

# Mitigating Age Biases in Resume Screening AI Models

Christopher G. Harris

Department of Mathematical Sciences, University of Northern Colorado, Greeley, CO 80639  
christopher.harris@unco.edu

## Abstract

As populations age, an increasing number of workers beyond the traditional retirement age are opting to continue working. Nevertheless, discrimination against older job seekers seeking new employment opportunities remains widespread. To address this issue, we enlisted a pool of crowdworkers to assess the resumes of IT job candidates and guess each candidate's age, race, and gender. Using this crowdsourced data, we trained an AI model and applied bias correction techniques from IBM's AI 360 and Microsoft's Fairlearn toolkits to correct for biases based on race, gender, and age. We analyzed the effectiveness of these tools in mitigating different types of bias in job hiring algorithms, explored why age may be more challenging to eliminate than other forms of bias, and discussed additional approaches to enhance fairness. Our results indicate that implicit age bias, or ageism, is prevalent in hiring decisions and more pervasive than other well-documented forms of bias, such as race and gender biases.

## Introduction

Nearly all job searches in the industrialized world today involve an online component – either to search for jobs or candidates, to upload a resume or CV, or to communicate with prospective employees or employers. The online job platform Indeed.com, for example, has more than 200 million visits per month, contains 100 million resumes, and adds 9.8 new jobs per second (Marino, 2022). Over eighty percent of job searches in the US are done online (Kolmar, 2022).

With much of the job search process moving online, it has become possible for both companies and job seekers to increase their search efficiency and reduce costs. However, the ease of applying for jobs online has led to a large volume of applicants, with the average online job posting receiving 250 applicants. Still, only 6-10 candidates are typically selected for interviews, providing a yield of only two to four percent (Clark, 2022). This low yield is mainly due to over fifty percent of online job applicants not meeting the job posting's criteria (Kolmar, 2022). To address this challenge, algorithms that leverage artificial intelligence (AI) and natural language processing (NLP) have become more prevalent. This combination of AI and NLP is well-

sued for screening job applicants and determining the salary and terms of selected candidates.

This use of AI algorithms also has another potential benefit. Biases against people of color, women, and other underrepresented groups have long plagued the traditional hiring process (Bogen & Rieke, 2018). Therefore, these algorithms have also been designed to eliminate (or substantially reduce) the biases held by hiring managers or human resources personnel (Raub, 2018). In theory, with AI, hiring decisions can be made on the merits of each candidate, not on biases based on race, gender, or any other protected status (Vrontis et al., 2021).

AI algorithms learn from their inputs, including biased human decisions or historical inequities. Unfortunately, these biases are often replicated by the algorithms themselves. For example, in 2015, Amazon abandoned an AI-powered recruiting tool that exhibited hiring biases against women for technical roles, as the algorithms were trained on resumes primarily from male candidates (Dastin, 2018). Even if features like gender, race, or sexual orientation are removed, this issue can still arise (Manyika, Silberg, & Presten, 2019). Amazon's hiring algorithm favored words such as "executed" or "captured," which were more common on men's resumes.

While algorithmic biases have been widely discussed in the media, another type of hiring bias, ageism, has not received as much attention. Many governments have raised retirement ages and limited early retirement options as people live longer. Labor laws in most developed countries prohibit firing workers based on their age (Martin et al., 2014), but several studies (e.g., Oude Mulders et al., 2014; Karpinska, Henkens, & Schippers, 2013) have shown that age discrimination during the hiring process prevents older job seekers from finding new jobs. Opportunities for older workers to obtain new employment are meager.

Ageism can occur due to either implicit or explicit biases. An *implicit bias* is a collection of attitudes that the holder is not consciously aware of having. An *explicit bias* is a collection of perspectives the holder is aware of and can express consciously (Daumeyer, 2019). Thus, age discrimination is typically viewed as explicit, whereas age bias is generally implicit. While the alteration of job screening and job hiring algorithms may fix explicit biases, it is far more challenging to fix implicit biases once they become part of the screening algorithm (Soleimani, 2021).

The paper is organized as follows. In the next section, we describe some related studies on age discrimination in the hiring process and biases in job search algorithms. We then describe our experimental methods for detecting biases and explain the results of this experiment. Next, we discuss how age biases are notoriously tricky to remedy and some strategies to address them. Last, we summarize our work in mitigating age bias.

## Background and Motivation

Although not as widely researched as other biases (e.g., gender and race), age bias has become a more prominent subject of study in social psychology. Age biases in the workplace can lead to adverse outcomes for employees (e.g., poor job morale and performance), employers (e.g., higher turnover, more lawsuits), and a degraded employee-employer relationship (Zacher & Steinvik, 2015).

### Ageism in society

Ageism exists across multiple facets of society, including in e-commerce (McIntosh, 2021), healthcare (Wiens, Price & Sjoding, 2020), and housing (Rich, 2014). Ageism is not unique to industrialized countries – it also exists across cultures. A sixteen-year study of 911,982 participants from 68 countries examined implicit and explicit age bias and differences in subjective age and feelings of warmth toward younger and older adults (Ackerman & Chopik, 2021). After reviewing numerous attributes, such as GDP, education level, and gender, this study found that living in countries labeled as highly collectivistic (according to Hofstede's cultural dimensions (Hofstede, 2005) was associated with less implicit and explicit age bias than highly individualistic ones. However, ageism exists in all cultures studied. Implicit biases against older adults were also far more prominent than explicit biases, making age biases and discrimination harder to monitor and control.

### Ageism in traditional hiring decisions

Previous research indicates that as much as 90 percent of all discriminatory decisions occur before the interview (Bovenkerk, 1992). Rarely are judgments considered completely objective and/or neutral regarding recruitment processes. Even implicit biases and will the most rely on professional and well-intended HR personnel can have their expert "gut feelings," which certainly applies to ageism (Rivera, 2015).

Stereotypes about the elderly are a primary hindrance to hiring older workers and have been present for centuries. (Karaoglu, Hargittai & Nguyen 2021). When job tasks historically involved physical labor, this age bias reflected the inability of older workers to remain productive under demanding physical conditions. However, few jobs today have challenging physical requirements, yet this bias remains.

A study (Richardson, 2013) asked 156 participants to evaluate a hypothetical job applicant's work-related competency and the likelihood of being hired. The study found that applicants over 54 were the least likely to be hired. They indicated this bias was primarily due to the raters' unwarranted concerns about the trainability and sociability of older workers and concluded that there was an explicit bias against older workers. Another more recent study (Neumark, Burn & Button, 2019). evaluated 40,000 job applications and found strong evidence of age biases in the hiring process, finding it was particularly acute for women and those near retirement age.

One may be led to believe that age discrimination and biases can be hidden by concealing age cues, but research indicates the opposite effect. Derous and Decoster (2017) discovered that anonymous resume screening does not dissuade age-discriminatory effects. Using a mixed factorial design, 610 HR professionals showed hiring discrimination against older applicants based on implicit age cues in resumes. Moreover, concealing one's date of birth led to overall lower ratings due to an impression of deception by withholding information.

### Ageism in algorithmic hiring decisions

Although a large body of literature has examined sexism and racism in AI recruitment algorithms, the unique effects of ageism by these systems have gained relatively little scholarly attention to date (Zavoronkov et al., 2019). AI-based recruitment can propagate unconscious biases in organizations' hiring processes (Stuart & Norvig, 2016). The vocabulary and method of expression in humans have been found to determine age and other demographical information (Schwartz et al., 2013). Since resumes and cover letters are the primary methods applicants communicate, the words that appear in these documents are significant factors in whether a candidate is selected (Upadhyay & Khandelwal., 2018). Research has shown that applying AI algorithms to everyday human language reproduces existing societal biases (Carlsson & Eriksson, 2019).

Consequently, these biases are introduced to the algorithmic decision-making process when the resume text is algorithmically analyzed. This allows these language features to be recognized, even when age is not explicitly indicated (Upadhyay & Khandelwal., 2018). We examine the language used in a dataset of information technology (IT) related job listings and resumes and see how it varies for older and younger workers.

Awareness of a job position's availability is also a significant issue since it limits the ability of elderly applicants to apply for positions due to technology challenges (e.g., Carlsson & Eriksson, 2019; Kaufmann et al., 2017; Caliskan et al., 2017). Algorithms used to reduce the cost of job advertisements have been found to show fewer opportunities to women than to men and older applicants than

younger ones, creating a selection bias (Lambrecht & Tucker, 2019). One component of the digital divide is the gap between those with access to digital technology and those without, which skews toward older adults.

In addition to the language used in job postings and more limited access to online jobs, AI algorithms' tasks in the job screening process also disproportionately affect older people. For example, AI algorithms can skim data posted on social media and obtain access to an applicant's values, attitudes, personality traits, and age (Faliagka, 2012). This data can be fed into the job candidate selection algorithm as features further aggravating potential age biases. To examine this in more detail, we conduct an experiment which is described in the next section.

## Experimental Methods

Field experiments are the most credible evidence of age-related discrimination in hiring (Richardson, 2013). We conducted a study to examine the relative effects of age, gender, and race discrimination and the impact of AI algorithms in job screening and hiring.

### Job Listing Data

We obtained a job listing for a mid-level IT position randomly selected from a dataset of 50,000 online job listings on Reed UK<sup>1</sup>. We limited our search to full-time positions that provided an annual salary range in the greater London area.

We then obtained candidate CVs by scraping individual CVs from the livecareer.com website. We randomly selected 12 CVs that listed as an objective a role like the job posting we obtained from Reed UK. Each of the 12 CVs was vetted by three executive recruiters (with an average of 11 years of recruiting experience for IT jobs) and determined to be suitable for selection for an initial interview for the job position we had randomly selected.

For each of the 12 CVs, we modified some information on the CV (see Table 1) to represent the following:

- *three races*: White (W), Black (B), or Asian (A)
- *two genders*: Male (M) or Female (F)
- *two age ranges*: Young with an average age of  $35 \pm 5$  years (Y) or Old with an average age of  $55 \pm 5$  years (O).

Each combination of (race) x (gender) x (age range) was applied to each of the 12 CVs, providing for 144 different combinations, which allowed for a comparative assessment across these three protected categories.

### Evaluation by Human Raters

We hired 144 human raters through Amazon Mechanical Turk (AMT) to make hiring decisions on employees. In

this within-group study design, all human raters were fluent in English and physically located in North America, Europe, or Australia. Each rater was shown the 12 CVs, but where the race, gender, and age range were randomly assigned to each CV, but each rater was offered one of each combination using a Latin square design. Raters were asked to indicate whether they would invite the candidate in for an interview for the job position based on the CV information alone. We did not provide race and gender as explicit information but provided subtle but observable clues as to the race/gender/age of each candidate (see Table 1 for some examples). Raters were shown text from the UK's Equality Act 2010, which describes the legal protections against discrimination in hiring against gender, race, and age.

Bias	Examples of clues given to raters
Race	Use of first and last names that tend to be viewed as more ethnic, as determined by external data [36] Membership in ethnic organizations (e.g., Member of the Asian Students Association, African Federation of Students) as activities while at university Work with charities that involve specific races as extracurricular activities or interests
Gender	Use of first and last names that tend to be viewed as more masculine or feminine [37] Membership in organizations viewed as traditionally masculine (e.g., member of the men's rugby team) or feminine (e.g., member of women's volleyball team)
Age	Cumulative job experience exceeding 20 years Graduation date from university was more than 20 years ago

Table 1. Sample Of Race, Gender, And Age Clues in Resumes

To reduce the impact of the learning curve for each rater, we implemented a calibration process where they compared their evaluations to the majority decision of three HR experts. This process was done using a separate set of CVs for training purposes. Before starting the main task, raters had to match the majority decision of the three experts in four out of five consecutive evaluations.

## Results

### Evaluation by Human Raters

Because the data met the assumptions for parametric tests, we conducted a three-factor ANOVA test on the decision to invite the candidate for an interview based on age, gender, and race. We took the average selection percentage for each of the 12 combinations of age, gender, and race. The results are given in the first column of Table 2.

We found no interactions between the three factors; however, there was a statistically significant difference between ages as determined by one-way ANOVA ( $F(9,134) = 4.501, p < 0.0001$ ).

After completing the task of evaluating resumes, the raters were asked to fill out a questionnaire that inquired whether they were aware of the age, race, and gender of the

<sup>1</sup> <https://www.kaggle.com/datasets/jobspikr/50000-job-board-record-from-reed-uk>

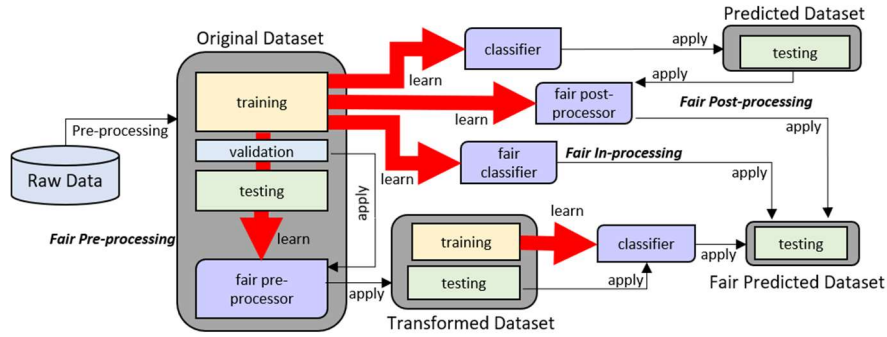


Fig. 1. Data flow of pre-processing, in-processing, and post-processing methods of the AI Fairness 360 Toolkit.

candidates based on the information provided in the resumes. The results showed that a nearly equal percentage of raters were aware of these factors, with 69% for age, 67% for race, and 72% for gender. Despite this, the implicit bias of the raters led to a significantly lower selection rate for older workers compared to other protected classes.

### Evaluation by AI Algorithm

The evaluations by the human raters validated that age was prominent in the decision to invite candidates for an interview. Similar to the work by Harris (2022), we wanted to see that these biases would carry over to an AI algorithm. We built a deep neural network to analyze the ratings based on the human raters' outputs to accomplish this. We had the algorithm evaluate 32 unseen resumes as our test set and used the results from the human raters as our training set. We developed a five-layer dense model with four hidden layers and one output layer. For an activation function, we used *relu* for the hidden layers, a "normal" kernel initializer, and a linear activation function in the output layer. For hyperparameters, our model used a dropout of 0.2 and a learning rate of 10-4. We divided the 1728 outputs into 2/3 train and 1/3 test.

Bias Category		Candidate Selection %		
Factor	Category	Human Rater	AI Algorithm	$\Delta$
Race	Asian (A)	0.620	0.636	+0.016
	Black (B)	0.646	0.643	- 0.003
	White (W)	0.661	0.648	- 0.013
Gender	Female (F)	0.597	0.622	+0.025
	Male (M)	0.688	0.663	- 0.025
Age	Young (Y)	0.753	0.758	+0.006
	Old (O)	0.531	0.525	- 0.006
Average		0.642	0.642	0.000

Table 2. Average Percentage of Candidates Selected, By Race, of Human Raters and Human-Trained AI Algorithm

Not surprisingly, the biases did carry over (see the AI Algorithm column in Table 2). Again, as a result of the AI algorithm, we found no interactions between the three factors; however, there was a statistically significant difference between ages as determined by one-way ANOVA

( $F(9,134) = 4.663, p < 0.0001$ ). This observation is substantiated by several other studies (e.g., Büsch, 2009; Köchling & Wehner, 2020; Silberg & Manyika, 2019; Raub, 2018). We observed that age bias became slightly more pronounced in the AI algorithm than in the human rater model (changes in all three factors were not statistically significant).

### Mitigation by AI Fairness Algorithms

A fairness metric is satisfied if a model's classification results are not dependent on a given sensitive attribute. A fairness-oriented conceptualization of construct validity is provided by (Jacobs and Wallach 2021). This conceptualization assists in analyzing fairness in sociotechnical contexts. We applied two open-source toolkits to mitigate age bias in job screening. We looked at three metrics:

- *Statistical Parity Difference* - measures the difference between the majority and protected classes receiving a favorable outcome. A value of 0 is considered fair.
- *Disparate Impact* – measures the difference in favorable proportions in the predicted labels metric as a ratio. A value of 1 indicates demographic parity.
- *Theil Index* - measures inequality as an entropic distance. It is the same as redundancy in information theory, which is the maximum possible entropy of the data minus the observed entropy. A value of 0 is considered fair.

### AI Fairness 360 Toolkit

First, we use the AI Fairness 360 toolkit (Bellamy et al., 2019). This toolkit provides several methods that involve pre-processing (modifying the training data), in-processing (modifying the learning algorithm), and post-processing (changing the predictions). Fig. 1 provides an overview of these methods.

Harris (2020) conducted an experiment assessing 150 applicants' suitability for three job openings. He examined the AI Fairness 360 toolkit and identified the most effective techniques for HR assessment: disparate impact reduction for pre-processing, adversarial debiasing in the in-

processing, and equalized odds post-processing in the post-processing. He achieved the best outcomes by using all three approaches in tandem, which we also adopted in our study. Table 3 demonstrates how these selected methods can decrease but not completely eliminate this bias (for a more detailed explanation of each fairness metric, refer to Bellamy et al., 2019 and Dwork, 2012). The numbers in italics indicate fairness based on the definitions by Bellamy et al., 2019.

Method Applied	Change in Selection Percentage	Statistical Parity Difference	Disparate Impact	Theil Index
<i>None</i>	Y = 0.758, O = 0.525, $ \Delta  = \mathbf{0.000}$	0.308	1.44	0.33
<i>A. disparate impact reduction (pre-processing)</i>	Y = 0.692, O = 0.591, $ \Delta  = \mathbf{0.066}$	0.146	1.17	0.22
<i>B. adversarial debiasing (in-processing)</i>	Y = 0.701, O = 0.582, $ \Delta  = \mathbf{0.057}$	0.170	1.24	0.24
<i>C. equalized odds (post-processing)</i>	Y = 0.683, O = 0.600, $ \Delta  = \mathbf{0.075}$	0.121	1.38	0.27
<i>Combination of methods A, B, and C</i>	Y = 0.654, O = 0.629, $ \Delta  = \mathbf{0.104}$	<i>0.069</i>	<i>1.07</i>	<i>0.16</i>

Table 3. Reduction in Age Bias by IBM AI Fairness 360

From Table 3, we can see that combining the three methods increases the fairness due to age but does not eliminate it. We believe that with more samples in our training data, the fairness metrics would not be reduced as substantially based on how the AI Fairness toolkit works. With a more significant variance in the sample resumes, as would be expected in a real-world scenario, we believe the toolkit would be far less practical.

### Microsoft Fairlearn Toolkit

We employed the Microsoft Fairlearn toolkit (Bird et al., 2020) to ensure demographic parity, which means that a protected class, such as age, gender, or race, should not influence a model's classification. In other words, all these attributes should receive a positive outcome at equal rates.

To achieve this, we utilized three methods from Fairlearn: Correlation Remover, Exponentiated Gradient, and Adversarial Fairness Classifier. The Correlation Remover is a pre-processing algorithm that eliminates correlation between sensitive and non-sensitive features through linear transformations, making it ideal for our task due to the simplicity of our features. The Exponentiated Gradient method is an in-processing approach that reduces unfair classification, as Agarwal et al. (2018) described. The Adversarial Fairness Classifier uses an optimization algorithm based on the work of Zhang et al. (2018) and trains a TensorFlow neural network classifier that minimizes training error while preventing an adversarial network from infer-

ring sensitive features. Table 4 shows the results obtained from applying these methods.

Method Applied	Change in Selection Percentage	Statistical Parity Difference	Disparate Impact	Theil Index
<i>None</i>	Y = 0.758, O = 0.525, $ \Delta  = \mathbf{0.000}$	0.308	1.44	0.33
<i>A. Correlation Remover (pre-processing)</i>	Y = 0.702, O = 0.581, $ \Delta  = \mathbf{0.056}$	0.171	1.26	0.26
<i>B. Exponentiated Gradient (in-processing)</i>	Y = 0.717, O = 0.566, $ \Delta  = \mathbf{0.041}$	0.207	1.33	0.30
<i>C. Adversarial Fairness Classifier (in-processing)</i>	Y = 0.696, O = 0.587, $ \Delta  = \mathbf{0.062}$	0.163	1.19	0.24
<i>Combination of methods A, B, and C</i>	Y = 0.671, O = 0.612, $ \Delta  = \mathbf{0.087}$	<i>0.101</i>	<i>1.12</i>	<i>0.21</i>

Table 4. Reduction in Age Bias by Microsoft Fairlearn

In Table 4, we can see that similar to the AI Fairness 360 toolkit, the combination of the three methods we applied using the Fairlearn toolkit improves fairness but does not entirely remove it. The results we obtained were an improvement but did not eliminate the biases as reflected in the three metrics we evaluated (statistical parity difference, disparate impact, and Theil index) as much as the AI Fairness 360 toolkit did.

Next, we looked at one of these metrics, statistical parity difference, as applied to each of the three protected attributes. This metric gives us a relative difference to see the amount of bias in our initial (untreated) model and how each toolkit can resolve these biases using its best model, which is a combination of methods explored in Table 3 and Table 4. This information is provided in Table 5.

Attribute	Statistical Parity Difference			
	Untreated	AI Fairness 360	MS Fairlearn	$\Delta$
<i>Age</i>	0.308	0.069	0.101	$ \Delta  = 0.239$
<i>Gender</i>	0.206	0.038	0.049	$ \Delta  = 0.168$
<i>Race</i>	0.153	0.018	0.026	$ \Delta  = 0.135$
<i>Average</i>	0.222	0.042	0.059	$ \Delta  = 0.181$

Table 5. Effectiveness of the Best AI Mitigation Techniques on Various Attributes

Table 5 shows that the combination of methods we applied from the AI Fairness 360 toolkit consistently reduced the biases for all three attributes better than the combination of techniques we used from Fairlearn; however, it did not eliminate them. Moreover, the applied methods were better able to reduce biases from race and gender than from age, demonstrating the challenge of age biases in human-trained data.

## Mitigating Age Biases

The previous section highlights the challenges of reducing age bias in job selection, even with advanced techniques. Possible solutions include changing the mindset of older applicants, altering how algorithms utilize information provided by older applicants, and motivating more senior candidates to stay current in rapidly evolving fields like IT and healthcare. Despite studies showing that older workers are perceived as less productive and perform less effectively than younger colleagues (Kudins, 2022), their presence can enhance productivity through mentorship. To recognize the advantages of older job seekers, job positions that benefit from mentor-mentee relationships should be rephrased. Including mentorship as a job search attribute can level the playing field and provide opportunities for qualified older workers who would relish the chance to serve as mentors.

According to other experts, limiting the job experience listed on a resume to only the past 15-20 years can be beneficial, especially if the applicant has had multiple job changes or roles throughout their career. The implementation of this limitation is not difficult from an algorithmic standpoint. While this algorithmic cap may obscure meaningful early career experiences, it may offer more advantages than disadvantages. Additionally, omitting the dates of university degrees can help shift attention away from using them to determine an applicant's age.

To improve their chances of being selected, job applicants in rapidly evolving industries like IT can take additional steps. They can start by highlighting their knowledge of the latest developments in their field and using relevant buzzwords and jargon. These are the criteria that AI filters often use to screen candidates. Focusing on specific achievements is also beneficial since most AI filters can recognize them. Lastly, older job applicants usually have more industry contacts and connections developed over time, which younger job applicants may be unable to match.

Some experts suggest that equal opportunities for older job applicants can be ensured by maintaining better statistics, which would only add complexity to an already complicated hiring process. Employers can always argue that they did not have complete information about the candidate pool when selecting candidates. However, in the US and many other job markets, the proportion of workers aged 50 or older is increasing and currently accounts for one-third of the workforce (Whalen, 2021). This demographic shift will necessitate many employers to reconsider their algorithmic hiring methods and approaches.

## Conclusion

This paper explores how the increasing use of online job searches has amplified age discrimination in job searching. Several studies have shown that the algorithms used by these searches tend to discriminate against older job seekers. While hiring biases based on race and gender have been extensively researched, age bias is more difficult to address because of the reliance on resume and CV scanning by algorithms prioritizing keywords over other valuable contributions of older job seekers.

We conducted a study to show how pervasive age bias is and how it can become part of an AI algorithm. In the survey, we asked 144 crowdworkers to determine if a candidate was someone they would bring in for an interview. Changing the age, race, and gender of the applicants, we found a statistically significant implicit bias against older applicants. This bias also became part of the AI algorithm developed using input from these crowdworkers.

While using the Microsoft Fairlearn and AI Fairness 360 toolkits to mitigate age bias in job selection yielded positive results, it should be noted that the process required significant effort and would not be practical for most companies in the real world. This additional required effort highlights the significant challenges that job searching and applicant screening algorithms face in minimizing implicit biases. For companies to hire the best workers, reducing these biases in the algorithms used for applicant screening is essential.

Various societal, technological, and educational approaches can be used to achieve this. Changing the perspective of older applicants is one such approach, as it can help to shift the perception that older workers are less productive or less valuable. Encouraging older applicants to stay current with new developments in their field, particularly in fast-changing industries like IT and healthcare, can also help showcase their continued relevance and expertise.

From a technological standpoint, screening algorithms can be adjusted to limit their focus on age-related features, such as the length of job history or dates of university degrees. This adjustment can help to ensure that valuable experiences early in an applicant's career are not overlooked. Additionally, providing clear guidelines for appropriate terminology and buzzwords can help older applicants tailor their resumes to properly match the screening algorithms' expectations.

While these approaches can help overcome some implicit biases that older applicants face in most industrial societies, it is essential to recognize that they may not be a perfect solution. Implicit biases can be deeply ingrained and difficult to eradicate, and ongoing efforts will be needed to ensure that older workers are not unfairly disadvantaged in the job market.

## References

- Ackerman, L. S., & Chopik, W. J. (2021). Cross-cultural comparisons in implicit and explicit age bias. *Personality and Social Psychology Bulletin*, 47(6), 953-968.
- Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna M. Wallach. (2018) A reductions approach to fair classification. In ICML, volume 80 of Proceedings of Machine Learning Research, 60–69. PMLR.
- Bellamy, Rachel KE, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, et al. (2019) "AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias." *IBM Journal of Research and Development* 63, no. 4/5: 4-1.
- Bird, S., Dudík, M., Edgar, R., Horn, B., Lutz, R., Milan, V., ... & Walker, K. (2020). Fairlearn: A toolkit for assessing and improving fairness in AI. Microsoft, Tech. Rep. MSR-TR-2020-32.
- Bogen, M., & Rieke, A. (2018). Help wanted: An examination of hiring algorithms, equity, and bias.
- Bovenkerk, F. 1992. Testing Discrimination in Natural Experiments: a manual for international comparative research on discrimination on the grounds of 'race' and ethnic origin. Geneva: ILO.
- Büsch, V. and Dittrich, Dennis Alexis Valin and Königstein, Manfred, Bad Hiring Chances for Older Workers (2009). Available at SSRN: <https://ssrn.com/abstract=1434351> or <http://dx.doi.org/10.2139/ssrn.1434351>
- Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183-186.
- Carlsson, M., & Eriksson, S. (2019). Age discrimination in hiring decisions: Evidence from a field experiment in the labor market. *Labour Economics*, 59, 173-183.
- Clark, C. (2022) "How Many Interviews For One Job? (Stats From a Recruiter)" <https://careersidekick.com/interviews-per-job/>
- Dastin, J. (2018). Amazon scraps secret AI recruiting tool that showed bias against women, in *Ethics of Data and Analytics* pp. 296-299. Auerbach Publications.
- Daumeyer, N. M., Onyeador, I. N., Brown, X., & Richeson, J. A. (2019). Consequences of attributing discrimination to implicit vs. explicit bias. *Journal of Experimental Social Psychology*, (84), 103812.
- Derous, E., & Decoster, J. (2017). Implicit age cues in resumes: Subtle effects on hiring discrimination. *Frontiers in Psychology*, 8, 1321.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. In Proceedings of the 3rd Innovations in theoretical computer science conference (pp. 214-226).
- Faliagka, E., Ramantas, K., Tsakalidis, A., & Tzimas, G. (2012). Application of machine learning algorithms to an online recruitment system. In Proc. International Conference on Internet and Web Applications and Services (pp. 215-220).
- Harris, C.G. (2020). Mitigating Cognitive Biases in Machine Learning Algorithms for Decision Making. In Companion Proceedings of the Web Conference pp. 775-781
- Harris, C.G. (2022) Age Bias: A Tremendous Challenge for AI Algorithms in the Job Candidate Screening Process. In IEEE SSIT international symposium on technology and society 2022 (ISTAS 2022). IEEE
- Hofstede, G., Hofstede, G. J., & Minkov, M. (2005). *Cultures and organizations: Software of the mind* (Vol. 2). New York: McGraw-hill.
- Jacobs, A. Z. and Hanna Wallach. Measurement and fairness. In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21, 375–385. New York, NY, USA, 2021. Association for Computing Machinery. URL: <https://doi.org/10.1145/3442188.3445901>, doi:10.1145/3442188.3445901.
- Karaoglu, G., Hargittai, E., & Nguyen, M. H. (2021). Inequality in online job searching in the age of social media. *Information, Communication & Society*, 1-19.
- Karpinska K., Henkens K., Schippers J. (2013). Hiring retirees: impact of age norms and stereotypes. *Journal of Managerial Psychology*, (28), 886–906.
- Kaufmann, M. C., Krings, F., Zebrowitz, L. A., & Sczesny, S. (2017). Age bias in selection decisions: The role of facial appearance and fitness impressions. *Frontiers in psychology*, 8, 2065.
- Köchling, A., & Wehner, M. C. (2020). Discriminated by an algorithm: A systematic review of discrimination and fairness by algorithmic decision-making in the context of HR recruitment and HR development. *Business Research*, 13(3), 795-848.
- Kolmar, C (2022) "15 incredible job search statistics: what job seekers need to know". <https://www.zipppia.com/advice/job-search-statistics/>
- Kudins, J. (2022). Economic usefulness of older workers in terms of productivity in the modern world. *Insights into Regional Development*, 4(1), 40-56.
- Lambrecht, A., & Tucker, C. (2019). Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Management Science*, 65(7), 2966-2981.
- Manyika, J., J. Silberg, & B. Presten (2019) "What Do We Do About the Biases in AI?" *Harvard Business Review*. October 25, 2019.
- Marino, J. (2022) "Must-Know Job Website Statistics (And How to Leverage Them)." <https://www.huemanrpo.com/resources/blog/must-know-job-website-statistics>
- Martin G. et al. (2014). In the name of meritocracy: managers' perceptions of policies and practices for training older workers. *Ageing and Society*, (34), 992–1018.
- McIntosh, J., Du, X., Wu, Z., Truong, G., Ly, Q., How, R., ... & Kanij, T. (2021). Evaluating Age Bias. In E-commerce. In 2021 IEEE/ACM 13th International Workshop on Cooperative and Human Aspects of Software Engineering (CHASE) (pp. 31-40). IEEE.
- Neumark, D., Burn, I., & Button, P. (2019). Is it harder for older workers to find jobs? New and improved evidence from a field experiment. *Journal of Political Economy*, 127(2), 922-970.
- Oude Mulders J. et al. (2014). How likely are employers to re-hire older workers after mandatory retirement? A vignette study among managers. *De Economist*, (162), 415–431.
- Raub, M. (2018). Bots, bias and big data: artificial intelligence, algorithmic bias and disparate impact liability in hiring practices. *Ark. L. Rev.*, (71) 529.

Rich, J. (2014) What Do Field Experiments of Discrimination in Markets Tell Us? A Meta-Analysis of Studies Conducted since 2000, IZA Discussion Papers, No. 8584, Institute for the Study of Labor (IZA), Bonn

Richardson, B., Webb, J., Webber, L., & Smith, K. (2013). Age discrimination in the evaluation of job applicants. *Journal of Applied Social Psychology*, 43(1), 35-44.

Rivera, L.A. (2015). "Go with Your Gut: Emotions and Evaluation in Job Interviews." *American Journal of Sociology* 120(5):1339-1389

Schwartz, H.A., Eichstaedt, J.C., Kern, M.L., Dziurzynski, L., Ramones, S.M., Agrawal, M., Shah, A., Kosinski, M., Stillwell, D., Seligman, M.E. and Ungar, L.H.(2013). Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one*, 8(9), p.e73791.

Silberg, J., & Manyika, J. (2019). Notes from the AI frontier: Tackling bias in AI (and in humans). McKinsey Global Institute, 1-6.

Soleimani, M., Intezari, A., Taskin, N., & Pauleen, D. (2021). Cognitive biases in developing biased Artificial Intelligence recruitment system. In *Proceedings of the 54th Hawaii International Conference on System Sciences*. p. 5091.

Stuart, R., & Norvig, P. (2016). *Artificial Intelligence: A Modern Approach*. Third Edition. Prentice Hall Press Upper Saddle River.

Upadhyay A, K. & Khandelwal., K. (2018) Applying artificial intelligence: implications for recruitment, *Strategic HR Review*, 17(5), 255-258

Vrontis, D., M. Christofi, V. Pereira, S. Tarba, A. Makrides & E. Trichina. (2021). Artificial intelligence, robotics, advanced technologies, and human resource management: a systematic review, *The International Journal of Human Resource Management*, DOI: 10.1080/09585192.2020.1871398

Whalen, J.R. (2021) How older workers can combat ageism while job hunting. <https://www.wsj.com/podcasts/your-money-matters/how-older-workers-can-combat-ageism-while-job-hunting>

Wiens, J., Price, W. N., & Sjoding, M. W. (2020). Diagnosing bias in data-driven algorithms for healthcare. *Nature medicine*, 26(1), 25-26.

Zacher, H., & Steinvik, H. R. (2015). Workplace age discrimination. *The encyclopedia of adulthood and aging*, 1-5.

Zhang, B. H. Blake Lemoine, and Margaret Mitchell. (2018) Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 335–340. 2018.

Zhavoronkov, A., Mamoshina, P., Vanhaelen, Q., Scheibye-Knudsen, M., Moskalev, A., & Aliper, A. (2019). Artificial intelligence for aging and longevity research: Recent advances and perspectives. *Ageing research reviews*, 49, 49-66.