I Whitepaper

Removing Bias from Talent Decisions with Artificial Intelligence.



Artificial Intelligence (AI) is the simulation of human processes by machines. Like most new technologies, AI is responsible for both good and bad outcomes.

However, the negative impacts have received disproportionate attention, particularly around the potential for AI to perpetuate bias. This concern is both well-founded and well-documented.

This fear of biased AI rests largely upon the fact that the foundational source of bias in AI is the

human behavior it is simulating, or the biased dataset used to train the algorithm. Al learns from humans, so if you don't like what the Al is doing, then you certainly won't like what humans are doing.



In the case of hiring and talent mobility specifically, the status quo is deeply flawed for four main reasons:



Unconscious human bias leads to unfair hiring decisions.



Unconscious human bias leads to unfair internal mobility decisions.



Significant pools of talent are being ignored.



Traditional talent assessment tools are already biased.

1

Unconscious human bias leads to unfair hiring decisions.

Applicants' resumes are typically reviewed before an in-person or digital interview by recruiters. Countless studies have shown that this process is wrought with significant unconscious bias against women, minorities and older workers.

In fact, a PNAS paper that aggregated 30 years of resume studies found that **white applicants continue to receive 36% more callbacks for interviews than Black applicants, and 24% more than Hispanics**. This indicates that there has largely been no progress on racial discrimination in hiring over the last several decades.



Unconscious human bias leads to unfair internal mobility decisions.

73% of senior executives are white. The remainder are 21% Asian, and just 3% Latino/a, and 2% Black. Relative to white men, before receiving a promotion, Black men must work longer periods of time after leaving school and Latinos must accrue significantly more years with their current employer. Additionally, relative to white men, Black women and Latinas must have more prior job-specific experience and more overall work experience before receiving a promotion — all else equal.

One study even shows that if performance evaluations are biased such that women and minorities receive lower ratings than they deserve, controlling for those ratings still caused an understatement of the amount of promotion discrimination experienced in this organization.

65% of Black professionals acknowledge these difficulties to advancing, but only 16% of white professionals agree with that statement, unaware of the vast differences and discrimination in the workplace for people of color.

Significant pools of talent are being ignored.

On average, any open role receives 250 applications, due in large part to the success of LinkedIn and other online platforms that support sweeping hiring campaigns. This translates into millions of applicants for a few thousand open roles, which makes it impossible for this process to be managed manually.

So, recruiters focus on "the top" 10% - 20% they presume show the most promise: talent coming from lvy League campuses, passive candidates from competitors seeking to fill positions, or employee-referral programs. However, not surprisingly, top colleges and employee-referral programs are far less diverse than the broader pool of applicants who are submitting resumes.

Organizations tend to ignore talent internally as well, since decisions are largely biased towards talent that has previously worked in a similar role and is looking to advance within the same department. We as human beings have trouble seeing candidate potential beyond work experience or hard skills, when in fact there is a plethora of well-fit talent internally from other areas of the business.



Traditional talent assessment tools are already biased.

The most fundamental component of hiring that induces bias is the information used to make hiring decisions in the first place. Resumes and other cognitive-ability or personality tests are entirely self-reported, meaning applicants can present themselves or answer questions in a way that is inaccurate and misleading.

A federal regulation loophole states that a hiring tool or methodology can be biased so long as it is "job-related". This means that if people who are successful in a role show particular characteristics, and all "successful employees" are white (for example) due to the history of biased human hiring and promotion practices, then it is almost guaranteed that your job-related talent assessment tools will favor white applicants and disadvantage minorities. A Black female from a non-lvy League university who is fortunate enough to make it into the talent pipeline, and whose resume is reviewed and passes the human recruiter evaluation, may then be asked to take a biased assessment...after all that!

It's no wonder then that we struggle to hire and promote a diverse workforce. Today's chronic lack of diversity is due to human paradigms in place today, not AI. Without a different approach, these biases will only continue, and even worsen.

AI holds the greatest promise for eliminating bias in hiring for two main reasons:



Unlike humans, Al can be de-biased.

2

Al can assess the entire pipelineof talent, unlike time-constrained humans, who tend to implement processes that only shrink the pipeline and narrow definitions of potential.

Unlike humans, AI can be de-biased.

Some AI tools used in the talent space have certainly been shown to be flawed. However, unlike their biased human counterparts, at least the AI can be corrected — the beauty of AI is that we can design it to meet certain specifications. AI practitioners such as OpenAI and the Future of Life Institute are already compiling a set of design principles for making AI ethical and fair.

One key principle in this movement is that AI must be designed to be audited, and if there are any biases found they must be removed. An AI audit should function not unlike the safety tests of a new car before someone drives it. If standards are not satisfactorily met, the defective technology must be corrected before it is allowed into production.

At pymetrics, we proactively debias each of our algorithms to ensure we are compliant with EEOC regulations around fairness optimization. In addition, once we deploy a model into the world, we conduct frequent checks to ensure we are driving fairness when real applicants are being assessed as the applicant pool can highly impact hiring flow as well as distribution of demographics.



Al can assess the entire pipelineof talent, unlike time-constrained humans, who tend to implement processes that only shrink the pipeline and narrow definitions of potential.

Companies today admit that just a small portion of the often thousands of applicants are ever reviewed. As a result, recruiters who are time constrained only look to the top schools to find talent, further worsening job prospects for those who are ethnically and socioeconomically-diverse. Technologists and lawmakers alike ought to work together to create tools and policies that enable and even mandate that the entire pipeline be reviewed. Furthermore, technologies can enable us to assess skills that are less tangible but still incredibly job-relevant. Human beings can only make decisions based on hard skills ad past performance — and even then we tend to be poor judges of fit.

Without relying on technology, we will be missing out on strong-performing talent based on more innate aptitudes and potential.

Using an automated top-of-funnel process can significantly reduce the bias induced by shrinking the initial talent pipeline. By having every applicant play the same set of neuroscience based games, pymetrics is able to assess people at scale and determine their match to a job within seconds.

Most importantly, every applicant actually gets considered and assessed for the job, and receives a personalized report outlining their behavioral strengths and opportunities that they can leverage in future interviews. In addition, our matching algorithms enable individuals to not only be assessed for the role they applied to, but for any other role that uses pymetrics. At scale, this fundamentally starts to change the way millions of people are efficiently and accurately assessed for jobs — in fact, it already has!

Fortunately, some precedents for such standards have already been set. The California State Assembly passed a resolution to leverage unbiased technology to promote diversity in hiring, and the San Francisco District Attorney is using "blind sentencing" Al in criminal justice proceedings. Why shouldn't the same standards be applied to existing talent assessment tools? Amazon was nationally lambasted for months for its male-biased hiring algorithm. Yet as it stands in the United States today, employers are legally allowed to use traditional and biased assessments that discriminate against women and/or minorities. Most people are unaware that biased assessments are legal and prominently used in this way.

If we are going to call for unbiased AI — which we undoubtedly should — we ought to call for the elimination of all biased traditional assessments for good.

By having every applicant play the same set of neuroscience based games, pymetrics is able to assess people at scale and determine their match to a job within seconds.

Key considerations when assessing Al-driven talent technologies.



pymetrics.ai

Fair datasets.

Whenever people data is involved, bias is involved. Structural injustices and unbalanced representation of the population can be reflected in your dataset, causing discrimination against certain protected classes. Because such people data has sociological meaning that a technologist might not be as well versed in uncovering, consider working with an expert in this field (such as an IO psychologist) that can help determine how to best proceed based on your goals. They can help you assess whether the data you're using is truly fair, as oftentimes seemingly unrelated data points can serve as proxies for determining other factors such as gender or race. For example, the quality of an applicants' high school may seem like an innocuous data point, but school districts are often tightly interwoven with historic patterns of bias.

To address this issue, pymetrics uses objective, behavioral data to drive hiring decisions based on an applicant's underlying behavioral tendencies.

Through this approach, the data we feed into our algorithms is already more fair to begin with, unlike self-reported surveys that are often skewed to present the persona that the candidate thinks will best impress a recruiter. The specific behavioral measures used by pymetrics are vetted for use across cultures and geographies, and the models used to predict job success are always built and tested on a diverse dataset.

Audited algorithms.

If there has ever been a time to think about how Al can help to equalize employment opportunities in the ensuing economic recovery, it is now.

We can't let careless AI deployment deepen the disparity. In the US, the rule for fair hiring, legally defined as the 4/5ths rule, holds that relative to the highest passing group, other groups should be passing at the 80% rate. Neither of the two most common employment selection tools in use today — cognitive tests and resume reviews — meet this benchmark. The former yields just 32 applicants of color for every 100 white applicants selected, and the latter yields 74 applicants of color for every 100 white applicants. Their effectiveness at predicting job fit has been in question too. **Audited Al tools outperform more traditional tools in every aspect.** Always create an audit trail you can map back to and be able to explain to yourself, your team, and the end user why a decision has been made.

At pymetrics, not only do we comply with the 4/5ths rule to ensure equal employment for all gender and ethnic groups, but we are paving the way for new regulations to address algorithmic use in hiring.

We have been actively involved in the proposal of a New York City Council bill that would require firms to disclose when they use software to assess talent, and vendors would have to ensure that their tech doesn't discriminate. Additionally, we recently had our platform audited for fairness by Northeastern University, the results of which are soon to be published in the ACM Conference on Fairness, Accountability and Transparency (FAccT), leading academic source for algorithmic fairness.

Mutual transparency.

It's essential to be mindful of whether the data you're using throughout the recruiting process is appropriate, job relevant, and even legal to be considering as an employer. When partnering with a vendor, transparency about these details on both sides is key. The more AI is used in hiring decisions, the more important it is to be able to explain what those algorithms are actually doing. Be clear on what you're trying to accomplish and what has and biased patterns in your team's decision-making. Ensure that the vendor's technology is relevant to your goals and your specific organization — there are lots of shiny objects out there that make it difficult to differentiate between those with scientific validity from those that do not. Do your research, ask lots of questions, and be upfront about what you're looking to accomplish and why.

pymetrics was founded on the belief that providing a transparent hiring process and explaining the decisions our technology makes is always the gold standard.

Starting with a job analysis, pymetrics assesses what knowledge, skills, and abilities are important to succeed in a role. This gives us an objective view of what matters, and how our behavioral attributes map to these competencies. We then report on these attributes to recruiters and candidates, using it to explain a candidate's fit-to-role for the former, and to provide an overview of strengths and opportunities at work for the latter.

In summary.

While it is impossible to correct human bias, it is demonstrably possible to identify and correct bias in Al. If we take the critical steps necessary to address the concerns being raised, and do so in a way that is fair, actively audited, and transparent, we can successfully harness the power of technology to diversify the workplace going forward.

To find out more about pymetrics and our capabilities, please connect with our team.



pymetrics.ai