occupation	$IIT$	$O : \mathbb{N} \in [0, 11]$
Suitability	$IIT$	$S \in \{0.25, 0.5, 0.75, 1\}$
Blind score	$T$	$T^U : \mathbb{R} \in [0, 1]$
Gender-biased score	$T$	$T^G : \mathbb{R} \in [0, 1]$
Ethnicity-biased score	$T$	$T^E : \mathbb{R} \in [0, 1]$

We include the possible values of each attribute, as well as its nature as Input and/or Target

## FairCVtest: Description

### General Learning Framework

The multimodal model represented by its parameters vector  $\mathbf{w}^F$  ( $F$  for fused model [95]) is trained using features learned by  $M$  independent models  $\{\mathbf{w}^1, \dots, \mathbf{w}^M\}$  where each model produces  $n_i$  features  $\mathbf{x}^i = [x_1^i, \dots, x_{n_i}^i] \in \mathbb{R}^{n_i}$ . Without loss of generality, the Fig. 3 presents the learning framework for  $M = 3$ . The learning process is guided by a Target function  $T$ , and a learning strategy that minimizes the error between the output  $O$  and the Target function  $T$ . In our framework, where  $\mathbf{x}^i$  is data obtained from the resume,  $\mathbf{w}^i$  are models trained specifically for different information domains (e.g., images, text) and  $T$  is a score within the interval  $[0, 1]$  ranking the candidates according to their merits. A score close to 0 corresponds to the worst candidate, while the best candidate would get 1. The learning strategy is traditionally based on the minimization of a loss function defined to obtain the best performance. The most popular approach for supervised learning is to train the model  $\mathbf{w}$  by minimizing a loss function  $\mathcal{L}$  over a set of training samples  $\mathcal{S}$ :

$$\min_{\mathbf{w}^F} \sum_{\mathbf{x}^j \in \mathcal{S}} \mathcal{L}(O(\mathbf{x}^j | \mathbf{w}^F), T^j). \quad (2)$$

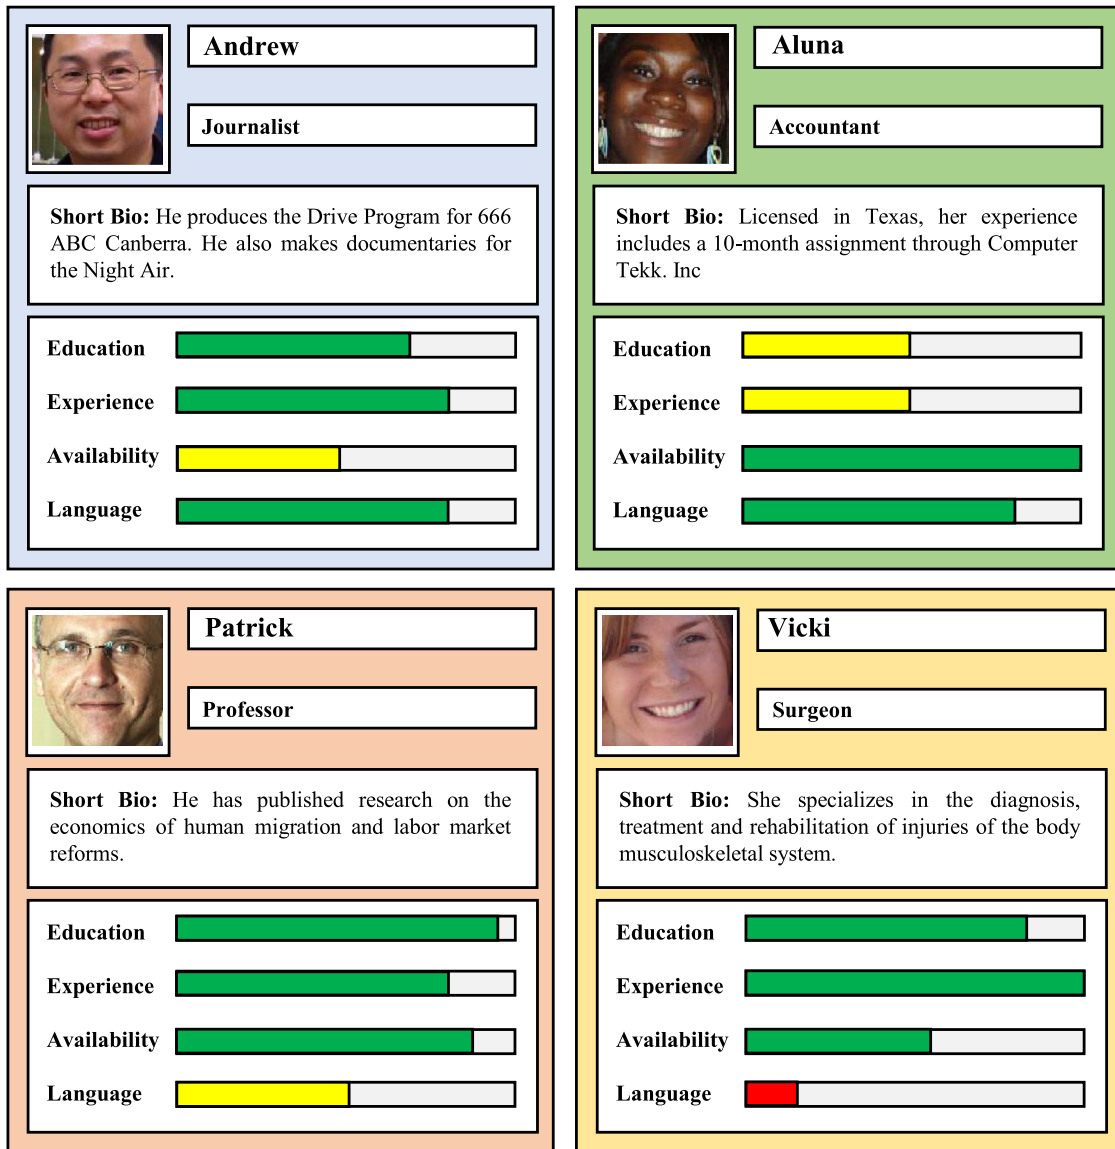


Fig. 2 Visual examples of the FairCVdb synthetic resumes, including a face image, a name, an occupation, a short biography and the candidate competencies

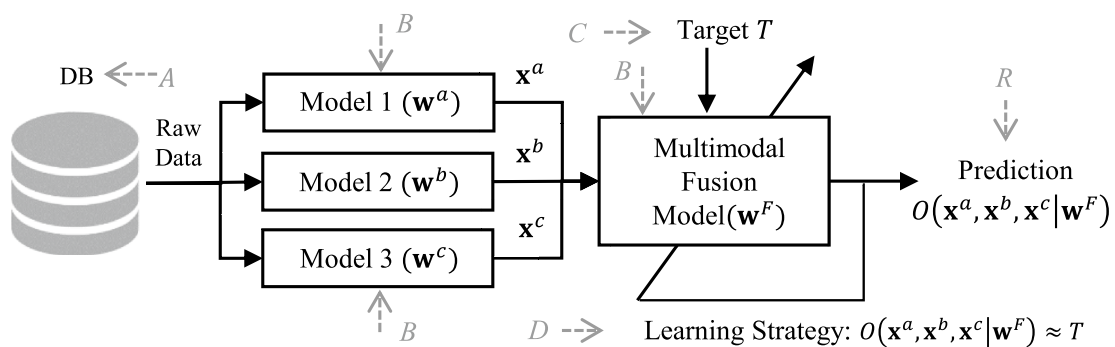
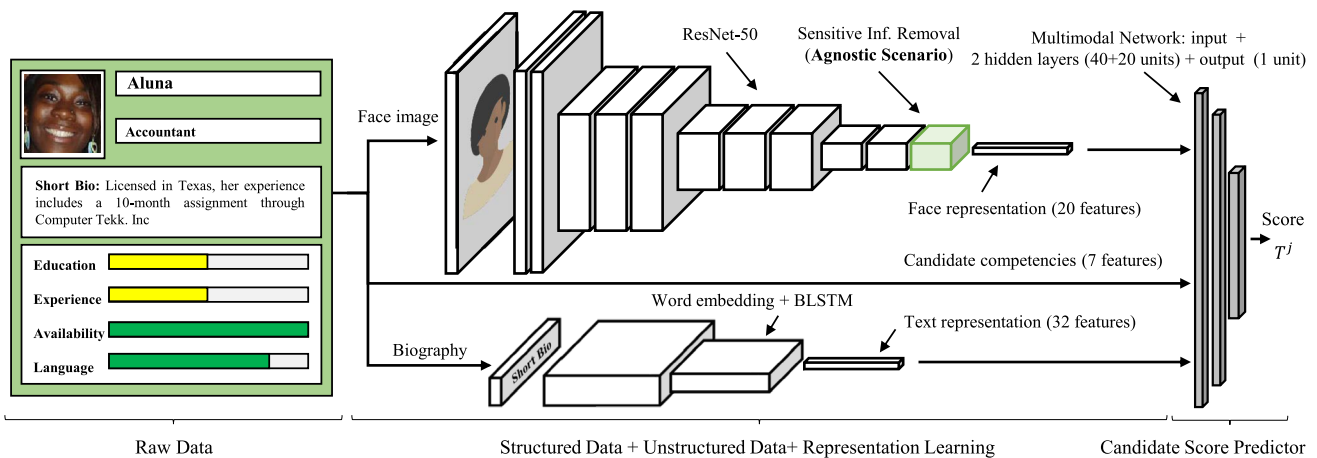


Fig. 3 Block diagram of the automatic multimodal learning process and 6 (A to E) stages where bias can appear



**Fig. 4** Multimodal learning architecture, composed by a Convolutional Neural Network (ResNet-50 [98]), a BLSTM and a fully connected network used to fuse the features from different domains

(image, text and structured data). Note that, in the agnostic scenario, we include a sensitive information removal module [71] over the ResNet network to generate agnostic face embeddings

Biases can be introduced in different stages of the learning process (see Fig. 3): in the Data used to train the models (A), the Preprocessing or Feature generation (B), the Target function (C), and the Learning strategy (D). As a result of the biases introduced at of these points (A to D), we may obtain biased Results (R). In this work, we focus on the Target function (C) and the Learning strategy (D). The Target function is critical as it could introduce cognitive biases from biased processes.

**FairCVtest: Multimodal Learning Architecture for Automatic CV Analysis**

Figure 4 summarizes the learning architecture proposed to study the different scenarios of FairCVtest. We designed the candidate score predictor as a multimodal neural network with three input branches: (i) face image, (ii) text biography, and (iii) candidate competencies. The learning architecture includes two specific models to process the face image and text data from the biography, before fusing the information from all three modalities.

**Face Analysis Model**

We use the face image from each profile, and the pre-trained model ResNet-50 [98] as feature extractor to obtain feature embeddings from the applicants’ face attributes. ResNet-50 is a popular Convolutional Neural Network composed with 50 layers including residual or “shortcuts” connections to improve accuracy as the net depth increases (i.e., solving the “vanishing gradient” problem). ResNet-50’s last convolutional layer outputs embeddings with 2048 features, so we added a fully connected layer to perform a bottleneck that

compresses these embeddings to just 20 features (maintaining the competitive face recognition performance), so that its size approximates to that of the candidates competencies. Note that our face model was trained exclusively for the task of face recognition. However, although gender or ethnicity information were not intentionally employed during the training process, this information is part of the face attributes. Therefore, an AI system trained on these face embeddings could detect the protected attributes without being explicitly trained for this task.

**Text Analysis Model**

The second branch is aimed to extract a text representation from the bios, using a bidirectional LSTM layer composed by 32 units and hyperbolic tangent activation. This branch receives as input a sequence of word vectors. We use the fastText<sup>8</sup> word embeddings [99] to represent each word in the biographies as 300-dimensional word vectors. Note that these word vectors were trained on a different Common Crawl subset than the one used to extract the biographies of [31].

**Multimodal Model**

The face and text features obtained from its respective models are combined with the candidate competencies to feed the multimodal network. This network is composed by two hidden layers, with 40 and 20 neurons respectively and ReLU activation, and only one neuron with sigmoid activation in the output layer. Note that, as the target functions  $T$

<sup>8</sup> <https://fasttext.cc/docs/en/english-vectors.html>.

in FairCVdb are real valued scores within the interval  $[0, 1]$ , we treat this task as a regression problem. A binary classifier can be obtained by thresholding the predicted scores (i.e., switching from a scoring tool to a selection tool), as we will show in “[Fairness in Recruitment Tools: Learning Demographic Parity](#)”.

### Privacy-Enhancing Representation Learning

With the aim of generating another representation, agnostic with regard to gender and ethnicity, we use the method proposed in [71], called SensitiveNets. This method was proposed to improve the privacy in face biometrics, by incorporating an adversarial regularizer capable of removing sensitive information from pre-trained feature embeddings without losing performance in the main task. Thus, two different face representations are available for each profile, one containing gender and ethnicity sensitive information, and a second one “blind” or agnostic to these attributes. In order to remove sensitive information from the learned space, Eq. 2 is replaced by

$$\min_{\mathbf{w}^F} \sum_{\mathbf{x}^j \in S} \mathcal{L}(O(\mathbf{x}^j | \mathbf{w}^F), T^j) + \Delta \quad (3)$$

where  $\Delta$  is an adversarial regularizer introduced to measure the amount of sensitive information available in the learned space represented by  $\mathbf{w}^j$ :

$$\Delta = \log\{1 + |0.9 - P(\text{Male} | \mathbf{x}^j)|\}. \quad (4)$$

The probability  $P$  is the output of a classifier trained to detect the sensitive attribute in the learned space (e.g., Gender in this example). In other words,  $P$  is the probability of observing *Male* features in the learned space after the sensitive information suppression (see [71] for details).

### Scenarios and Protocols

In order to evaluate how and to what extent an algorithm is influenced by biases that are present in the FairCVdb target function, we use the FairCVdb dataset previously introduced in “[FairCVdb: Dataset for Multimodal Bias Research](#)” to train a recruitment system under three different scenarios. The proposed testbed (FairCVtest) consist of FairCVdb, the trained recruitment systems, and the related experimental protocols.

We present three different versions of the recruitment system, with slight differences in the input data and the target function aimed at studying gender/ethnicity biases in multimodal learning. The three scenarios included in FairCVtest were all trained using the candidate competencies, a face

representation, and a short bio, with the following particular configurations:

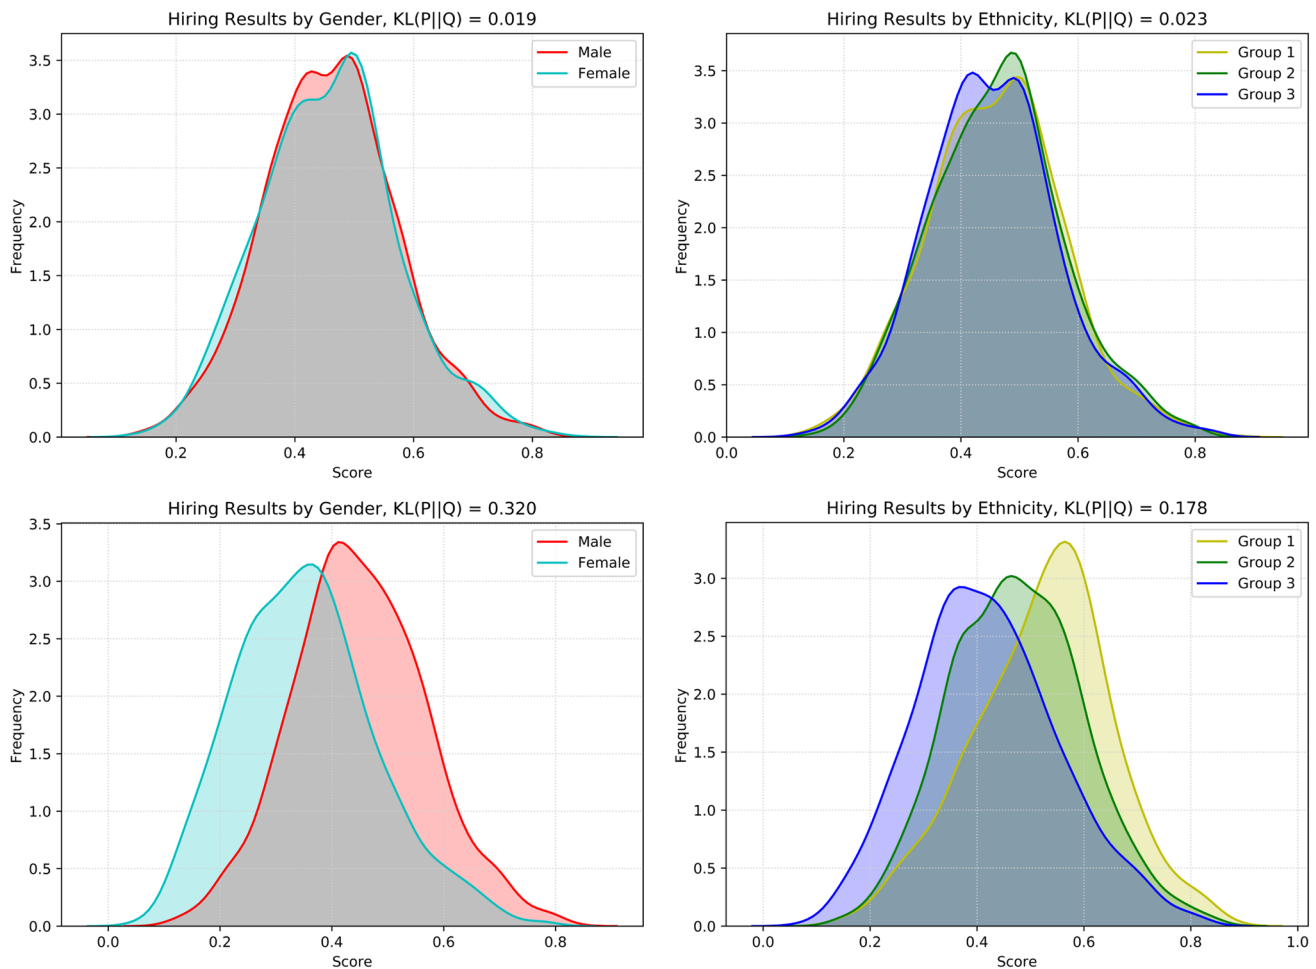
- *Neutral*: Training with Unbiased scores  $T^U$ , the original face representation extracted with ResNet-50 [98], and the biography with explicit gender indicators.
- *Biased*: Training with Biased scores  $T^{(G/E)}$ , the original face representation, and the biography with explicit gender indicators.
- *Agnostic*: Training with Biased scores  $T^{(G/E)}$ , the gender and ethnicity agnostic representation learned with [71], and the “blind” biography.

The experiments performed in next section will try to evaluate the capacity of the recruitment AI in each scenario to detect protected attributes (e.g., gender, ethnicity) without being explicitly trained for this task.

## Experiments and Results

In this section we will train and evaluate different recruitment models, aimed to predict a score from the candidate resumes. Each recruitment tool follows the configuration of one of the scenarios exposed in “[Scenarios and Protocols](#)”, and was trained for 16 epochs using Adam optimizer ( $\alpha = 0.001$ ,  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ ), batch size of 128, and mean absolute error as loss metric.

In Fig. 5 we can observe the distribution of the scores predicted from our validation set, by gender or ethnicity, in both Neutral and Biased scenarios. As a measure of the bias’ impact in the classifier, we compute the Kullback–Leibler divergence  $\text{KL}(P \parallel Q)$  between demographic distributions. In the gender case, we define  $P$  as the male score distribution and  $Q$  as the female’s one, while in the ethnicity setup we make 1-1 comparisons (i.e., G1 vs G2, G1 vs G3 and G2 vs G3) and report the average divergence. In the Neutral Scenario (see top row in Fig. 5), there is no difference between demographic groups, as can be corroborated with the KL divergence tending to zero in both cases (KL = 0.019 in the gender case, KL = 0.023 in the ethnicity one). As expected, using the unbiased scores  $T^U$  as target function and a balanced training set leads us to an unbiased classifier, even in the presence of data containing demographic information (as we will see in “[Privacy in Recruitment Tools: Removing Sensitive Information](#)”). On the other hand, the demographic difference is clearly visible in the Biased scenarios. This difference is most notorious in the gender case (see bottom-left plot in Fig. 5), with the KL divergence rising to 0.320, compared to its low value in the Neutral setup. Attending to the Ethnicity-biased Scenario, the average KL divergence rises to 0.178. However, the difference between groups 1 and 3



**Fig. 5** Hiring score distributions by gender (left) and ethnicity (right). The top row presents hiring score distributions in the Neutral Scenario, while the bottom presents them in the Gender- and Ethnicity-biased Scenarios

is close to that seen between male–female classes, with a KL divergence around 0.317. Note that gender or ethnicity are not inputs of our model, but rather the system is able to detect this sensitive information from some of the input features (i.e., the face embedding, the biography, or the competencies). Therefore, despite not having explicit access to demographic attributes, the classifier is able to detect this information and find its correlation with the biases introduced in the scores, and so it ends up reproducing them.

The third scenario provided by FairCVtest, which we call Agnostic Scenario, aims to prevent the system to inherit data biases. As we introduced in “[Scenarios and Protocols](#)”, the Agnostic Scenario uses a gender blind version of the biographies, as well as a face embedding where sensitive information has been removed using the method of [71]. Figure 6 presents the hiring score distributions in this setup. As we can see, the gender distributions are close to the ones observed in the Neutral Scenario (see top-left plot in Fig. 5),

despite using gender-biased labels during training. In the ethnicity case, we can observe a slight difference between groups, much smoother than the one we saw in the Biased Scenario (see bottom-left plot in Fig. 5), as can be confirmed with the KL divergence (i.e., 0.061, compared to the biased case where this value is around 0.178). However, this gap on the scores between demographic groups still has margin to decrease to a level similar to that of the Neutral Scenario. The difference observed in the behavior of gender and ethnicity agnostic cases can be explained by the fact that we removed almost all gender information from the input (i.e., face embedding and biography), but for the ethnicity we only took measures on the face embedding, not on the competencies. Thus, competencies are acting as a soft proxy for the ethnicity group.

Note that our agnostic approximation does not seek to make the system capable of detecting whether a score is unfair, nor to compensate such bias, but rather blind it to sensitive attributes with the aim of preventing the model to establish a

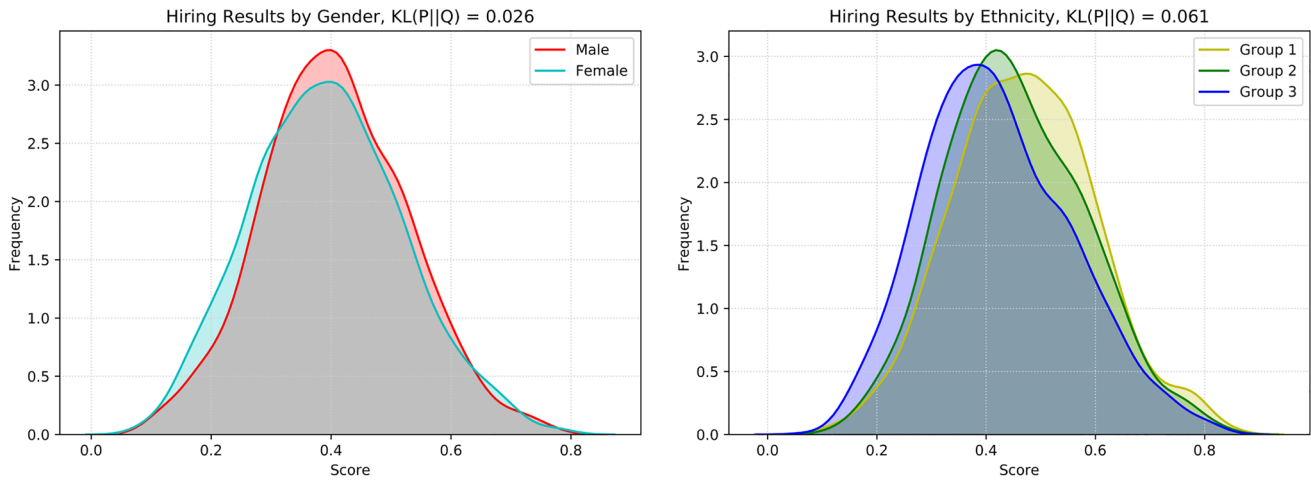


Fig. 6 Hiring score distributions by gender (left) and ethnicity (right), in the Agnostic Scenario

correlation between the demographic groups and score biases. This fact can be corroborated with the training loss, which has a higher value in the Agnostic Scenario (0.035 for gender, 0.044 for ethnicity) than in the Biased Scenario (0.49 for gender, 0.64 for ethnicity). By removing sensitive information from the input, the model is not able to learn what motivates the difference in the scores between individuals with similar competencies, as it is blind to the demographic group, and therefore its output does not approximate correctly the biased target function after training.

### Fairness in Recruitment Tools: Learning Demographic Parity

Now that we have analyzed the effect of data biases in the score distributions, in this section we evaluate their impact in the final decision of a screening process. A screening tool is used to assess a set of individuals according to certain criteria to select a subset of the “best” ones. The outcome of such process could be a list of selected candidates (e.g., applicants selected for an individual interview) or a top-*k* ranking that measures the relative quality of the *k* best individuals from the set. We propose an experiment to simulate a screening process with FairCTest, using the recruitment tools that we trained in the previous section. For each scenario, we predict the scores from a pool including the 4800 resumes of our validation set, and select the top-1000 candidates (i.e., the candidates with the highest scores) among them. By selecting the 1000 candidates with the highest scores, we establish a thresholding rule to classify the candidates in two categories, therefore switching from a regression task to a binary classification task.

We will measure fairness in each scenario using the demographic parity criterion. This criterion requires a classifier’s decision to be statistically independent of a protected attribute (i.e., gender or ethnicity in our experiments). As we’re

working with balanced groups, the criterion implies that all demographic groups should have the same rate of appearance in the top. We can measure demographic parity between two groups through the *p*% score as

$$p\% = \min \left( \frac{P(\hat{y} = 1 | z = 0)}{P(\hat{y} = 1 | z = 1)}, \frac{P(\hat{y} = 1 | z = 1)}{P(\hat{y} = 1 | z = 0)} \right) \tag{5}$$

where  $\hat{y}$  is a trained classifier’s prediction, and *z* is a binary protected attribute. The *p*% score calculates how far off the equality the model’s decisions are. According to the U.S. Equal Employment Opportunity Commission “4/5 rule” [100], the positive rate of a protected group should not be less than 4/5 of that of the group with the higher positive rate. Otherwise, the protected group could be suffering disparate impact. Hence, we will set this rate as an indicator that a model is biased.

Table 3 presents the top-1000 candidates in each scenario, by gender and ethnicity group. In the ethnicity case, we compute three *p*% scores per model by doing 1-1 comparisons with the three ethnic groups. As we can observe, in the Neutral Scenario the classifier shows no demographic bias, with both gender and ethnicity groups having a balanced representation in the ranking. This can be corroborated with the *p*% score, which reach values higher than 90% in all cases. In the Biased scenario the groups Male and Ethnic Group 1 are significantly favored and the difference between groups is now clearly visible. In the gender case, almost 70% of the individuals in the top belong to the Male group, which reduces the *p*% score nearly to 40%. On the other hand, the first ethnic group represents almost half of the top, with the third one exhibiting a 21.6%. For both G2 and G3 the *p*% score points out unfair treatment (i.e., a value under 80%) with respect to G1 (see *p*<sub>1</sub>% and *p*<sub>2</sub>% in Table 3). Finally, in the Agnostic Scenario, the demographic differences were

**Table 3** Distribution of the top 1000 candidates in each Scenario of FairCVtest, by gender and ethnicity group

Scenario	Gender		$p\%$	Ethnicity			$p_1\%$	$p_2\%$	$p_3\%$
	Male (%)	Female (%)		Group 1 (%)	Group 2 (%)	Group 3 (%)			
Neutral	51.90	48.10	92.68	34.20	35.00	30.80	97.71	90.06	88.00
Biased	72.90	27.10	37.17	50.80	30.40	18.80	59.84	37.01	61.84
Agnostic	52.80	47.20	89.39	36.70	32.70	30.60	89.10	83.38	93.58

We include the  $p\%$  score (see Eq. 5) as a measure of the difference between groups. In the ethnicity case,  $p_1\%$  for G1 vs G2,  $p_2\%$  for G1 vs G3 and  $p_3\%$  for G2 vs G3

**Table 4** Accuracy of different classification algorithms, trained with feature embeddings extracted by the recruitment tool in each scenario (SVM = Support Vector Machines, RF = Random Forests, NN = Neural Networks)

Scenario	Gender classification			Ethnicity classification		
	SVM (%)	RF (%)	NN (%)	SVM (%)	RF (%)	NN (%)
Neutral	65.04	62.25	63.92	54.13	51.94	50.29
Biased	89.50	88.46	86.37	76.40	75.88	74.31
Agnostic	54.13	51.94	52.94	48.85	48.13	49.71

significantly reduced with respect to the Biased one, with male and female rates showing even more balance than in the Neutral Scenario. The reduction of the gap among ethnic rates is enough to overcome the limit in the  $p\%$  score, but still leaves room for improvement with a difference of nearly 6% between G1 and G3. This is not surprising, as we already observed in Fig. 6 a slight difference between the score distribution of each ethnicity group.

### Privacy in Recruitment Tools: Removing Sensitive Information

We have observed in the previous sections the impact of demographic biases in both the score distribution and the selection rates in different scenarios. In these experiments, the difference between groups was a consequence of the biases introduced in the target function. However, as can be seen in the Agnostic Scenario, by removing gender and ethnicity information from the input we can prevent the model to reproduce those biases, as it cannot see which factor determines the score penalty for some individuals.

Since the key of our Agnostic Scenario is the removal of sensitive information, in this section we will analyze the demographic information extracted by the hiring tool in each scenario. To this aim, we use multimodal feature embeddings extracted by the recruitment tool to train and evaluate the performance of both gender and ethnicity classifiers. We obtain these embeddings as the output of the first dense layer of our learning architecture (see “Scenarios and Protocols”), in which the information from different data domains has already been fused. For each scenario, we train 3 different classification algorithms,

namely Support Vector Machines (SVM), Random Forests (RF), and Neural Networks (NN).

Table 4 presents the accuracies obtained by each classification algorithm in the three scenarios of FairCVtest. The results show a different behavior between scenarios and demographic traits. As expected, the setup in which most sensitive information can be extracted (gender and ethnicity in this work) is the Biased one for both attributes. The SVM classifier obtains the higher validation accuracies, with almost 90% in the gender case and 76.40% in the ethnicity one. Note that none of these values reach state-of-art performances (i.e., neither the ResNet-50 model nor the hiring tools were not explicitly trained to classify those attributes), but both of them warn of large amounts of sensitive information within the embeddings. On the other hand, both Neutral and Agnostic scenarios show lower accuracies than the Biased configuration. However, we can see a gap in performance between them, with all the classifiers showing higher accuracy in the Neutral Scenario. This fact demonstrates that, despite training with the Unbiased scores  $T^U$  which have no relationship with any demographic group membership, the embeddings extracted in the Neutral Scenario contain some sensitive information. Using the gender blinded bios and the face embeddings in which demographic information has been removed, we reduced the amount of latent sensitive information within the agnostic embeddings. This reduction leads us to almost random-choice accuracies in the gender case (i.e., in a binary task, the random-choice classifier’s accuracy is 50%), but in the ethnicity one the classifiers fall far from this limit (i.e., 33% corresponding to 3 ethnic groups), since there is still some information related to that sensitive attribute in the candidate competencies.

## Conclusions

The development of Human-Centric Artificial Intelligence applications will be critical to ensure the correct deployment of AI technologies in our society. In this paper, we have revised the recent advances in this field, with particular attention to available databases proposed by the research community. We have also presented FairCVtest, a new experimental framework (publicly available<sup>9</sup>) on AI-based automated recruitment to study how multimodal machine learning is affected by biases present in the training data. Using FairCVtest, we have studied the capacity of common deep learning algorithms to expose and exploit sensitive information from commonly used structured and unstructured data.

The contributed experimental framework includes FairCVdb, a large set of 24,000 synthetic profiles with information typically found in job applicants' resumes from different data domains (e.g., face images, text data and structured data). These profiles were scored introducing gender and ethnicity biases, which resulted in gender and ethnicity discrimination in the learned models targeted to generate candidate scores for hiring purposes. In this scenario, the system was able to extract demographic information from the input data, and learn its relation with the biases introduced in the scores. This behavior is not limited to the case studied, where the bias lies in the target function. Feature selection or unbalanced data can also become sources of biases. This last case is common when datasets are collected from historical sources that fail to represent the diversity of our society.

We discussed recent methods to prevent undesired effects of algorithmic biases, as well as the most widely used databases in the bias and fairness research in AI. We then experimented with one of these methods, known as SensitiveNets, to improve fairness in this AI-based recruitment framework. Our agnostic setup removes sensitive information from text data at the input level, and apply SensitiveNets to remove it from the face images during the learning process. After the demographic "blinding" process, the recruitment system did not show discriminatory treatment even in the presence of biases in training data, thus improving equity among different demographic groups.

The most common approach to analyze algorithmic discrimination is through group-based bias [14]. However, recent works are now starting to investigate biased effects in AI with user-specific methods, e.g., [75, 101]. We plan to update FairCVtest with such user-specific biases in addition to the considered group-based bias. Other future work includes extending our testbed to other multimodal setups like smartphone-based interaction with application to authentication [102], behavior understanding [103], and

remote monitoring/assessment [104]. Finally, we also foresee worthy research in the extension of the presented bias-assessment [105] and bias-reduction methods [71] based on recent advances in biometric template protection [106] and distributed privacy preservation [107].

**Acknowledgements** This work has received funding from different projects, including BBforTAI (PID2021-127641OB-I00 MICINN/FEDER), HumanCAIC (TED2021-131787B-I00), TRESPASS-ETN (MSCA-ITN-2019-860813), and PRIMA (MSCA-ITN-2019-860315). The work of A. Peña is supported by a FPU Fellowship (FPU21/00535) by the Spanish MIU. Also, I. Serna is supported by a FPI Fellowship from the UAM.

**Funding** This work has received funding from different projects, including BBforTAI (PID2021-127641OB-I00 MICINN/FEDER), HumanCAIC (TED2021-131787B-I00), TRESPASS-ETN (MSCA-ITN-2019-860813), and PRIMA (MSCA-ITN-2019-860315). The work of A. Peña is currently supported by a FPU Fellowship (FPU21/00535) by the Spanish MIU and was supported by Madrid Government (PRICIT(2020/00334/001)) during the elaboration of this work. Also, I. Serna is supported by a FPI Fellowship from the UAM.

**Data availability** The FairCVdb is available in the following link <https://github.com/BiDALab/FairCVtest>. Please contact the corresponding author for more information.

## Declarations

**Conflict of Interest** On behalf of all the authors, the corresponding author states that there is no conflict of interest.

**Ethical Approval** This article does not contain any studies with human participants performed by any of the authors.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Barocas S, Selbst AD. Big data's disparate impact. *Calif Law Rev.* 2016;104:671–732.
2. Acien A, Morales A, Vera-Rodriguez R, Bartolome I, Fierrez J. Measuring the gender and ethnicity bias in deep models for face recognition. In: *Proceedings of Iberoamerican Congress on pattern recognition (IbPRIA)*, Madrid, Spain; 2018.
3. Drozdowski P, Rathgeb C, Dantcheva A, Damer N, Busch C. Demographic bias in biometrics: a survey on an emerging challenge. *IEEE Trans Technol Soc.* 2020;1:89–103.
4. Nagpal S, Singh M, Singh R, Vatsa M, Ratha NK. Deep learning for face recognition: pride or prejudiced? 2019. [arXiv:1904.01219](https://arxiv.org/abs/1904.01219).

<sup>9</sup> <https://github.com/BiDALab/FairCVtest>.




5. Zhao J, Wang T, Yatskar M, Ordonez V, Chang K. Men also like shopping: reducing gender bias amplification using corpus-level constraints. In: Proceedings of conference on empirical methods in natural language processing; Copenhagen, Denmark: Association for Computational Linguistics; 2017. p. 2979–89.
6. Noble SU. Algorithms of oppression: how search engines reinforce racism. New York: NYU Press; 2018.
7. Sweeney L. Discrimination in online ad delivery. *Queue*. 2013;11:10–29.
8. Ali M, Sapiezynski P, Bogen M, Korolova A, Mislove A, Rieke A. Discrimination through optimization: how Facebook’s ad delivery can lead to skewed outcomes. In: Proceedings of the ACM conference on human–computer interaction; NY, USA: Association for Computing Machinery; 2019.
9. Angwin J, Larson J, Mattu S, Kirchner L. Machine bias. New York: ProPublica; 2016.
10. Evans M, Mathews AW. New York regulator probes United Health algorithm for racial bias. *Wall Street J*. 2019.
11. Knight W. The Apple Card didn’t ‘see’ gender—and that’s the problem. *Wired*; 2019.
12. Buolamwini J, Geburu T. Gender shades: intersectional accuracy disparities in commercial gender classification. In: Proceedings of the ACM conference on fairness, accountability, and transparency; NY, USA: PMLR; 2018.
13. Wang M, Deng W. Mitigating bias in face recognition using skewness-aware reinforcement learning. In: IEEE conference on computer vision and pattern recognition (CVPR); Seattle, USA: IEEE; 2020. p. 9322–31.
14. Serna I, Morales A, Fierrez J, Cebrian M, Obradovich N, Rahwan I. Algorithmic discrimination: formulation and exploration in deep learning-based face biometrics. In: Proceedings of the AAAI workshop on SafeAI; NY, USA: CEUR Workshop Proceedings; 2020.
15. Balakrishnan G, Xiong Y, Xia W, Perona P. Towards causal benchmarking of bias in face analysis algorithms. In: European conference on computer vision (ECCV); Glasgow, UK: Springer-Verlag; 2020. p. 547–63.
16. Bogen M, Rieke A. Help wanted: examination of hiring algorithms, equity, and bias. Technical report; 2018. <https://www.upturn.org>
17. Black JS, van Esch P. AI-enabled recruiting: what is it and how should a manager use it? *Bus Horiz*. 2020;63:215–26.
18. Dastin J. Amazon scraps secret AI recruiting tool that showed bias against women. London: Reuters; 2018.
19. Bertrand M, Mullainathan S. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *Am Econ Rev*. 2004;94:991–1013.
20. Raghavan M, Barocas S, Kleinberg J, Levy K. Mitigating bias in algorithmic hiring: evaluating claims and practices. In: Conference on fairness, accountability, and transparency; NY, USA: Association for Computing Machinery; 2020. p. 469–81.
21. Schumann C, Foster JS, Mattei N, Dickerson JP. We need fairness and explainability in algorithmic hiring. In: Proceedings of the 19th international conference on autonomous agents and multiagent systems; Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems; 2020. p. 1716–20.
22. Sánchez-Monedero J, Dencik L, Edwards L. What does it mean to ‘solve’ the problem of discrimination in hiring? Social, technical and legal perspectives from the UK on automated hiring systems. In: Conference on fairness, accountability, and transparency; NY, USA: Association for Computing Machinery 2020. p. 458–68.
23. Goodman B, Flaxman S. EU regulations on algorithmic decision-making and a “Right to explanation.” *AI Mag*. 2016;38:50–7.
24. Cheng L, Varshney KR, Liu H. Socially responsible ai algorithms: issues, purposes, and challenges. *J Artif Intell Res*. 2021;71:1137–81.
25. Baltrušaitis T, Ahuja C, Morency L. Multimodal machine learning: a survey and taxonomy. *IEEE Trans Pattern Anal Mach Intell*. 2019;41:423–43.
26. Peña A, Serna I, Morales A, Fierrez J. Bias in multimodal AI: testbed for fair automatic recruitment. In: IEEE CVPR workshop on fair, data efficient and trusted computer vision; 2020.
27. Richards C, Bouman WP, Seal L, Barker MJ, Nieder TO, T’Sjoen G. Non-binary or genderqueer genders. *Int Rev Psychiatry*. 2016;28(1):95–102.
28. Keyes O. The misgendering machines: trans/hci implications of automatic gender recognition. In: Proceedings of the ACM on human–computer interaction 2(CSCW); 2018. p. 1–22.
29. Larrazabal AJ, Nieto N, Peterson V, Milone DH, Ferrante E. Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. *Proc Natl Acad Sci*. 2020;117:12592–4.
30. Speicher T, Ali M, Venkatadri G, Ribeiro F.N, Arvanitakis G, Benevenuto F, Gummadi K.P, Loiseau P, Mislove A. Potential for discrimination in online targeted advertising. In: Conference on fairness, accountability and transparency; 2018. p. 5–19.
31. De-Arteaga M, Romanov R, Wallach H, Chayes J, Borgs C, et al. Bias in bios: a case study of semantic representation bias in a high-stakes setting. In: Conference on fairness, accountability, and transparency; 2019. p. 120–8.
32. Bolukbasi T, Chang K, Zou JY, Saligrama V, Kalai AT. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. *Adv Neural Inf Process Syst*. 2016;29:4356–64.
33. Bengio Y, Courville A, Vincent P. Representation learning: a review and new perspectives. *IEEE Trans Pattern Anal Mach Intell*. 2013;35(8):1798–828.
34. Bau D, Zhu J, Strobel H, Lapedriza A, Zhou B, Torralba A. Understanding the role of individual units in a deep neural network. *Proc Natl Acad Sci*. 2020;117(48):1–8.
35. Yosinski J, Clune J, Nguyen A, Fuchs T, Lipson H. Understanding neural networks through deep visualization. In: International conference on machine learning (ICML) deep learning workshop, Lille, France; 2015.
36. Geirhos R, Rubisch P, Michaelis C, Bethge M, Wichmann FA, Brendel W. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In: International conference on learning representations (ICLR), New Orleans, Louisiana, USA; 2019.
37. Bach S, Binder A, Montavon G, Klauschen F, Müller K, Samek W. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS One*. 2015;10(7):1–46.
38. Selvaraju R, Cogswell M, et al. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In: IEEE international conference on computer vision (CVPR), Honolulu, Hawaii, USA. IEEE; 2017. p. 618–26.
39. Ortega A, Fierrez J, Morales A, Wang Z, de la Cruz M, Alonso CL, Ribeiro T. Symbolic AI for XAI: evaluating LFIT inductive programming for explaining biases in machine learning. *Computers*. 2021;10(11):154.
40. Hendricks LA, Akata Z, Rohrbach M, Donahue J, Schiele B, Darrell T. Generating visual explanations. In: European conference on computer vision (ECCV), Amsterdam, The Netherlands. Berlin: Springer; 2016. p. 3–19.
41. Montavon G, Samek W, Müller K. Methods for interpreting and understanding deep neural networks. *Digit Signal Process*. 2018;73:1–15.

42. Erhan D, Bengio Y, Courville A, Vincent P. Visualizing higher-layer features of a deep network, vol. 1341(3). Montreal: University of Montreal; 2009.
43. Simonyan K, Vedaldi A, Zisserman A. Deep inside convolutional networks: visualising image classification models and saliency maps. In: International conference on learning representations (ICLR) workshop, Banff, Canada; 2014.
44. Mahendran A, Vedaldi A. Understanding deep image representations by inverting them. In: IEEE conference on computer vision and pattern recognition (CVPR), Boston, MA, USA. IEEE; 2015. p. 5188–96.
45. Nguyen A, Yosinski J, Clune J. Multifaceted feature visualization: uncovering the different types of features learned by each neuron in deep neural networks. In: International conference on machine learning (ICML) deep learning workshop, New York, NY, USA; 2016.
46. Nguyen A, Dosovitskiy A, Yosinski J, Brox T, Clune J. Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. In: Conference on neural information processing systems (NIPS), Barcelona, Spain; 2016. p. 3395–403.
47. Nguyen A, Clune J, Bengio Y, Dosovitskiy A, Yosinski J. Plug & Play generative networks: conditional iterative generation of images in latent space. In: IEEE conference on computer vision and pattern recognition (CVPR), Honolulu, Hawaii, USA. IEEE; 2017.
48. Karnin ED. A simple procedure for pruning back-propagation trained neural networks. *Trans Neural Netw.* 1990;1(2):239–42.
49. Zurada JM, Malinowski A, Cloete I. Sensitivity analysis for minimization of input data dimension for feedforward neural network. In: International symposium on circuits and systems (ISCAS), vol. 6; 1994. p. 447–50.
50. Zeiler D, Fergus R. Visualizing and understanding convolutional networks. In: European conference on computer vision (ECCV), Zurich, Switzerland. Berlin: Springer; 2014. p. 818–33.
51. Springenberg JT, Dosovitskiy A, Brox T, Riedmiller M. Striving for simplicity: the all convolutional net. In: International conference on learning representations (ICLR), San Diego, CA, USA; 2015.
52. Zhang Q, Cao R, Shi F, Wu YN, Zhu S. Interpreting CNN knowledge via an explanatory graph. In: AAAI conference on artificial intelligence, vol. 32. New Orleans: AAAI Press; 2018.
53. Adebayo J, Gilmer J, Muelly M, Goodfellow I, Hardt M, Kim B. Sanity checks for saliency maps. In: Advances in neural information processing systems (NIPS), vol. 31. Montréal: Curran Associates Inc.; 2018. p. 9525–36.
54. Szegedy C, Zaremba W, Sutskever I, Estrach JB, Erhan D, Goodfellow I, Fergus R. Intriguing properties of neural networks. In: International conference on learning representations (ICLR), Banff, Canada; 2014.
55. Pang WK, Percy L. Understanding black-box predictions via influence functions. In: International conference on machine learning (ICML), vol. 70. Sydney: PMLR; 2017. p. 1885–94.
56. Nguyen A, Yosinski J, Clune J. Deep neural networks are easily fooled: high confidence predictions for unrecognizable images. In: IEEE conference on computer vision and pattern recognition (CVPR). IEEE; 2015. p. 427–36.
57. Su J, Vargas DV, Sakurai K. One pixel attack for fooling deep neural networks. *Trans Evol Comput.* 2019;23(5):828–41.
58. Quadrianto N, Sharmanska V, Thomas O. Discovering fair representations in the data domain. In: IEEE conference on computer vision and pattern recognition (CVPR) (2019). p. 8227–36.
59. Sattigeri P, Hoffman SC, Chenthamarakshan V, Varshney KR. Fairness GAN: generating datasets with fairness properties using a generative adversarial network. *IBM J Res Dev.* 2019;63:1–9.
60. Odena A, Olah C, Shlens J. Conditional image synthesis with auxiliary classifier GANs. In: International conference on machine learning (ICML), Sydney, Australia; 2017. p. 2642–51.
61. Calmon FP, Wei D, Vinzamuri B, Ramamurthy KN, Varshney KR. Optimized pre-processing for discrimination prevention. In: Proceedings of the 31st international conference on neural information processing systems; 2017. p. 3995–4004.
62. Ramaswamy VV, Kim SS, Russakovsky O. Fair attribute classification through latent space de-biasing. In: IEEE conference on computer vision and pattern recognition; 2021. p. 9301–10.
63. Jia S, Lansdall-Welfare T, Cristianini N. Right for the right reason: training agnostic networks. In: Advances in intelligent data analysis XVII; 2018. p. 164–74.
64. Ganin Y, Ustinova E, Ajakan H, Germain P, Larochelle H, Laviolette F, Marchand M, Lempitsky V. Domain-adversarial training of neural networks. *J Mach Learn Res.* 2016;17:1–35.
65. Wang M, Deng W, Hu J, Tao X, Huang Y. Racial faces in the wild: reducing racial bias by information maximization adaptation network. In: IEEE international conference on computer vision (ICCV); 2019. p. 692–702.
66. Romanov A, De-Arteaga M, Wallach H, Chayes J, Borgs C, et al. What’s in a name? Reducing bias in bios without access to protected attributes. In: Proceedings of the 2019 conference of the North American chapter of the Association for Computational Linguistics: human language technologies; 2019. p. 4187–95.
67. Deng J, Guo J, Xue N, Zafeiriou S. ArcFace: additive angular margin loss for deep face recognition. In: IEEE conference on computer vision and pattern recognition (CVPR); 2019. p. 4690–99.
68. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, et al. Generative adversarial nets. *Adv Neural Inf Process Syst.* 2014;27:2672–80.
69. Alvi M, Zisserman A, Nellaker C. Turning a blind eye: explicit removal of biases and variation from deep neural network embeddings. In: European conference on computer vision (ECCV); 2018.
70. Kim B, Kim H, Kim K, Kim S, Kim J. Learning not to learn: training deep neural networks with biased data. In: IEEE conference on computer vision and pattern recognition (CVPR); 2019. p. 9012–20.
71. Morales A, Fierrez J, Vera-Rodriguez R, Tolosana R. SensitiveNets: learning agnostic representations with application to face recognition. *IEEE Trans Pattern Anal Mach Intell.* 2021;43(6):2158–64.
72. Schroff F, Kalenichenko D, Philbin J. FaceNet: a unified embedding for face recognition and clustering. In: IEEE conference on computer vision and pattern recognition (CVPR); 2015. p. 815–23.
73. Berendt B, Preibusch S. Exploring discrimination: a user-centric evaluation of discrimination-aware data mining. In: IEEE international conference on data mining workshops; 2012. p. 344–51.
74. Pedreshi D, Ruggieri S, Turini F. Discrimination-aware data mining. In: Proceedings of the 14th ACM SIGKDD international conference on knowledge discovery and data mining; 2008. p. 560–8.
75. Zhang Y, Bellamy R, Varshney KR. Joint optimization of AI fairness and utility: a human-centered approach. In: Proceedings of the AAAI/ACM conference on AI, ethics, and society; 2020. p. 400–6.
76. Yang K, Stoyanovich J. Measuring fairness in ranked outputs. In: Proceedings of the 29th international conference on scientific and statistical database management; 2017. p. 1–6.
77. Celis LE, Straszak D, Vishnoi NK. Ranking with fairness constraints. In: Proceeding of the international colloquium on automata, languages, and programming; 2018. p. 1–15.
78. Zehlike M, Bonchi F, Castillo C, Hajian S, Megahed M, Baeza-Yates R. FA\*IR: a fair top-k ranking algorithm. In: Proceedings

- of the 2017 ACM on conference on information and knowledge management; 2017. p. 1569–78.
79. Dua D, Graff C. UCI machine learning repository; 2017. <http://archive.ics.uci.edu/ml>.
  80. Moro S, Cortez P, Rita P. A data-driven approach to predict the success of bank telemarketing. *Decis Support Syst*. 2014;62:22–31.
  81. Zhao J, Wang T, Yatskar M, Ordonez V, Chang K. Gender bias in coreference resolution: evaluation and debiasing methods. In: *Conference of the North American chapter of the association for computational linguistics: human language technologies, vol. 2*; 2018.
  82. Liu Z, Luo P, Wang X, Tang X. Deep learning face attributes in the wild. In: *International conference on computer vision (ICCV)*; 2015.
  83. Rothe R, Timofte R, Van Gool L. Dex: deep expectation of apparent age from a single image. In: *IEEE international conference on computer vision workshops (CVPRW)*; 2015. p. 10–5.
  84. Ricanek K Jr, Tesafaye T. Morph: a longitudinal image database of normal adult age-progression. In: *International conference on automatic face and gesture recognition*; 2006. p. 341–5.
  85. Karkkainen K, Joo J. FairFace: face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In: *IEEE winter conference on applications of computer vision*; 2021. p. 1548–58.
  86. Merler M, Ratha N, Feris SR, Smith JR. Diversity in faces. 2019. [arXiv:1901.10436](https://arxiv.org/abs/1901.10436).
  87. Robinson J.P, Livitz G, Henon Y, Qin C, Fu Y, Timoner S. Face recognition: too bias, or not too bias? In: *IEEE conference on computer vision and pattern recognition workshops (CVPRW)*; 2020.
  88. Hupont I, Fernández C. DemogPairs: quantifying the impact of demographic imbalance in deep face recognition. In: *IEEE international conference on automatic face and gesture recognition*; 2019.
  89. Torralba A, Efros AA. Unbiased look at dataset bias. In: *IEEE conference on computer vision and pattern recognition (CVPR)*; 2011.
  90. Serna I, Morales A, Fierrez J, Cebrian M, Obradovich N, Rahwan I. SensitiveLoss: improving accuracy and fairness of face representations with discrimination-aware deep learning. *Artif Intell*. 2022;305:103682
  91. Garg N, Schiebinger L, Jurafsky D, Zou J. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proc Natl Acad Sci*. 2018;115:3635–44.
  92. Guo Y, Zhang L, Hu Y, He X, Gao J. MS-Celeb-1M: a dataset and benchmark for large-scale face recognition. In: *European conference on computer vision (ECCV)*; 2016.
  93. Bendick M Jr, Jackson CW, Romero JH. Employment discrimination against older workers: an experimental study of hiring practices. *J Aging Soc Policy*. 1997;8:25–46.
  94. Cowgill B. Bias and productivity in humans and algorithms: theory and evidence from resume screening. *Columbia Business School, Columbia University*. 2018;29.
  95. Fierrez J, Morales A, Vera-Rodriguez R, Camacho D. Multiple classifiers in biometrics. Part 1: fundamentals and review. *Inf Fusion*. 2018;44:57–64.
  96. Gonzalez-Sosa E, Fierrez J, Vera-Rodriguez R, Alonso-Fernandez F. Facial soft biometrics for recognition in the wild: recent works, annotation and COTS evaluation. *IEEE Trans Inf Forensics Secur*. 2018;13:2001–14.
  97. Ranjan R, Sankaranarayanan S, Bansal A, Bodla N, Chen J, Patel VM, Castillo CD, Chellappa R. Deep learning for understanding faces: machines may be just as good, or better, than humans. *IEEE Signal Process Mag*. 2018;35:66–83.
  98. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: *IEEE conference on computer vision and pattern recognition (CVPR)*; 2016. p. 770–8.
  99. Mikolov T, Grave E, Bojanowski P, Puhusch C, Joulin A. Advances in pre-training distributed word representations. In: *Proceedings of the international conference on language resources and evaluation (LREC 2018)*; 2018.
  100. Biddle D. Adverse impact and test validation: a practitioner's guide to valid and defensible employment testing. London: Routledge; 2017.
  101. Bakker M, Valdes HR, Tu DP, Gummadi KP, Varshney KR, et al. Fair enough: improving fairness in budget-constrained decision making using confidence thresholds. In: *AAAI workshop on artificial intelligence safety, New York, NY, USA*; 2020. p. 41–53.
  102. Acien A, Morales A, Vera-Rodriguez R, Fierrez J, Delgado O. Smartphone sensors for modeling human-computer interaction: general outlook and research datasets for user authentication. In: *IEEE conference on computers, software, and applications (COMPSAC)*; 2020.
  103. Acien A, Morales A, Fierrez J, Vera-Rodriguez R, Bartolome I. BeCAPTCHA: detecting human behavior in smartphone interaction using multiple inbuilt sensors. In: *AAAI workshop on artificial intelligence for cyber security (AICS)*; 2020.
  104. Hernandez-Ortega J, Daza R, Morales A, Fierrez J, Ortega-Garcia J. edBB: biometrics and behavior for assessing remote education. In: *AAAI workshop on artificial intelligence for education (AI4EDU)*; 2020.
  105. Serna I, DeAlcala D, Morales A, Fierrez J, Ortega-Garcia J. IFBiD: inference-free bias detection. In: *AAAI workshop on artificial intelligence safety (SafeAI)*. CEUR, vol. 3087; 2022.
  106. Gomez-Barrero M, Maiorana E, Galbally J, Campisi P, Fierrez J. Multi-biometric template protection based on homomorphic encryption. *Pattern Recognit*. 2017;67:149–63.
  107. Hassanpour A, Moradikia M, Yang B, Abdelhadi A, Busch C, Fierrez J. Differential privacy preservation in robust continual learning. *IEEE Access*. 2022;10:24273–2428.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

## Authors and Affiliations

Alejandro Peña<sup>1</sup>  · Ignacio Serna<sup>1</sup> · Aythami Morales<sup>1</sup> · Julian Fierrez<sup>1</sup> · Alfonso Ortega<sup>1</sup> · Ainhoa Herrarte<sup>1</sup> · Manuel Alcantara<sup>1</sup> · Javier Ortega-Garcia<sup>1</sup>

Ignacio Serna  
ignacio.serna@uam.es

Aythami Morales  
aythami.morales@uam.es

Julian Fierrez  
julian.fierrez@uam.es

Alfonso Ortega  
alfonso.ortega@uam.es

Ainhoa Herrarte  
ainhoa.herrarte@uam.es

Manuel Alcantara  
manuel.alcantara@uam.es

Javier Ortega-Garcia  
javier.ortega@uam.es

<sup>1</sup> Universidad Autónoma de Madrid, 28049 Madrid, Spain