Tyler Schoulte
Experiential Ethics
Summer 2020

Tell Me About Yourself: A.I. Edition

The world is constantly changing, to become more efficient. Technology has been creeping its way into every aspect of life, from smart homes to self-driving cars. But one aspect, in particular, has been a growing industry: using artificial intelligence and machine learning to hire job applicants. There are several ways in which this technology can be used in the process of hiring an applicant. However, this essay will focus on an emerging new use case: replacing the human interviewer with a program that asks all of the questions, records all of the answers, and completes all of the analysis. In this essay, I will discuss how such a technology is currently being used, touch on how it could be used in the future, but most importantly, address the question: should this technology be used in the first place? Are the limitations and drawbacks of using such a system worth its merits? These discussions will hopefully lead you to the same conclusion I argue; that these systems are currently humanity's best chance at eliminating hiring bias.

It's worth noting that this paper will use artificial intelligence (AI) and machine learning (ML) interchangeably. While these technologies are subtly different, the abstraction exists as a means to refer to any system that is judging human applicants with computer software. Before discussing whether this artificially intelligent system *should* exist, it's best to discuss *why* it already exists in the first place.

Monetarily it is logical to replace recruiters with computers, so long as the cost of paying for such software-as-a-service is cheaper than the ousted recruiters. But there are deeper reasons as well. Some companies need to filter through an extensive number of applicants. While this sounds tedious and time-consuming, it is actually in the company's best interest to have more applicants who apply to be filtered through. The larger the applicant pool, the better the talent a company can pick from. As an example, Former Google HR chief Laszlo Bock says Google receives over 50,000 resumes a week.[1] Without some sort of computerized method of filtering, Google would struggle to filter these applicants without driving their recruiters insane reading resumes all day, every day. For this reason, artificially intelligent systems already scan and read resumes to filter applicants. However, this comes with caveats that I will return to later.

Outside of resume-readers, other systems are also already commonplace in the industry. As briefly mentioned, there are already several types of AI on the market that are monitoring job applicants. LinkedIn has a feature called LinkedIn Recruiter, which uses machine learning to determine which LinkedIn accounts to suggest to recruiters. Every day, depending on what you feature on your profile, what you post, and who you follow or respond to on LinkedIn, the platform is deciding whether to recommend you or millions of others to recruiters' feeds.[2] Your behavior on a 'social media site' determines whether a recruiter might reach out to you.

Finally, there is also a popular challenge, typically used amongst computer science majors, and that challenge is a coding challenge. While not typically involving machine learning for grading, HackerRank coding challenges or eSkill assessments are meant as a way to allow many candidates to apply to a role and take a competency test that is scored by the computer. If

[1] Zetlin, M. (2018, April 09). Google Automatically Rejects Most Resumes for Common Mistakes You've Probably Made Too. Retrieved August 20, 2020, from https://www.inc.com/minda-zetlin/google-resume-mistakes-laszlo-bock-job-hiring-employment.html

[2] Guo, Q. (2020, April 17). The AI Behind LinkedIn Recruiter search and recommendation systems. Retrieved August 20, 2020, from https://engineering.linkedin.com/blog/2019/04/ai-behind-linkedin-recruiter-search-and-recommendation-systems

the grading algorithm determines the candidate has the basic skills required for the job, they can proceed in the recruitment process.

What's important to realize is that companies are already using algorithms and machine learning in their recruitment processes. They use machine learning to grade candidates, find the candidates, read the resumes of candidates. But the most egregious part is they are using these simple methods like resume screening to eliminate tons of people before they even are given the chance for an interview. Perhaps for people who miss a spelling error, their resume is thrown out in the hyper-competitive recruiting industry. For disproportionate minorities who do not have resources or peers to do in-depth resume reviews, this leaves them stonewalled, never seeing the day for an interview. Companies like Google have to filter 50,000 candidates per week down to less than 100 (Google has a 0.2% hiring rate, significantly lower than that of, say, Harvard, which has a 4.7% acceptance rate[3]). An absurdly high percentage of candidates are turned away after a quick resume scan, never having a chance to tell their story or argue their merit. But it doesn't have to be that way. There is another way.

The next evolutionary step for recruiting was for machine learning to administer the interview. Currently, HireVue is the most popular artificial interviewing platform, and by a large margin. HireVue has 700 reputable corporate partners, including the likes Carnival Cruise Line, Goldman Sachs, Staples, and Under Armor.[4] These diverse partners trust HireVue to manage their recruiting by filtering candidates that match their desired preferences for candidates and

---

[3] Nisen, M. (2014, October 23). Here's why you only have a 0.2% chance of getting hired at Google. Retrieved August 17, 2020, from https://qz.com/285001/heres-why-you-only-have-a-0-2-chance-of-getting-hired-at-google

[4] Customers: HireVue Customer Stories. (2020, April 01). Retrieved August 16, 2020, from https://www.hirevue.com/customers

skillsets. This process saves money, saves recruiters time, and ultimately makes it possible for HireVue to interview every candidate who applies. No more having an applicant's worth being reduced to a single-sided piece of paper! A spelling mistake isn't the end of the road! The possibilities are endless….

Slowing down a little, there are admittedly great challenges in building an unbiased artificially intelligent system to recruit. HireVue has a great product in terms of marketing stats: the company makes the claim "We help our clients lead their industries with 50% faster growth, 29% less turnover, and 13% more top performers." Additionally, Forbes has ranked HireVue as one of the "Top Ten Most Promising Companies".[5] Marketing hype aside, the biggest ethical question is whether the product is designed to limit politics. In all honesty, achieving a completely unbiased system is, in my view, impossible! But that doesn't mean that I believe it shouldn't be created or implemented.

To justify my reasoning for implementing an imperfect system, we need to analyze further the current system. While there are great challenges when creating an ML computer interviewing system, there are also great challenges when recruiting using human interviewers. Most of the mistakes humans make in judgment, computers can deliberately be trained to avoid with ease, or not need to be trained to avoid at all because, unlike other humans, robots are not

[5] About HireVue: What We Do and How We Got Here. (2020, March 05). Retrieved August 20, 2020, from https://www.hirevue.com/company/about-us

attracted to looks or charm.[6] These biases subconsciously cloud interviewer judgement frequently, with the seven main types of human interview biases including:[7]

1. **Cultural Noise:** Interviewee responds to cater to the interviewer's interpreted preferences, rather than the truth.
2. **Stereotyping Bias:** Stereotyping based on the member belonging to a certain group, such as men vs. women or white vs. Hispanic.
3. **Generalization Bias:** The interviewer assumes a one-off mannerism is a reoccurring behavior, i.e., lack of confidence or saying 'um'.
4. **Halo/Horn Bias:**
   a. **Halo:** Single strong point of interviewee overshadows failures
   b. **Horn:** Single failure of interviewee overshadows successes
5. **Recency Bias:** The interviewer more clearly recalls recently interviewed candidates over earlier candidates.
6. **Contrast Effect:** Previous candidates being stronger or weaker than the current interviewee can make the current interviewee perceivably stronger or weaker *in contrast*, even if they are underqualified or qualified respectively.
7. **Gender and Racial Bias:** Stereotyping, but due to gender or racial reason specifically.

---

[6] Silberg, J., & Manyika, J. (2020, July 22). Tackling bias in artificial intelligence (and in humans). Retrieved August 15, 2020, from https://www.mckinsey.com/featured-insights/artificial-intelligence/tackling-bias-in-artificial-intelligence-and-in-humans

[7] Gomes, O. (2019, July 10). What is Interview Bias & How to Avoid it While Hiring. Retrieved August 18, 2020, from https://blog.talview.com/what-is-interview-bias-how-to-avoid-it-while-hiring

These biases are not just philosophy on the human mind, either. A management study done by Stephen Cohen at Princeton University showed how human employers can stereotype and believe that someone is fit for a role just based on gender norm bias.[8] In the paper, Cohen experimented with showing recruiters different resumes which were nearly identical, except for a male vs. female name. For roles thought to be traditionally female, such as assistant, there was a 28% drop in the acceptable males vs. females. However, for the traditionally male role of technician, 86% percent of males were qualified against 53% of qualified females. Again, these resumes were nearly identical in qualifications, but the subconscious interview biases skewed the results. In the real world, and not just a simulation, this leads to cyclical loops of gender gaps for certain fields. When that gap is created, fewer of the minority gender are likely to pursue an education in that field, continuing the cycle, as well as the bias.

TABLE 3

Percent of Applicants Rated as Acceptable or Unacceptable for Hiring*

Job position:	**Editorial Assistant**		(Female-oriented)	
Recruiter's decision:	Acceptable		Unacceptable	
Applicant's sex:	Male	Female	Male	Female
Percentage:	58%	86%	42%	14%
Job position:	**Personnel Technician**		(Male-oriented)	
Recruiter's decision:	Acceptable		Unacceptable	
Applicant's sex:	Male	Female	Male	Female
Percentage:	86%	53%	14%	47%

Many of these biases would not exist with artificially intelligent systems. These systems can be trained on equal amounts of female and male data to prevent issues like this from occurring. A computer doesn't care if you are male or female, only if you sound knowledgeable in the field. Companies like HireVue have released software that has been used for over 10 million interviews, with over 1 million more interviews being conducted every 90 days in over

[8] Cohen, S. L. (1976). The basis of sex-bias in the job recruitment situation. *Human Resource Management, 15*(3), 8-10. doi:10.1002/hrm.3930150303

180 countries.[9] We sit now at a paradigm shift, where computer science and automation once again take over another job; the interviewer.

There are concerns that a hiring system such as HireVue isn't fully transparent into how the algorithms work, and thus, have doubts about whether a system might potentially favor one race over the other, a deep voice over a soft voice, etc.[10] To address these concerns, the first step to moving forward is to address the root cause of the problem. In a 2020 paper titled "A Framework for Understanding Unintended Consequences of Machine Learning", MIT researchers Suresh and Guttag concluded that "Commonly used terms such as 'training data bias' can be too broad to be useful, and context-specific fixes don't have the shared terminology to generalize and communicate the problem to a wider audience." [11]. In the paper, the researchers determined six different biases that machine learning models can have, and by using their guidelines, they hope that developers of systems can more accurately address issues by more quickly determining the actual cause of the fault. What Suresh and Guttag ultimately promote is a roadmap to assure quality and rectify issues directly, rather than grouping it all into a "training data problem". Using these roadmaps, developers can more effectively rectify issues their models have.

However, first biases must be known to follow a roadmap to fix them. That is why another group of researchers from Adobe and Georgia Tech came to similar conclusions in 2020

---

[9] HireVue Surpasses Ten Million Video Interviews Completed Worldwide. (2020, May 19). Retrieved August 18, 2020, from https://www.hirevue.com/press-release/hirevue-surpasses-ten-million-video-interviews-completed-worldwide

[10] Burke, L. (2019, November 4). As AI-assessed job interviewing grows, colleges try to prepare students. Retrieved August 18, 2020, from https://www.insidehighered.com/news/2019/11/04/ai-assessed-job-interviewing-grows-colleges-try-prepare-students

[11] Suresh, H., & Guttag, J. V. (2019). A framework for understanding unintended consequences of machine learning. arXiv preprint arXiv:1901.10002.

with their paper titled "Designing Tools for Semi-Automated Detection of Machine Learning Biases: An Interview Study".[12] Failures in machine learning were not being categorized correctly, so it took researchers longer to address the biases. Rather than examining the inputs of the black box and trying to resolve the issue by changing the inputs, semi-automated tools needed to be created to evaluate the data. This tool would be tasked with trying to group the data in a fashion that would determine if the training data was biased and if so, remark how such training data could be altered or aided to fix the model.

One machine learning scholar interviewed in the paper remarked "just imagine we have a hunch that something is wrong, it is really hard just to go and pinpoint those [biases] one by one." Instead, having a semi-automated system aware of a very large collection of biases for disabilities, speech impediments, skin color, gender. This means that every iteration of the software doesn't have to have the developers following their hunches, but can instead rely on pre-labeled data and results that there are no biases being output. The makers of this large collection of biases is left as an experiment to the reader, but that is where I believe ethicists could join forces with the engineers. Ethicists have been studying biases for quite some time now, as evidenced by several types that have been stated in this paper. If these two parties teamed up, they could ultimately work together to create the product, as well as create the auditing system to ensure the product is producing tolerable results.

Ultimately, artificially intelligent recruiting systems require work still. The systems will never be perfect in my opinion. But it doesn't have to be perfect. As a society, we are faced with a trolley problem. Currently, bias is happening every day in the hiring world. There were almost

[12] Law, P. M., Malik, S., Du, F., & Sinha, M. (2020). Designing Tools for Semi-Automated Detection of Machine Learning Biases: An Interview Study. arXiv preprint arXiv:2003.07680.

24 million managers in 2014 according to the Bureau of Labor Statistics.[13] To believe that 24 million hiring managers from a wide variety of locales and backgrounds would be unbiased is a pipedream. It might be possible to attempt to make them aware of their subconscious biases. But perhaps it would be easier to group 100 engineers and 100 ethicists to build a system that, while not 100% perfect, had so much thought put into it that it is less biased than the 24 million humans. But perhaps that's where the trolley problem comes in. If we accept people are going to be run over by the trolley, do we flip the switch to reduce bias to give control to the computer? Or do we let the trolley keep going, potentially harming more people, but we know that it was due to direct human fault? The technology might not be value-neutral, but neither are humans.

So, tell me about yourself.

[13] Zanini, G. (2016, September 06). Excess Management Is Costing the U.S. $3 Trillion Per Year. Retrieved August 20, 2020, from https://hbr.org/2016/09/excess-management-is-costing-the-us-3-trillion-per-year

Bibliography

About HireVue: What We Do and How We Got Here. (2020, March 05). Retrieved August 20, 2020, from https://www.hirevue.com/company/about-us

Burke, L. (2019, November 4). As AI-assessed job interviewing grows, colleges try to prepare students. Retrieved August 18, 2020, from https://www.insidehighered.com/news/2019/11/04/ai-assessed-job-interviewing-grows-colleges-try-prepare-students

Cohen, S. L. (1976). The basis of sex-bias in the job recruitment situation. *Human Resource Management, 15*(3), 8-10. doi:10.1002/hrm.3930150303

Customers: HireVue Customer Stories. (2020, April 01). Retrieved August 16, 2020, from https://www.hirevue.com/customers

Gomes, O. (2019, July 10). What is Interview Bias & How to Avoid it While Hiring. Retrieved August 18, 2020, from https://blog.talview.com/what-is-interview-bias-how-to-avoid-it-while-hiring

Guo, Q. (2020, April 17). The AI Behind LinkedIn Recruiter search and recommendation systems. Retrieved August 20, 2020, from https://engineering.linkedin.com/blog/2019/04/ai-behind-linkedin-recruiter-search-and-recommendation-systems

HireVue Surpasses Ten Million Video Interviews Completed Worldwide. (2020, May 19). Retrieved August 18, 2020, from https://www.hirevue.com/press-release/hirevue-surpasses-ten-million-video-interviews-completed-worldwide

Law, P. M., Malik, S., Du, F., & Sinha, M. (2020). Designing Tools for Semi-Automated Detection of Machine Learning Biases: An Interview Study. arXiv preprint arXiv:2003.07680.

Nisen, M. (2014, October 23). Here's why you only have a 0.2% chance of getting hired at Google. Retrieved August 17, 2020, from https://qz.com/285001/heres-why-you-only-have-a-0-2-chance-of-getting-hired-at-google

Silberg, J., & Manyika, J. (2020, July 22). Tackling bias in artificial intelligence (and in humans). Retrieved August 15, 2020, from https://www.mckinsey.com/featured-insights/artificial-intelligence/tackling-bias-in-artificial-intelligence-and-in-humans

Suresh, H., & Guttag, J. V. (2019). A framework for understanding unintended consequences of machine learning. arXiv preprint arXiv:1901.10002.

Zanini, G. (2016, September 06). Excess Management Is Costing the U.S. $3 Trillion Per Year. Retrieved August 20, 2020, from https://hbr.org/2016/09/excess-management-is-costing-the-us-3-trillion-per-year

Zetlin, M. (2018, April 09). Google Automatically Rejects Most Resumes for Common Mistakes You've Probably Made Too. Retrieved August 20, 2020, from https://www.inc.com/minda-zetlin/google-resume-mistakes-laszlo-bock-job-hiring-employment.html