



Don't Judge Me by My Face: An Indirect Adversarial Approach to Remove Sensitive Information From Multimodal Neural Representation in Asynchronous Job Video Interviews

Leo Hemamou, Arthur Guillon, Jean-Claude Martin, Chloé Clavel

► To cite this version:

Leo Hemamou, Arthur Guillon, Jean-Claude Martin, Chloé Clavel. Don't Judge Me by My Face: An Indirect Adversarial Approach to Remove Sensitive Information From Multimodal Neural Representation in Asynchronous Job Video Interviews. 2021 9th International Conference on Affective Computing and Intelligent Interaction (ACII), Sep 2021, Nara, Japan. pp.1-8, 10.1109/ACII52823.2021.9597443 . hal-04276017

HAL Id: hal-04276017

<https://hal.science/hal-04276017>

Submitted on 11 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Don't Judge Me by My Face: An Indirect Adversarial Approach to Remove Sensitive Information From Multimodal Neural Representation in Asynchronous Job Video Interviews

Léo Hemamou^{*†‡}, Arthur Guillon^{*}, Jean-Claude Martin[†] and Chloé Clavel[‡]

^{*}EASYRECRUE, Paris, France

Email: {l.hemamou,art.guillon}@gmail.com

[†]LIMSI-LISN, CNRS, Paris-Sud University, Paris-Saclay University / F-91405 Orsay, France

Email: Jean-Claude.Martin@limsi.fr

[‡]Télécom-Paris, IP-Paris, / F-75013 Paris, France

Email: chloe.clavel@telecom-paristech.fr

Abstract—Use of machine learning for automatic analysis of job interview videos has recently seen increased interest. Despite claims of fair output regarding sensitive information such as gender or ethnicity of the candidates, the current approaches rarely provide proof of unbiased decision-making, or that sensitive information is not used. Recently, adversarial methods have been proved to effectively remove sensitive information from the latent representation of neural networks. However, these methods rely on the use of explicitly labeled protected variables (e.g. gender), which cannot be collected in the context of recruiting in some countries (e.g. France). In this article, we propose a new adversarial approach to remove sensitive information from the latent representation of neural networks without the need to collect any sensitive variable. Using only a few frames of the interview, we train our model to not be able to find the face of the candidate related to the job interview in the inner layers of the model. This, in turn, allows us to remove relevant private information from these layers. Comparing our approach to a standard baseline on a public dataset with gender and ethnicity annotations, we show that it effectively removes sensitive information from the main network. Moreover, to the best of our knowledge, this is the first application of adversarial techniques for obtaining a multimodal fair representation in the context of video job interviews. In summary, our contributions aim at improving fairness of the upcoming automatic systems processing videos of job interviews for equality in job selection.

Index Terms—component, formatting, style, styling, insert

I. INTRODUCTION

Job interviews are omnipresent in recruitment procedures. Although they are more generally conducted face-to-face, new forms of job interviews have emerged, such as the asynchronous video interview (AVI). For this form of interview, candidates connect to a web platform and answer a series of questions by recording themselves in a video monologue. Later on, recruiters have the opportunity to watch these videos, evaluate the candidates, and decide whether or not to invite them for a face-to-face interview. The growing interest in

AVIs has naturally led to the development of algorithms and proprietary tools to automatically assess and rank candidates through the use of machine learning [1], [2]. The adoption of this practice has met scepticism, both from lawmakers and the public, as sensitive information (e.g. gender or race) could be unintentionally used by these algorithms leading to unfair selection results. Vendors of algorithmic pre-employment assessment technologies generally claim unbiased predictions, but these claims are rarely backed with published studies or audits [3]. Moreover, those companies only guarantee fairness regarding the *impact disparity* (i.e. if outcomes differ across subgroups) but not regarding the *treatment disparity* (i.e. if subgroups are treated differently). It therefore seems necessary to improve these decision support models for a more accurate and unbiased process.

In this article, we aim to ensure equal treatment by preventing the automatic analysis system from using protected variables of the candidates. Despite these variables not being present in the input of the algorithm, we show that they leak into the hidden representation of the neural network. Sensitive information of the candidates can be retrieved from the inner layers of the network. Thus there is no way of guaranteeing that it is not used by the algorithm to predict the candidates' interview performance or 'hireability' as it is commonly referred to in the literature [1], [2], [4], [5]. Recently, adversarial learning have been successfully proposed to obtain fair representations. To the initial prediction task is added a secondary adversarial task which consists of predicting the protected variables (e.g. gender and ethnicity of the candidates). The overall objective then becomes maximization of the hireability prediction task while minimizing the ability of the adversary branch to predict the protected variables. However, these techniques rely explicitly on these variables to train the adversary branch. This dependency raises two issues

in the context of recruitment. Firstly, the collection of these variables is prohibited by the law in some countries. Secondly, recruitment suffers from multiple potential biases, like physical attractiveness, obesity or age, and explicit labeling of each protected variable by recruiters could become impossible in practice.

In Section III, We propose, two indirect adversarial approaches, which do not explicitly rely on the protected variables of the candidate but only on images of their face, recorded during the video interview. These images are correlated with the protected information, but are available, as indeed automatic analysis of AVIs requires explicit consent from the candidates for the processing of their videos. By relying on a state of the art neural network designed to predict hireability, we propose two new architectures based on this approach.

In Section IV, we compare the performance of our approach to standard adversarial techniques that explicitly use the protected variables. We show that it succeeds in removing this information from the latent layer of the network without degrading its performance on the main task or using sensitive candidate information. We believe that our contribution represents an important first step in the field of fairness for AVIs. While our data allows us to assess the impact of our approach only on gender and ethnicity features, we are of the opinion that the methods described are likely to deliver similar results for other types of biases.

II. SOCIAL CONTEXT AND RELATED WORKS

A. Discrimination in Job Interviews

Discrimination in job interviews is a recurring problem at the intersection of industrial and organizational psychology, social science and ethics. Both the law and the theory of selection process indicate that decisions should only be made in relation to useful dimensions regarding the necessary skills of the job. However, literature has shown that factors such as gender, ethnicity, physical appearance or obesity do have an influence, even though taking a decision on these factors is expressly forbidden by law in both Europe and the United States [6]. Moreover, in the context of AVIs, having access to the image of the candidates early in the selection process has a potential impact on increasing the influence of aesthetics [7], gender or ethnicity of the candidates [8]. Approaches to mitigate discrimination in job interviews may differ from country to country [9]. For example, in the US, the 4th/5th rule proposed by the EEOC (Equal Employment Opportunity Commission) states that the ratio of the least favored group compared to the most favored group must not be less than 0.8. In comparison, this rule cannot be used in some European countries where it is forbidden to collect sensitive attributes such as race [10], preventing approaches based on statistical analysis.

These policies have an impact on the methods for measuring fairness used in the machine learning field [11]. For instance, *classification parity* ensures that performance measures are

equal between each of the protected groups, whereas *anti-classification* ensures that no protected variable (race, gender, facial attractiveness, etc.) is used to make a decision.

B. Automatic evaluation of AVIs

Originally investigated in face-to-face settings [12], automatic analysis of video interviews has benefited from the emerging trend of AVIs. Indeed, the possibility of obtaining a large dataset as well as the structure of these interviews has allowed the use of machine learning [1], [2] and more recently deep learning [5], [13]–[15] to assess hireability. However, little work has been done to investigate the question of fairness in these systems [3], which could produce biased outputs or unfair treatment towards minority groups. In fact, systematic bias could be the result of bias in the dataset, inadequate representation in the construction of the model, or under-representation of minority populations. Thus, some researchers have studied biases that might exist within the dataset [13], [14] or in evaluating the fairness of the model’s output [15]. These studies are limited to ensuring that there is no bias in the system output or in the dataset, and they do not propose any method in the case of a biased pipeline. Furthermore, they do not ensure that sensitive information does not leak into the hiring decision of the system, which would contradict laws about equal treatment. In this paper, we adopt the framework of previous studies on the topic, including descriptor extraction through widely-used libraries (OpenFace [16] and OpenSMILE [17]) and obtaining a latent representation through a GRU-based neural architecture [18]. We show that this methodology does indeed leak sensitive information when inferring the hiring decision.

C. Fairness via adversarial methods

Two major problems tackled by fairness in machine learning are group and individual fairness. *Group fairness* designates the goal of producing equal or comparable results for populations belonging to different groups, according to some protected variable. Some approaches have been specifically designed to tackle this problem. For example, by taking into account the protected variables during inference [19], researchers obtain a perfect fair classifier in terms of group fairness. However, the application of such methods can be limited, as positive discrimination is forbidden in several countries (e.g. France, UK or Germany).

Individual fairness, on the other hand, consists of giving a similar treatment to individuals that are similar regarding the task at hand. To this end, a recent family of approaches in machine learning focus on removing sensitive information from neural representation during training time through adversarial methods [19]–[21]. The idea is to obtain a representation without any information which could potentially lead to the recovering of the protected variables. Interestingly, these methodologies are also linked to privacy methods where network designers try to protect their system from attackers trying to retrieve personal information from latent representations. Thus, adversarial learning has been used in the context

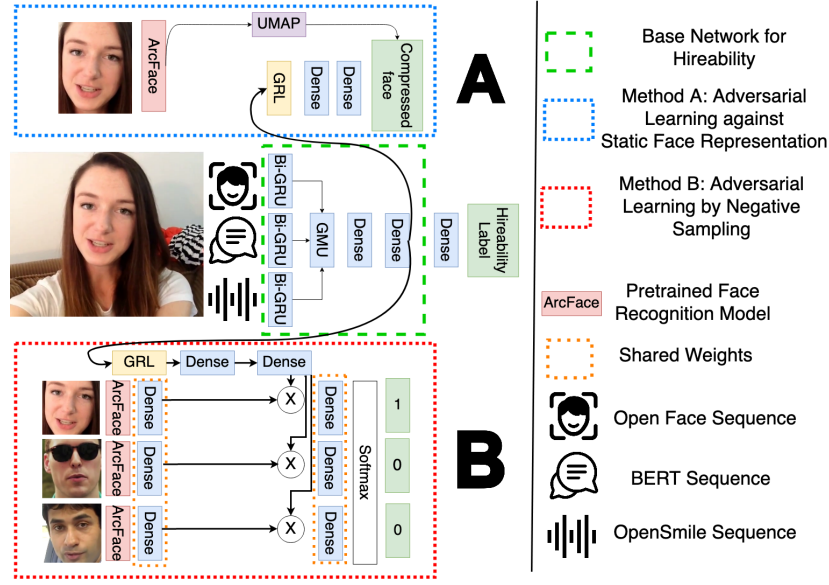


Fig. 1. Architecture of the base network trained for hireability (green, middle of the figure) and the two proposed adversarial methods: static face representation (method A, blue, top) and negative sampling (B, red, bottom).

of speech processing [22], deep visual recognition [21], or multimodal (prosodic and verbal content) emotion recognition [23]. However these approaches always rely on the explicit usage of the protected variables. In the setting of this paper, this would rely on asking the users for this information, a policy which would be inappropriate during a selection process or even illegal in several countries (collecting information about ethnicity is for example illegal in France [24]).

In this article, we evaluate the utility of adversarial methods regarding protection of gender and ethnicity in the context of hireability prediction from multimodal monologue videos without the need of collecting additional sensitive information.

III. PROPOSED CONTRIBUTION: AN INDIRECT ADVERSARIAL APPROACH

We present a new adversarial approach, available in two forms, to remove personal information of AVI candidates from the latent layer of a deep network trained for hireability. To this end, instead of using the candidates' protected features directly, we use a representation of their face, which is extracted from the AVI. In this article, we integrate our approach with an architecture inspired by the state of the art of automatic AVI analysis. We first present the data used, the features and the base architecture in Section III-A. The two new adversarial architectures are presented in Section III-B.

A. Main task: base network for predicting hireability

The data used in this article are videos from the dataset “ChaLearn First Impressions” [13]. To our knowledge, this dataset is the only one which reproduces conditions of AVIs (monologue video), is annotated for hireability, provides sensitive information (gender and ethnicity) of the “candidates” and is publicly available. We also sought industrial

partners for collaboration but found no company willing to explicitly collect sensitive information from real candidates.

Dataset description. The ChaLearn dataset consists of $n = 10000$ video clips extracted from over 3000 YouTube videos of people facing the camera and speaking in English. The number of clips appearing in the data for each video varies from 1 to 5 and their length is 15 s. Each video was annotated for a binary hireability variable by Amazon Mechanical Turk workers, who were asked whether or not they would invite the persons to a job interview. Note that because annotators are not trained recruiters, we expect first impressions to have a greater impact on hireability inference, making it more difficult to remove sensitive information from the representation. Although the ChaLearn dataset comes with a proposed 3-way split between training, validation and test sets, we found that 84 % of clips from the test set have at least one clip in the train set extracted from the same Youtube video. This overlap is potentially problematic, since it could allow classifiers to obtain good results on the hireability prediction task by learning speakers' specific features, e.g. facial or audio, a strategy that makes no sense from the point of view of recruitment. For this reason, we use our own 3-way split with no overlap.

Sensitive information. Each video was annotated with gender (male or female) and ethnicity (asian, caucasian or african-american). In this article, we consider these annotations (gender and ethnicity) as the protected variables which should not be used during the hireability inference. As such, they are only used for the evaluation of the proposed algorithms. Table I displays the average hireability label of each protected class in the complete dataset as well as in all subsets, along with the distribution of the members of these classes.

Features for hireability prediction. Following the method-

Protected variable	Complete dataset	Training set (6991 clips)	Validation set (1448 clips)	Test set (1559 clips)
Gender - Female	0.560	0.560 (3053)	0.575 (702)	0.541 (783)
Gender - Male	0.495	0.482 (3938)	0.524 (746)	0.519 (776)
Disparate Impact	0.883	0.860	0.911	0.959
Ethnicity - Asian	0.570	0.584 (236)	0.718 (32)	0.444 (63)
Ethnicity - Caucasian	0.541	0.537 (5998)	0.545 (1273)	0.554 (1327)
Ethnicity - African-Americans	0.434	0.424 (757)	0.559 (143)	0.374 (171)
Disparate Impact	0.761	0.726	0.759	0.675

TABLE I

SIZE AND MEAN VALUE OF THE HIREABILITY LABEL FOR THE CHALEARN FIRST IMPRESSIONS DATASET, ACCORDING TO OUR TRAINING/VALIDATION/TEST SETS AND THE CLASSES FORMED BY THE TWO PROTECTED VARIABLES.

ology from automatic analysis of interview literature [1], [5], [14], we use three different modalities: vocal cues of the video, facial expressions of the candidate and verbal content. Each of these modalities is extracted from the video as a time series. We extract audio frame-level features from each video using the OpenSmile ComParE [25] feature set, which is standard in the affective computing community. This results in a 130-d time series to represent audio. Frame-level facial expression features are extracted using the OpenFace library [16]. These include: position and rotation of the head, presence and intensity of facial Action Units and direction of the gaze resulting in a vector of dimension 52. Finally, as manual transcripts of each video clip are available, we compute a BERT [26] encoding of each transcript using the second last layer of the base uncased BERT model provided by HuggingFace. We build a representation of the transcript by sequencing the embeddings word by word, obtaining a 768-d time series.

Base network for hireability. Our approach was not designed with a specific architecture in mind and could be used on any deep learning algorithm trained for hireability. In conducting experiments for this article, we use an architecture adapted from that published previously under the name HireNet [5]. Moreover, in addition to the monomodal networks from [5], we also propose to study the impact of our approach on a simple multimodal representation, by fusing the three modalities through the use of a Gated Multimodal Unit (GMU, [27]). The GMU works by projecting all modalities on the same space and learning the contribution of each through a gated mechanism.

Figure 1 presents the base network as well as the two methods we propose. The base multimodal network is shown in the middle of the figure, which depicts the three modalities and the GMU, together with several dense layers and the binary hireability label. On the figure, the green box depicts the inner layers of the network. In the following, we denote by H the latent representation of the top layer of this subnetwork, and θ_H the corresponding learned parameters. Likewise, we denote by θ_D the parameters of the dense layer responsible for the hireability decision. Monomodal networks are not shown, but are obtained by removing the GMU layer and stacking the dense layers on top of the corresponding bidirectional GRU unit.

Static representations of the candidates’ faces. Our

approach uses static representation of the candidates’ face, captured during the AVI, to remove private information from the hidden layer of the network. We use OpenFace confidence levels to extract 5 frames where facial action units are minimal (to obtain the most neutral face). Then, a face representation is obtained using a pre-trained neural network for face recognition tasks, namely ArcFace [28]. We obtain a 512-dimensional vector per candidates by averaging face representations belonging to each candidate. In Figure 1, usage of these static faces representations is shown by boxes labeled “ArcFace”.

B. Indirect adversarial approach

In the context of recruitment, the goal is to predict the hireability Y of the candidate from their video interview X without using protected features Z . In the adversarial setting, this amounts to learning a latent representation H from which Z cannot be predicted. However, unlike in a classical adversarial framework, Z is not accessible. Instead, we have access to the candidate’s face W , which is correlated to Z but contains a lot of additional information. For this reason, it is necessary to project W on a simpler space that respects the protected classes formed by the variable Z (candidates of the same gender, ethnicity, glasses, hair color, tatoos, etc). In this paper, we propose two methods to compress W and learn an independent representation H .

Our two proposed methods employ adversarial training. More precisely, we aim to maximize performance for the hireability objective while minimizing an adversarial objective. The complete model is trained by optimizing the following min-max objective:

$$\min_{\theta_H, \theta_D} \max_{\theta_A} \mathcal{L}_T(\theta_H, \theta_D) - \lambda \mathcal{L}_A(\theta_H, \theta_A) \quad (1)$$

where θ_H and θ_D are the parameters of the base networks defined in the previous section, θ_A is the parameter of the adversarial branch, $\lambda \geq 0$ is a trade-off hyperparameter between the hireability objective and the adversarial objective, \mathcal{L}_T is the loss function for the hireability prediction, i.e. binary cross-entropy, and \mathcal{L}_A is the loss function for the adversarial branch, detailed in the next sections.

Our adversarial approach is implemented through a Gradient Reversal Layer (GRL, [29]), which is a special layer with no weight vector. It behaves as the identity function during the forward pass, but negates the gradient during the backward

pass, resulting in the network *unlearning* the corresponding information. Figure 1 present both methods (top and bottom of the figure).

First method: static representation of the faces. The first method we propose, shown at the top of Figure 1, uses the previously computed ArcFace representation of the candidates' faces as a target for the adversarial branch. However, instead of using the raw ArcFace representation (a 512-d vector), we extract a compressed representation W' using the UMAP algorithm [30]. UMAP is a dimensionality reduction method that projects the data while respecting local distances for similar points and global distances for dissimilar points. By preserving intra and inter-cluster distances, our goal is to obtain a low-level representation that respects the majority of inter-individual differences within the data, which we hope will coincide with the gender and ethnicity of the candidates. In this case, the adversarial objective function is the classical mean squared loss function :

$$\mathcal{L}_A = \frac{1}{N} \sum_{i=1}^N (W'_i - f(H_i))^2 \quad (2)$$

where f denotes the dense layers on top of the GRL.

Second method: negative sampling. We propose a second method, independent of the first, shown at the bottom of the Figure 1. Rather than relying on a compressed representation of the candidates' faces, we sample $k - 1$ negative faces representation from other candidates. Like the previous method, a GRL is grafted just after the processing of the interview representation H . Then the task of the adversarial branches is to identify which face is associated with the interview representation H . More precisely we compute for each pair $(H, W_l)_{l=1..k}$ the value of interest p_l :

$$p_l = \frac{\exp(g(H, W_l))}{\sum_{j=1}^k \exp(g(H, W_j))} \quad (3)$$

where g is a learnt function behaving as a similarity function between H and W_l . Finally, we obtain the following adversarial loss function :

$$\mathcal{L}_A = \frac{1}{N} \sum_{i=1}^N \sum_{l=1}^k \delta_{il} \log(p_l) = \frac{1}{N} \sum_{i=1}^N \log(p_i) \quad (4)$$

where $\delta_{il} = 1$ if face l matches the i th interview, 0 otherwise.

In practice, as shown at the bottom of the Figure 1, g consists of three sub-components. The first one is the encoding of the faces (left of the frame on the figure). The second one is, like the previous method, a branch based on a GRL in charge of encoding the representation of the interviews. The third one consists of a Hadamard product between the face and the interview representation followed by a dense layer. We chose k equal to 5 resulting in batches of one positive and 4 negative pairs.

IV. EXPERIMENTS

In order to assess our approach, we compare the algorithms from Section III with a standard adversarial baseline which explicitly uses the protected variables of gender and ethnicity.

A. Evaluation metrics and baseline classifiers

Hireability metrics. Evaluation metrics used for the hireability task are the accuracy (ACC) and the Area under the ROC curve (AUC). These evaluation metrics are well suited for binary classification and have been widely used in automatic analysis of AVIs [2], [13].

Fairness metrics. In order to assess the presence of sensitive information in the latent representation of the base network, we use a standard privacy-inspired metric: first, we train two *diagnostic classifiers*, Logistic Regression (LR) and XGBoost [31] (XGB) to recover the protected variable (gender or ethnicity) from the latent representation of the network, using the same training, validation and test sets. We use the AUC of these classifiers as a metric. In the case of a multi-class problem, we report the macro one-vs-rest AUC. This protocol is widely used in the privacy and fairness literature [20], [23]. Good performance of the diagnostic classifiers means that sensitive information is still contained in the latent representation.

The second metric we use is the Disparate Impact (DI):

$$DI = \frac{Pr(Y = 1|Z = \text{unprivileged})}{Pr(Y = 1|Z = \text{privileged})} \quad (5)$$

Although it is not explicitly optimized by our proposed approach, it enables us to monitor any potential change in the DI as this metric is used in human resources and in the fair machine learning literature [9]. A DI closer to 1 induces a fairer selection rate across groups.

Standard supervised adversarial learning. We compare our approach to a standard adversarial method which explicitly uses either gender or ethnicity variables. These baselines have the same architecture as the first method from Section III-B, except that the adversarial head is replaced by a sigmoid (gender) or a softmax (ethnicity) to predict the class, and are called "Supervised" in the following.

Size of the representations. As pointed out in section Section III-B, the dimensions of the compressed representation W could be critical. In order to investigate this aspect, we run our experiments by varying this dimension. Specifically, we evaluate both approaches by using either dimensions 2 (highly compressed representation) or 16 (low dimensional representation).

B. Experiments results

Table II presents the results of the experiments, averaged over 3 runs. Proposed methods appear under the name "Static faces" and "Negative sampling", followed by the size of the representation. The results are detailed for each modality and each hireability classifier, allowing comparison of the base "unprotected" network to the adversarially trained networks. The Hireability column gives the performance for the main task. Columns AUC Gender and AUC Ethnicity present the

	Hireability		Gender	Ethnicity	DI Gender	DI Ethnicity
	AUC	ACC	AUC	AUC		
Language modality						
Unprotected	0.613	0.584	0.584	0.500	0.984	0.867
Supervised gender	0.611 $\pm 9.8e-4$	0.582 $\pm 6.4e-4$	0.567 $\pm 2.5e-2$	-	0.957 $\pm 1.1e-2$	-
Supervised ethnicity	0.609 $\pm 1.3e-5$	0.576 $\pm 2.9e-3$	-	0.522 $\pm 1.9e-2$	-	0.832 $\pm 2.7e-2$
Static faces 2	0.611 $\pm 1.0e-3$	0.580 $\pm 9.7e-4$	0.578 $\pm 5.0e-3$	0.503 $\pm 6.7e-3$	0.986 $\pm 1.2e-2$	0.818 $\pm 2.9e-2$
Static faces 16	0.612 $\pm 1.1e-3$	0.580 $\pm 2.3e-3$	0.573 $\pm 4.0e-3$	0.506 $\pm 1.3e-2$	0.963 $\pm 1.0e-2$	0.834 $\pm 1.6e-2$
Negative sampling 2	0.604 $\pm 5.8e-3$	0.578 $\pm 1.3e-3$	0.527 $\pm 2.2e-2$	0.527 $\pm 4.9e-3$	0.974 $\pm 1.0e-2$	0.828 $\pm 1.1e-1$
Negative sampling 16	0.609 $\pm 1.7e-3$	0.578 $\pm 5.8e-3$	0.565 $\pm 2.3e-2$	0.500 ± 0.0	0.973 $\pm 2.0e-2$	0.845 $\pm 4.1e-2$
Audio modality						
Unprotected	0.695	0.642	0.850	0.506	0.837	0.913
Supervised gender	0.692 $\pm 1.3e-4$	0.638 $\pm 6.7e-3$	0.584 $\pm 6.7e-3$	-	0.954 $\pm 8.6e-3$	-
Supervised ethnicity	0.695 $\pm 4.4e-4$	0.640 $\pm 2.3e-3$	-	0.511 $\pm 7.1e-3$	-	0.892 $\pm 1.1e-2$
Static faces 2	0.674 $\pm 5.8e-3$	0.597 $\pm 2.5e-2$	0.593 $\pm 1.4e-2$	0.501 $\pm 2.5e-2$	0.880 $\pm 3.1e-2$	0.946 $\pm 2.0e-2$
Static faces 16	0.695 $\pm 5.6e-4$	0.642 $\pm 6.4e-4$	0.811 $\pm 2.6e-2$	0.505 $\pm 1.2e-2$	0.847 $\pm 3.7e-3$	0.903 $\pm 1.9e-2$
Negative sampling 2	0.678 $\pm 4.2e-3$	0.645 $\pm 9.6e-3$	0.593 $\pm 4.0e-2$	0.514 $\pm 8.4e-3$	0.947 $\pm 3.7e-2$	0.856 $\pm 3.2e-2$
Negative sampling 16	0.674 $\pm 1.6e-2$	0.633 $\pm 9.8e-3$	0.642 $\pm 6.1e-2$	0.519 $\pm 2.2e-2$	0.956 $\pm 5.2e-2$	0.881 $\pm 4.3e-2$
Video modality						
Unprotected	0.700	0.647	0.745	0.661	0.591	0.708
Supervised gender	0.693 $\pm 3.5e-3$	0.647 $\pm 3.2e-3$	0.650 $\pm 1.1e-2$	-	0.769 $\pm 2.0e-2$	-
Supervised ethnicity	0.698 $\pm 1.2e-3$	0.654 $\pm 3.3e-3$	-	0.520 $\pm 1.7e-2$	-	0.744 $\pm 3.3e-3$
Static faces 2	0.700 $\pm 5.9e-4$	0.650 $\pm 4.0e-3$	0.665 $\pm 5.3e-3$	0.646 $\pm 9.1e-3$	0.704 $\pm 1.8e-2$	0.864 $\pm 6.7e-2$
Static faces 16	0.699 $\pm 8.3e-4$	0.651 $\pm 1.9e-3$	0.708 $\pm 1.6e-2$	0.643 $\pm 1.2e-3$	0.636 $\pm 9.2e-3$	0.775 $\pm 7.2e-3$
Negative sampling 2	0.658 $\pm 6.4e-3$	0.613 $\pm 4.3e-3$	0.659 $\pm 1.9e-2$	0.585 $\pm 1.3e-2$	0.971 $\pm 4.5e-2$	0.745 $\pm 9.0e-2$
Negative sampling 16	0.693 $\pm 2.2e-3$	0.643 $\pm 1.6e-3$	0.656 $\pm 1.3e-2$	0.630 $\pm 2.2e-2$	0.692 $\pm 2.6e-2$	0.707 $\pm 1.5e-2$
Multimodality						
Unprotected	0.741	0.682	0.762	0.582	0.695	0.656
Supervised gender	0.730 $\pm 5.2e-3$	0.675 $\pm 1.1e-2$	0.647 $\pm 5.6e-2$	-	0.851 $\pm 9.5e-2$	-
Supervised ethnicity	0.731 $\pm 2.7e-2$	0.669 $\pm 5.5e-3$	-	0.538 $\pm 3.2e-2$	-	0.636 $\pm 3.2e-2$
Static faces 2	0.730 $\pm 1.9e-3$	0.667 $\pm 8.4e-3$	0.644 $\pm 6.2e-2$	0.510 $\pm 3.0e-2$	0.847 $\pm 2.9e-2$	0.640 $\pm 7.2e-2$
Static faces 16	0.741 $\pm 1.1e-3$	0.683 $\pm 2.0e-3$	0.752 $\pm 2.2e-2$	0.552 $\pm 7.2e-3$	0.759 $\pm 3.9e-3$	0.652 $\pm 1.1e-2$
Negative sampling 2	0.706 $\pm 1.8e-2$	0.655 $\pm 1.6e-2$	0.615 $\pm 1.7e-2$	0.535 $\pm 1.4e-2$	0.911 $\pm 6.2e-2$	0.852 $\pm 2.2e-2$
Negative sampling 16	0.718 $\pm 4.2e-3$	0.656 $\pm 9.7e-4$	0.649 $\pm 2.3e-2$	0.509 $\pm 1.3e-2$	0.860 $\pm 1.0e-2$	0.772 $\pm 5.7e-2$

TABLE II

PERFORMANCE (AUC AND ACC) OF THE DEEP NETWORKS FOR THE HIREABILITY TASK, MAXIMUM PERFORMANCE (AUC) BETWEEN THE DIAGNOSTIC CLASSIFIERS (LR OR XGB) FOR THE IDENTIFICATION OF GENDER AND ETHNICITY FROM THE LATENT REPRESENTATION OF THE NETWORKS, AND DISPARATE IMPACT (DI) FOR GENDER AND ETHNICITY OF THE NETWORKS, FOR THE FOUR ARCHITECTURES.

maximum performance of the diagnostic classifiers, LR or XGB, for the task of predicting gender and ethnicity from the latent representation of the networks. Columns DI Gender and DI Ethnicity give the disparate impact of each classifier for the base network and each adversarially trained classifier. For each modality and fairness metric, best scores for indirect methods are shown in bold. We discuss the results as a series of questions and answers:

Do automatic predictions amplify the bias present in human annotation? We compare the DI columns for “unprotected” networks to the DI displayed in Table I, which is linked to the annotations of the dataset. Regarding gender, DI for audio, video and multimodal models is lower than with human annotations, hence the predictions are less fair than gold truth annotations. On the other hand, predictions based on language modality are fairer. Regarding ethnicity, the language and audio models show a higher disparate impact than actual human annotations, whereas video and multimodal models seem to degrade disparate impact. Interestingly, the multimodal model is the worst model regarding disparate impact whereas it presents the best performance regarding hireability.

Can we recognize gender and ethnicity given representations trained solely for hireability task? We discuss here the AUC scores of diagnostic classifiers for both gender and ethnicity on the unprotected networks. The high AUC values show that gender can be retrieved from the audio, video and multimodal models, but not from the language model. AUC scores for ethnicity are important only for video and, to a lesser extent, multimodal models, meaning that some information is retrievable. Finally, contrary to [23], our multimodal model does not show a higher leakage than our unimodal models.

What is the impact of supervised adversarial training on protected variables? As our work is the first to apply adversarial methods to AVIs, we comment on the effectiveness of classical adversarial methods. By comparing the AUC and accuracy values in the first columns between the unprotected method and the two supervised methods, we observe a marginal decrease in performance. For the audio, video and multimodal modalities, the decrease in AUC of the diagnostic classifiers shows that sensitive information is effectively removed from the latent representation. These scores provide a gold standard for the indirect methods we propose.

What is the impact of indirect adversarial methods on hireability performance and protected variables? We discuss here the columns Hireability, Gender and Ethnicity for the indirect methods we propose, Static faces and Negative sampling. For the audio modality, the indirect methods decrease the AUC of the diagnostic classifier for gender (except for the Static faces 16 method, discussed in the next question). This results in a score comparable to the supervised method in two cases, but also slightly affects the AUC for Hireability. Ethnicity is little affected, since the diagnostic classifiers have an AUC very close to 0.5 even for the Unprotected classifier. For the video modality, there is little gain for gender. However, for ethnicity, only the Negative sampling 2 method obtains an important drop in AUC from the diagnostic classifier, but this performance comes at the price of a drop in performance on the hireability prediction task. In the multimodal case, the indirect methods obtain good results with only a slight impact on the classifier’s performance.

What is this impact of dimensionality of the compressed representation? In all columns, there are differences in outcome among the 4 indirect adversarial approaches. In particular, we notice that approaches using a 2-dimensional representation for faces generally perform better than the same 16-dimensional approach. The diagnostic AUCs of the Unprotected and the Static faces 16 methods are very close, showing that the latter is particularly inefficient. On the contrary, the Negative sampling 2 method obtains generally low AUC diagnostic scores, and performs better than the other approaches for gender in language, for ethnicity in video and for the multimodal modality in general. However, it is also the model which obtains the worst performance on the Hireability task, particularly in the video and multimodal cases, reflecting the tradeoff between the two competing tasks.

What is the impact of indirect adversarial methods on the fusing step for the multimodal model? Modalities are fused by the GMU through a weighted sum of each modality. As the model is simple enough we can investigate the norm of the vector for each modality, which reflects their degree of contribution to the multimodal representation. We display in Figure 2 the boxplots of the contribution of each modality to the multimodal vector depending on indirect adversarial methods used. Here, we observe that, for the unprotected model, the visual modality contributes the most, followed by the audio and language modalities. We then observe that the adversarial training does have an impact on the fusing step for the multimodal model especially for the Static Face 2, NS 2 and NS 16 methods. This influence is reflected in a greater contribution of the language modality and in a smaller contribution of the visual modality. Interestingly, the Static Face 16 method does not show this influence reinforcing our thinking that this method performs worse than the other three. Overall, we can assume that adversarial training tends to decrease the contribution of the most problematic modality and strengthens the contribution of the modality least likely to leak sensitive information.

Does our approach improve Disparate Impact as well?




Fig. 2. Boxplots of the contribution of each modality constituting multimodal vector through the GMU depending on the different adversarial training methods. NS stands for Negative Sampling. The number stands for the dimension of the representation.

Although the approach we propose aims to reduce the potential disparity of treatment for automatic evaluation of AVIs, measuring the Disparate Impact shows that this improvement can also benefit group fairness. Thus, the last two columns of the Table II show that adversarial methods often improve the Disparate Impact (with the exception of the Supervised Ethnicity method) compared to the Unprotected network, and that this increase is often even greater for indirect approaches that have succeeded in removing sensitive information from the latent representation. Both types of fairness can thus coexist, to a certain extent.

V. CONCLUSION AND FUTURE WORK

Guaranteeing equal treatment of candidates is crucial during a job interview. Here, we evaluate the potential for discrimination in a state-of-the-art hiring system on a public dataset. Despite not using sensitive features as input, we show that gender and ethnicity are retrievable from the latent representation of the system. Taking into account the fact that some countries prohibit the collection of this protected information in the context of recruiting, we use adversarial learning to remove sensitive information from the latent representation using accessible information which can be easily extracted from the video interviews as a proxy – the face of the candidates. Comparing this to more standard adversarial baselines, we show that our approach succeeds in removing protected information from the latent representations of the models. Moreover, we observe experimentally that the proposed algorithms generally improve the disparate impact even if group fairness is not targeted by our methods. Overall, we would like to reiterate that automatic hiring decisions should not be made without any human intervention and we believe that our work is a first step towards better addressing ethical issues in automatic analysis of AVI. In future work, we aim to test our approach on different datasets (e.g. more realistic or balanced), or other critical tasks such as automatic emotion recognition.

ACKNOWLEDGMENT

I thank Peter Hagelund, Adrian Mihai and David Wright for their valuable assistance and William Coleman for proof-reading.

REFERENCES

- [1] S. Rasipuram and D. B. Jayagopi, "Automatic assessment of communication skill in interview-based interactions," *Multimedia Tools and Applications*, vol. 77, no. 14, pp. 18 709–18 739, 2018.
- [2] L. Chen, R. Zhao, C. W. Leong, B. Lehman, G. Feng, and M. E. Hoque, "Automated video interview judgment on a large-sized corpus collected online," in *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 10 2017, pp. 504–509. [Online]. Available: <http://ieeexplore.ieee.org/document/8273646/>
- [3] M. Raghavan, S. Barocas, J. Kleinberg, and K. Levy, "Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices," 6 2019. [Online]. Available: <http://arxiv.org/abs/1906.09208><http://dx.doi.org/10.1145/3351095.3372828>
- [4] L. S. Nguyen, D. Frauendorfer, M. S. Mast, and D. Gatica-Perez, "Hire me: Computational Inference of Hirability in Employment Interviews Based on Nonverbal Behavior," *IEEE Transactions on Multimedia*, vol. 16, no. 4, pp. 1018–1031, 6 2014. [Online]. Available: <http://ieeexplore.ieee.org/document/6746207/>
- [5] L. Hemamou, G. Felhi, V. Vandenbussche, J.-c. Martin, and C. Clavel, "HireNet: A Hierarchical Attention Model for the Automatic Analysis of Asynchronous Video Job Interviews," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 573–581, 7 2019. [Online]. Available: <https://aaai.org/ojs/index.php/AAAI/article/view/3832>
- [6] N. Schmitt, Ed., *The Oxford Handbook of Personnel Assessment and Selection*. Oxford University Press, 3 2012. [Online]. Available: <http://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199732579.001.0001/oxfordhb-9780199732579>
- [7] E. N. Torres and C. Mejia, "Asynchronous video interviews in the hospitality industry: Considerations for virtual employee selection," *International Journal of Hospitality Management*, vol. 61, pp. 4–13, 2017. [Online]. Available: <http://dx.doi.org/10.1016/j.ijhm.2016.10.012>
- [8] E. Kroll and M. Ziegler, "Discrimination due to Ethnicity and Gender: How susceptible are video-based job interviews?" *International Journal of Selection and Assessment*, vol. 24, no. 2, pp. 161–171, 6 2016. [Online]. Available: <http://doi.wiley.com/10.1111/ijsa.12138>
- [9] J. Sánchez-Monedero, L. Dencik, and L. Edwards, "What does it mean to 'solve' the problem of discrimination in hiring? Social, technical and legal perspectives from the UK on automated hiring systems," *FAT* 2020 - Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 458–468, 2020.
- [10] R. Lieberman, "A Tale of Two Countries: The Politics of Color Blindness in France and the United States," *French Politics, Culture & Society*, vol. 19, no. 3, 1 2001. [Online]. Available: <http://berghahnjournals.com/view/journals/fpcs/19/3/fpcs190303.xml>
- [11] S. Corbett-Davies and S. Goel, "The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning," no. Ec, 2018. [Online]. Available: <http://arxiv.org/abs/1808.00023>
- [12] I. Naim, M. I. Tanveer, D. Gildea, and M. E. Hoque, "Automated Analysis and Prediction of Job Interview Performance," *IEEE Transactions on Affective Computing*, vol. 9, no. 2, pp. 191–204, 4 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/7579163/>
- [13] H. J. Escalante, H. Kaya, A. A. Salah, S. Escalera, Y. Gucluturk, U. Guclu, X. Baro, I. Guyon, J. C. Jacques, M. Madadi, S. Ayache, E. Viegas, F. Gulpinar, A. S. Wicaksana, C. Liem, M. A. Van Gerven, and R. Van Lier, "Modeling, Recognizing, and Explaining Apparent Personality from Videos," *IEEE Transactions on Affective Computing*, vol. 3045, no. c, 2020.
- [14] C. W. Leong, K. Roohr, V. Ramanarayanan, M. P. Martin-Raugh, H. Kell, R. Ubale, Y. Qian, Z. Mladineo, and L. McCulla, "To Trust, or Not to Trust? A Study of Human Bias in Automated Video Interview Assessments," 2019. [Online]. Available: <http://arxiv.org/abs/1911.13248>
- [15] A. Singhanian, A. Unnam, and V. Aggarwal, "Grading video interviews with fairness considerations," 2020. [Online]. Available: <http://arxiv.org/abs/2007.05461>
- [16] T. Baltrusaitis, A. Zadeh, Y. C. Lim, and L. P. Morency, "OpenFace 2.0: Facial behavior analysis toolkit," *Proceedings - 13th IEEE International Conference on Automatic Face and Gesture Recognition, FG 2018*, pp. 59–66, 2018.
- [17] F. Eyben, F. Weninger, F. Gross, and B. Schuller, "Recent developments in openSMILE, the munich open-source multimedia feature extractor," *Proceedings of the 21st ACM international conference on Multimedia*

- *MM '13*, no. May, pp. 835–838, 2013. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2502081.2502224>

- [18] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, “Learning phrase representations using RNN encoder-decoder for statistical machine translation,” *arXiv preprint arXiv:1406.1078*, 2014.
- [19] T. Adel, I. Valera, Z. Ghahramani, and A. Weller, “One-Network Adversarial Fairness,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 2412–2420, 2019.
- [20] Q. Xie, Z. Dai, Y. Du, E. Hovy, and G. Neubig, “Controllable invariance through adversarial feature learning,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017.
- [21] T. Wang, J. Zhao, M. Yatskar, K. W. Chang, and V. Ordonez, “Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations,” *Proceedings of the IEEE International Conference on Computer Vision*, vol. 2019-Octob, pp. 5309–5318, 2019.
- [22] R. Aloufi, H. Haddadi, and D. Boyle, “Privacy-preserving Voice Analysis via Disentangled Representations,” 2020. [Online]. Available: <http://arxiv.org/abs/2007.15064>
- [23] M. Jaiswal and E. Mower Provost, “Privacy Enhanced Multimodal Neural Representations for Emotion Recognition,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, pp. 7985–7993, 4 2020. [Online]. Available: <http://arxiv.org/abs/1910.13212https://aaai.org/ojs/index.php/AAAI/article/view/6307>
- [24] J. Streiff-Fénart, “A French dilemma: Anti-discrimination policies and minority claims in contemporary France,” *Comparative European Politics*, vol. 10, no. 3, pp. 283–300, 7 2012. [Online]. Available: <http://link.springer.com/10.1057/cep.2012.14>
- [25] B. Schuller, S. Steidl, A. Batliner, A. Vinciarelli, K. Scherer, F. Ringeval, M. Chetouani, F. Weninger, F. Eyben, E. Marchi, M. Mortillaro, H. Salamin, A. Polychroniou, F. Valente, and S. Kim, “The INTER-SPEECH 2013 computational paralinguistics challenge: Social signals, conflict, emotion, autism,” *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, no. August, pp. 148–152, 2013.
- [26] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, vol. 1, no. Mlm, pp. 4171–4186, 2019.
- [27] J. Arevalo, T. Solorio, M. Montes-y Gómez, and F. A. González, “Gated multimodal networks,” *Neural Computing and Applications*, vol. 32, no. 14, pp. 10 209–10 228, 7 2020. [Online]. Available: <http://arxiv.org/abs/1702.01992http://link.springer.com/10.1007/s00521-019-04559-1>
- [28] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, “ArcFace: Additive angular margin loss for deep face recognition,” *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2019-June, pp. 4685–4694, 2019.
- [29] Y. Ganin and V. Lempitsky, “Unsupervised domain adaptation by backpropagation,” *32nd International Conference on Machine Learning, ICML 2015*, vol. 2, no. i, pp. 1180–1189, 2015.
- [30] L. McInnes, J. Healy, N. Saul, and L. Großberger, “UMAP: Uniform Manifold Approximation and Projection,” *Journal of Open Source Software*, vol. 3, no. 29, p. 861, 9 2018. [Online]. Available: <http://arxiv.org/abs/1802.03426http://joss.theoj.org/papers/10.21105/joss.00861>
- [31] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, vol. 13-17-Aug, pp. 785–794, 2016.
- [32] S. Poria, N. Mazumder, E. Cambria, D. Hazarika, L. P. Morency, and A. Zadeh, “Context-dependent sentiment analysis in user-generated videos,” *ACL 2017 - 55th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*, vol. 1, pp. 873–883, 2017.

APPENDIX

A. Formalization

In our model, we denote each instance by the triplet X, Y, Z where X denotes the candidate’s answer, Y the hireability label of this candidate and Z the protected variables which must not be used for inference of Y . Let $X_{\{L,V,A\}} \in \mathbb{R}^{T_{\{L,V,A\}} \times d_{\{L,V,A\}}}$ the sequence of low level descriptors describing the candidate’s answer respectively from the three modalities: language (L), video (V) and audio (A). T_m and d_m represent the sequence length and feature dimension of the modality $m \in \{L, V, A\}$ as each modality has a different sampling frequency and feature dimension. We consider our problem for the main task as a classification problem: Y is a binary variable (i.e. hireable or not hireable). Finally protected variable Z is a categorical variable (e.g. gender or ethnicity).

B. Architecture

The goal of our model is to maximize the performance of the main task (e.g. hireability prediction) while minimizing the performance to predict Z . In that sense, our system is composed of three parts: a) the main network which encodes information from the multimodal answer X to a latent representation H , b) the hireability classifier which classifies the answer based on the representation H and c) the adversarial branch network which tries to predict Z based on the representation H .

Base network for hireability. Previous literature [5] suggests the relevance of recurrent neural networks for modeling the sequentiality of job video interviews. Inspired by such approaches we encode each modality separately by a specific modality encoder in order to obtain a better representation of each element for the intra-modality representation [32]. For that purpose, we use a recurrent neural network component, namely a Bidirectional Gated Recurrent Unit (BiGRU).

$$z_t^m = BiGRU(x_t), t \in [1, T_m] \quad (6)$$

In order to obtain a fixed size vector for each modality and to account for salient moments in the candidate’s answer [5], we use an additive attention mechanism, described by the following equations:

$$\begin{aligned} u_t^m &= \tanh(W_A^m z_t^m + b^m) \\ \alpha_t^m &= \frac{\exp(u_p^{m\top} u_t^m)}{\sum_{t'} \exp(u_p^{m\top} u_{t'}^m)} \\ o^m &= \sum_t \alpha_t^m z_t^m \end{aligned} \quad (7)$$

where W_A^m are weight matrices, u_p^m and b^m are weight vectors and $u_p^{m\top}$ denotes the transpose of u_p^m respectively for $m \in \{L, A, V\}$.

Previous work on job interview analysis using deep learning approaches [5], [14] do not offer satisfactory multimodal representation of the inputs. We thus propose to improve the multimodal representation by fusing the three modalities through the use of a Gated Multimodal Unit (GMU, [27]). The GMU works by projecting all modalities on the same

space and learning the contribution of each through a gated mechanism. We propose to use it on top of the attention mechanism of each modality.

$$\begin{aligned}\tilde{o}^a &= \tanh(W_{Aproj}o^a) \\ \tilde{o}^l &= \tanh(W_{Lproj}o^l) \\ \tilde{o}^v &= \tanh(W_{Vproj}o^v) \\ \sigma^a &= \sigma(W_{Agating}[o^a, o^l, o^v]) \\ \sigma^v &= \sigma(W_{Lgating}[o^a, o^l, o^v]) \\ \sigma^l &= \sigma(W_{Vgating}[o^a, o^l, o^v])\end{aligned}$$

$$o^{multimodal} = \sigma^a * \tilde{o}^a + \sigma^l * \tilde{o}^l + \sigma^v * \tilde{o}^v \quad (8)$$

where W_{Aproj} , W_{Lproj} , W_{Vproj} , $W_{Agating}$, $W_{Lgating}$, $W_{Vgating}$ are weight matrices, σ is the sigmoid function and $[x, y]$ is the concatenation of x and y . Output of the multimodal unit is then given by $o^{multimodal}$, which represents the multimodal information of the answer.

On top of the multimodal representation, we choose to use a simple combination of two dense layers, which is a simplification of the architectures found in the literature and built on the structure of AVIs [5], [14]:

$$h_1 = \tanh(W_1^\top o^{multimodal} + b_1) \quad (9)$$

$$H = \tanh(W_2^\top h_1 + b_2) \quad (10)$$

where W_1 and W_2 are weight matrices and b_1 and b_2 are weight vectors. We denote by θ_H the parameters of the main network.

Hireability classifier. Once H is obtained, we use it as representation in order to infer hireability:

$$\hat{Y} = \sigma(W_v H + b_v) \quad (11)$$

where W_v is a weight matrix and b_v a weight vector. We denote by θ_D parameters of the hireability classifier.

As the problem we are facing is that of a binary classification, the loss function of the main task is:

$$\mathcal{L}_T = -\frac{1}{N} \sum_{i=1}^N Y_i \log \hat{Y}_i + (1 - Y_i) \log(1 - \hat{Y}_i)$$

where Y denotes the true labels of hireability of candidates.

Supervised Adversarial Branch Network. In order to force the network to forget sensitive information, we use an adversarial approach to train a second branch on top of the latent representation H , which learns to recover sensitive information about the candidates. This second branch is grafted on the main network through a Gradient Reversal Layer (GRL, [29]). A GRL is a special layer with no weight vector or matrix, which behaves as the identity function during the forward pass, but negates the gradient during the backward pass, resulting in the network *unlearning* the corresponding information. In order to obtain a fairer representation, we decide to use a deeper network than the hireability branch. Also, as type of

protected variables change, we use different activation or loss function. More precisely:

$$\hat{Z} = \text{softmax}(W_4 \sigma(W_3 H + b_3) + b_4)$$

where W_3 , W_4 are trainable weight matrices and b_3 , b_4 are trainable vectors. We denote by θ_A the parameters of the adversarial branch. Then, the loss of the adversarial branch is as follow :

$$\mathcal{L}_A = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^C Z_{ij} \log(\hat{Z}_{ij})$$

where C denotes the number of classes for categorical Z (e.g. 2 for gender and 3 for ethnicity in our case).

Following the notations used in the main article, the complete model is trained by optimizing the following min-max objective:

$$\min_{\theta_H, \theta_D} \max_{\theta_A} \mathcal{L}_T(\theta_H, \theta_D) - \lambda \mathcal{L}_A(\theta_H, \theta_A)$$

where $\lambda \geq 0$ is a trade-off hyperparameter between the hireability objective and the adversarial objective.

C. Indirect adversarial approach method 1 : static representation of the faces

We describe in this subsection the function f applied to the representation H referred in section 3.2 "First method: static representation of the faces" of the main article.

f consists in two dense layers, more precisely :

$$f(H) = W_6(W_5 H + b_5) + b_6$$

where W_5 , W_6 are trainable weight matrices and b_5 , b_6 are trainable vectors.

D. Indirect adversarial approach method 1 : static representation of the faces

We describe in this subsection the function g applied to each pair $(H, W_l)_{l=1..k}$ referred in Section 3.2 "Second method: negative sampling" of the main article.

g consists in three sub-components, more precisely :

$$\hat{W}_l = g_1(W_l) = \tanh(W_7 W_l + b_7)$$

g_1 is responsible for encoding the face representation in a low dimensional space. W_7 is a trainable weight matrix and b_7 is a trainable vector.

$$\hat{H} = g_2(H) = \tanh(W_9(W_8 H + b_8) + b_9)$$

g_2 is responsible for encoding the interview representation. W_8 , W_9 are trainable weight matrices and b_8 , b_9 are trainable vectors.

$$g(W_l, H) = g_3(\hat{W}_l, \hat{H}) = W_{10}(\hat{W}_l \odot \hat{H} + b_{10})$$

g_3 is responsible for estimate similarity measure between \hat{W}_l and \hat{H} . W_{10} is a trainable weight matrix and b_{10} is a trainable

vector, \odot stands for the Hadamard product (element-wise product).

We have followed this training strategy in order to train the model :

- 1) First we train the main network and the hirability branch (θ_H and θ_D).
- 2) Then, we train only the adversarial branch (θ_A)
- 3) The adversarial training is then organized by alternatively train the full network (θ_H, θ_D , and θ_A) or only the adversarial branch (θ_A)
 - a) Train the whole network (main network, hirability branch, adversarial branch on top of the GRL) for one epoch (θ_H, θ_D , and θ_A)
 - b) reset weights of the adversarial branch (θ_A)
 - c) train adversarial branch until any improvement on the validation loss (θ_A)

We reiterate this loop until the validation loss defined by

$$\min_{\theta_H, \theta_D} \max_{\theta_A} \mathcal{L}_T(\theta_H, \theta_D) - \lambda \mathcal{L}_A(\theta_H, \theta_A)$$

does not improve.

All the experiments, including the pretraining of base models for hireability, are ran 3 times and results are averaged. As mentionned in the attached code archive, all the experiments have been ran on a 32-core CPU server with 93 GB of RAM and P100 NVidia Tesla GPU with 16 GB of VRAM. The total runtime is between 3 weeks and 1 month.

E. Hyperparameters of the model

We describe here the hyperparameters (optimizer, layers sizes...) used for the architecture detailed in Section 4.2 in the main paper. In order to find the best hyperparameters, we conducted a grid search on the first task (hireability prediction), but observed little variation in the results. The final hyperparameters are displayed in Table III, using the same notations as in the paper.

Note that size of W_4 and b_4 , which describe the final dense layer of the adversarial branch, vary depending on the protected variable: 1 for gender (sigmoid output), 3 for ethnicity (softmax output).

During the adversarial training, we observed that the λ hyperparameter had a somewhat important effect on the performance. For the experiments, the grid used for the lambda value was $\{0.5, 1, 2, 5, 10\}$. For a given modality, the selected λ parameter is the one corresponding to the minimum final values of the loss function $\mathcal{L}_T - \mathcal{L}_A$ on the validation set, which represents a trade-off between the hireability and the privacy-inducing tasks. Table IV displays the value chosen for all the models.

F. Hyper parameters for diagnostic classifiers

Two diagnostic classifiers are used to retrieve the protected variables from the network latent representation, Logistic Regression (LR) and XGBoost. During our initial experiments, we conducted a grid search on the hyperparameters of the two algorithms, but observed little variation on the performance of

Hyper parameter	Value or Dimensions
Optimizer	Adam
Units of intra-modality Bi-GRU	16
W_A^m, u_p^m and b^m	[30,16],30,30
$W_{Aproj}, W_{Lproj}, W_{Vproj}$	[48,16],[48,16],[48,16]
$W_{Agating}, W_{Lgating}, W_{Vgating}$	[1,48],[1,48],[1,48]
W_1, b_1	[16,48],16
W_2, b_2	[16,16],16
W_v, b_v	[1,16],1
W_3, b_3	[30,16],30
W_4, b_4 depending on the output	[(1/3),30],[1/3]
W_5, b_5	[30,16],30
W_6, b_6 depending on the output	[(2/16),30],[2/16]
W_7, b_7 depending on the hidden size chosen	[(2/16),512],[2/16]
W_8, b_8	[32,16],32
W_9, b_9 depending on the output	[(2/16),32],[2/16]
W_{10}, b_{10} depending on the hidden size chosen	[1,(2/16)],1
Batch size	32
Regularization L2	10^{-4}
Dropout	0.2
Gradient Clip	1.0
Learning rate joint training	0.0001
Learning rate adversarial branch	0.003

TABLE III
HYPER PARAMETERS USED FOR TRAINING THE MODEL

XGBoost classifier. Because the hyperparameter search was very extensive for this classifier, we decided to select a set of average values for the parameters: loss reduction split γ is fixed to 0.1, regularization parameter $\alpha \in \{0.3, 0.5\}$ and number of estimators is 500. For LR, we observed more variation in the output depending on the regularization parameter. We thus run a hyperparameter search each time, selecting the best value in the range $10^{-4} \dots 10^4$, depending on the performance on the validation set, and both ℓ_1 and ℓ_2 norms were tested.

Although the ChaLearn dataset comes with a proposed 3-way split between training, validation and test sets, we found that 84 % of clips from the test set have at least one clip in the train set extracted from the same Youtube video. This overlap is potentially problematic, since it could allow classifiers to obtain good results on the hireability prediction task by learning speakers' specific features, e.g. facial or audio, a strategy that makes no sense from the point of view of recruitment. For this reason, we use our own 3-way split with no overlap. Table V presents the statistics of the protected variables for the new split.

G. Experiments on the original dataset

As an illustration of the differences between the two sets, we conduct two experiments. Firstly, we show that it is possible to obtain good results to the original ChaLearn challenge by using a naive baseline. For this, we use the original test set and, for each video, predict an hireability score by computing the average of the original hireability (binary) labels for candidates

Protected variable	Language	Audio	Video	Multimodal
λ Static faces 2	5.0	2.0	2.0	10.0
λ Static faces 16	10.0	5.0	2.0	1.0
λ Negative sampling 2	5.0	10.0	10.0	2.0
λ Negative sampling 16	10.0	5.0	2.0	1.0
λ Supervised gender	0.5	0.5	1.0	0.5
λ Supervised ethnicity	10.0	0.5	1.0	10.0

TABLE IV
SUMMARY OF LAMBDA CHOSEN DEPENDING ON MODALITY AND PROTECTED VARIABLES

Protected variable	Complete dataset	Training set (6991 clips)	Validation set (1448 clips)	Test set (1559 clips)
Gender - Female	0.560	0.560 (3053)	0.575 (702)	0.541 (783)
Gender - Male	0.495	0.482 (3938)	0.524 (746)	0.519 (776)
Disparate Impact	0.883	0.860	0.911	0.959
Ethnicity - Asian	0.570	0.584 (236)	0.718 (32)	0.444 (63)
Ethnicity - Caucasian	0.541	0.537 (5998)	0.545 (1273)	0.554 (1327)
Ethnicity - Afro-Americans	0.434	0.424 (757)	0.559 (143)	0.374 (171)
Disparate Impact	0.761	0.726	0.759	0.675

TABLE V
SUMMARY OF INITIAL BIAS ON THE PROPOSED SPLIT

	LR		XGBoost	
	ACC	AUC	ACC	AUC
Original split	0.680	0.742	0.745	0.811
Proposed split	0.650	0.695	0.622	0.672

TABLE VI
DIFFERENCES BETWEEN THE TWO SPLITS OF THE DATASET FOR
HIREABILITY PREDICTION USING ONLY STATIC FACES.

also present in training set. Missing values (candidates present in the test set only) are replaced with the average label on training and validation set. Comparing these values to the original labels of the test set, we obtain an AUC of 0.797 and an accuracy of 0.7215, which are comparable to the performance of most classifiers reported during the original challenge. Secondly, we extract face representations of the candidates and train two classifiers (LR and XGBoost) on each of the splits. We report evaluation metrics regarding hireability in table VI. Classifiers trained on the original split have a higher score (AUC of 0.811) than the score obtained by classifiers trained and evaluated on the proposed split (AUC of 0.695): we interpret this result as proof that the sharing of videos occurring in the original split makes the task easier and could potentially reduce the hirability task to a person's identity recognition task.

H. Hirability Performance depending on splits

Finally, we report in Table VII the performance of the 4 models studied in the paper on the original and proposed splits. We observe that the drop of performance between the two sets is not very important, except for the video modality. We interpret this as proof that the classifiers do not rely too much on information shared between the two sets and that they are pretty stable. This impression is strengthened by comparing these results to the Table VI.

	Language		Audio		Video		Multimodal	
	ACC	AUC	ACC	AUC	ACC	AUC	ACC	AUC
Original split	0.601	0.607	0.664	0.721	0.686	0.751	0.701	0.766
Proposed split	0.584	0.610	0.642	0.695	0.647	0.700	0.682	0.741

TABLE VII
RESULTS FOR HIRABILITY TASK