

A Silicon Valley Love Triangle: Hiring Algorithms, Pseudo-Science, and the Quest for Auditability

MONA SLOANE

New York University, ms11521@nyu.edu

EMANUEL MOSS

CUNY Graduate Center, Data & Society Research Institute, emanuel@datasociety.net

RUMMAN CHOWDHURY

Director of ML Ethics, Transparency and Accountability (META) at Twitter, rchowdhury@gmail.com

In this paper, we suggest a systematic approach for developing socio-technical assessment for hiring ADS. We suggest using a matrix to expose underlying assumptions rooted in pseudoscientific essentialized understandings of human nature and capability, and to critically investigate emerging auditing standards and practices that fail to address these assumptions.

CCS CONCEPTS • Social and Professional Topics

Additional Keywords and Phrases: Audit, Accountability, Hiring, Automated Decision Making Systems

ACM Reference Format:

Mona Sloane, Emanuel Moss, Rumman Chowdhury. 2021. A Silicon Valley Love Triangle: Hiring Algorithms, Pseudo-Science, and the Quest for Auditability. In *CHI '21: ACM CHI Virtual Conference on Human Factors in Computing Systems*, May 8-13, 2021, Yokohama, JN. ACM, New York, NY, USA.

1. Introduction

Prospective job applicants find their interactions with future workplaces increasingly algorithmically mediated through automated decision making systems (ADS) (Sanchez-Montero et al. 2020, Ajunwa 2021). These systems are typically NLP, computer vision systems or gamified assessments that claim to predict job performance based on intonation, written text, micro-expressions, or game performance, often grounded in pseudo-scientific practices (Aguera y Aracas et al. 2017, Stark and Hoey 2020) such as phrenology and physiognomy (Sloane, 2021).

Even though there is mounting evidence that such systems harbor bias across demographic categories, algorithmic and bureaucratic opacity (Pasquale 2015, Burrell 2016, Levendowski 2019) have led to slow responses from regulators. Indeed, many of these systems promise more equitable outcomes than ostensibly more biased human decision-makers. While recent work has examined how hiring managers operationalize concepts like "fairness" through their interactions with these algorithmic systems (van den Broek et al. 2019), more work is needed to trace how the claims that these systems can reveal the true potential of job candidates are shaping the hiring ecosystem, and how these claims, and not just the demographic parity they promise, can be made assessable.

2. Auditing efforts

In the meantime, a new “algorithmic auditing” industry is blossoming in the absence of clear regulatory definitions of the term (Brown et al. 2021; Koshiyama et al. 2021; Bandy 2021).

Even proposed bills, such as New York City’s bill 1894 (Cumbo 2020), which mandates annual “bias audits” of hiring ADSs, does not dictate what precisely these would entail. This leaves a regulatory grey area upon which the relevance of the bill hinges. That is, depending on the depth, breadth, and focus of these audits, the bill could either enact unprecedented, meaningful steps towards enabling transparency and accountability, or it can, at worst, further remove agency from impacted individuals by creating a bureaucratic shield for unscrupulous companies to hide behind.

Within the hiring industry, we may encounter additional difficulties. First, hiring ADSs cannot be treated as individual models, but must be assessed as a linear series of interdependent models. In other words, the macro ADS is only as transparent and accountable as its weakest, or most intransparent and unaccountable, model or application. Additionally, any downstream model can only be maximally optimized to the most biased prior model. As an example - assume an organization has a simple macro ADS consisting of an ADS to identify and reach out to potential candidates and an NLP ADS to rank resumes by fit to the job description. Let us also assume that a company goal is to improve gender diversity and correct for biases introduced by historical data or flawed algorithmic designs. In this limited case, there is likely to be an outsized focus on the biases introduced by language models and the potential for gender-based language biases. However, even if this model were assessed and corrected for language biases, it is only as unbiased as the model before it. If the outreach ADS discriminates against non-male candidates, then our downstream model can only perform to the ceiling set by the prior, biased, model.

If we are to add in a third model, let’s say an “emotion detection” ADS that determines candidate trustworthiness by video interview, the additional complexity introduced is that the “bias” we are now considering does not exist as a function of data or model choice, but in the epistemological roots of the system. That is, a system that operates “as intended,” if the intention is based on pseudo-scientific and flawed research into imagined links between biology and trustworthiness, is a fundamentally broken system (Sloane, 2021). Investigations into disparate impact, gender distributions in the data, and the like, cannot account for or correct these problems.

3. The Socio-Technical Matrix for Assessing ADS

Against that backdrop, it is clear that in order to develop effective regulation we need new ways of framing and understanding how applicants encounter technological systems that build on pseudo-scientific theories to control labor. In order to aid this effort, we propose a matrix for developing a holistic view on hiring ADS by combining information on its context, its goal, its data, its function, its assumption, and epistemological roots. This matrix can serve as a template for workers, researchers, policymakers and practitioners alike to create new mechanisms for literacy, accountability, and oversight of ADSs, including audit and impact assessment.

3.1 How the Matrix Works

The matrix is a research tool that can serve as a basis for developing holistic socio-technical assessment and audit methods for hiring ADS. In order to use the matrix, information on the hiring ADS needs to be collected. Companies using such systems are not obliged to disclose to candidates or the public that they are using hiring ADSs, or which hiring ADSs they are using. However, information on these hiring ADSs can be found on the Internet as vendors advertise their products and services via case studies, or in federal trademark filings (Levendowski 2019).

Some vendors offer a single hiring ADS to be used for a narrow purpose (such as the use of a resume parser to narrow down prospects), while other companies offer a suite of hiring tools. For the purposes of this paper, the unit of analysis for the matrix is the ADS, not the company. There are seven elements comprising the matrix: ADS, funnel stage, goal, data, function, assumption, and epistemological roots. These seven elements comprise the description of a hiring system that needs to be assembled in order to assess its claims.

The **Hiring ADS** and **Funnel Stage** identify what the hiring ADS is (eg. Hubert.ai, ZipRecruiter) and how it is intended to be used (eg. recruiting, screening, etc.).

The **Goal** of the hiring ADS should clearly state what it aims to do (eg. "Filter the top 1% of applicants while maintaining the diversity of the applicant pool"). These three elements can be derived from sales copy, but should be supplemented by interviews with developers and hiring managers who purchase and use the hiring ADS.

Automated hiring ADSs use **Data**. Some hiring ADSs, particularly those that use machine learning or make claims about using artificial intelligence, use data other than that provided by a job applicant to sort, rank, filter, and predict performance for applicants. The matrix should help identify the data provided by applicants is processed (such as resumes, facial images, voice recordings, chatbot histories, gameplay).

The **Function** of a hiring ADS is a plain-language description of how it processes data to make its claims (eg. "compares resumes of previously successful employees to current applicants to predict future success"). Information pertaining to the function of an hiring ADS, how it processes data, and access to the hiring ADS itself should be procured through arrangements with developers and the hiring managers that configure and operate that hiring ADS. The model, or the hiring ADS itself, can subsequently be inspected as a part of an audit or impact assessment by examining the machine learning model that pursues this function in the context of training data, parameter settings, performance characteristics, and its integration into the hiring funnel.

The **Assumptions** that undergird a hiring ADS should capture the logic by which a hiring ADS is seen as useful, and can sometimes be derived from sales copy, but should a more thorough understanding of a hiring ADS's assumptions can be gathered from interviewing developers who create a hiring ADS and hiring managers who use a hiring ADS. These assumptions take the form of: "This hiring ADS works by comparing the resumes of successful employees to new applicants, because successful hires have proven that the attributes documented in their resumes are good predictors of future success."

Establishing the **Epistemological Roots** of a hiring ADS requires archival and/or ethnographic research to outline how a hiring ADS is understood to produce useful knowledge

about an applicant. The use of resumes, for example, has a long history in which the resume document itself has been constructed as a reasonable proxy on which to base a hiring manager's judgements about an applicant that need to be examined in their historical context. Similarly, hiring ADSs that analyze tone of voice to discern personality characteristics have their epistemological roots in psychological profiles of discrete "personality types" and physiognomic approaches that [spuriously] link biological components of vocalization to personality (Semel 2020).

Table 1: Matrix Elements

Element	Data	Questions and Method
Hiring ADS	Name of hiring ADS	Question: What is the name of the hiring ADS? Method: Identify from sales copy
Funnel Stage	Select from Bogen and Reike 2018	Question: At what stage does this company's hiring ADS operate? Method: Identify from sales copy and align with funnel list
Goal	Narrative description	Question: What is the hiring ADS intended to be used for? Method: Identify from sales copy, interview developers and hiring managers who operate the hiring ADS
Data	Inventory of data types, datasets, benchmarking datasets	Question: What data, and what types of data, are used in training, testing, and operating the hiring ADS? Method: Interview developers and hiring managers who operate the hiring ADS, inspect data directly
Function	Narrative description, machine learning models, metadata about models	Question: How does the hiring ADS work? What is it optimizing for? Method: Interview developers and hiring managers who operate the hiring ADS, inspect models, metadata, and product directly
Assumption	Narrative description	Question: Why is the hiring ADS useful? What is the assumed relationship between data about an applicant and the goals of the hiring manager? How does the hiring ADS inform the hiring process? Method: Interview developers and hiring managers who operate the hiring ADS
Epistemological Roots	Narrative description	Question: Where do the assumptions made by the hiring ADS come from? What is their intellectual lineage? What are the critiques of this lineage? Method: Archival research, interview developers and hiring managers who operate the hiring ADS, ethnographic study of hiring managers and developers

3.2 Using the Matrix

In this section, we want to demonstrate how the matrix can be used. The landscape of automated tools used in the context of hiring is vast and emerging. Most companies do not just use one hiring ADS, but combine various ADSs at various stages of the talent scouting and hiring process. Therefore, to demonstrate the use of the matrix in this short paper, we focus

on the second stage of the “hiring funnel” (Bogen and Reike 2018): screening. This is the stage where candidates are assessed whether or not they match a job. This assessment can be based on a myriad of aspects. Here, we want to focus on the aspects that become the basis for predicting job fit in automated hiring ADSs: experience, skill, ability, and personality.

Experience assessment is the most basic form of assessment used in hiring and often focuses on using an analysis of education, previous positions, and years of experience as a proxy for job fit and future job performance. A standard way in which experience assessments happen is via parsing a resume.

Skill assessment is a form of standardized testing that sets out to measure a candidate’s knowledge and skills that are needed for a particular role. For example, a very common skill assessment for programmers are so-called “coding challenges” whereby applicants are presented with typical programming challenges and have to solve live in a job interview.

Ability assessments typically refer to cognitive abilities tests. They are different from skill assessments, because they do not assess skill that is learned, but are based on the assumption that there are hidden mental abilities, such abstract thinking, understanding of complex concepts or adaptability to change, that do not necessarily show in a resume, a cover letter, or an interview.

Personality tests set out to determine personality traits in an individual, such as introversion or extroversion. Aside from the infamous Myers-Briggs Type Indicator, a popular taxonomy is the OCEAN model that models the “Big 5 Personality Traits”: openness to experience, conscientiousness, extraversion, agreeableness, and neuroticism (John et al. 2008). In psychology, these personality traits are assumed to be stable (Costa and McCrae 1986). Personality tests have a long history in corporate management in American corporations, and automating them as part of candidate assessment can be seen as falling well into the general shift towards the automation of general managerial decision-making (Lussier, 2018).

To show how the matrix can serve as a way to unpack how ADS construct experience, skill, ability, and personality, we are using a selection of hiring ADSs offered by various companies.

Table 2: Matrix Examples

Hiring ADS	Hiretual	Codility	Pymetrics	Humantic
Funnel Stage	Screening	Screening	Screening	Screening
Goal	Experience	Skill	Ability	Personality
Data	Resume Professional profiles Social media profiles Proprietary database	Coding test exercises	Gameplay scores from applicants and workers	Resume LinkedIn profile Twitter profile
Function	Use profiling for job matching	Use test performance for	Use gameplay performance for	Use personality profiling for job matching

		screening candidates in / out	screening candidates in / out	
Assumption	Professional and social profile can be matched to job fit	Code test performance is a predictor of job skills	Gameplay is a predictor of job success	Personality is a good predictor for job fit
Epistemological Roots	Social Network Theory: The idea that who you are connected with reveals your identity.	Vocational Aptitude Testing: The idea that test scores predict ability.	Eugenics: The idea that intelligence and ability are innate and can be revealed through testing.	Physiognomy: The idea that personality traits are discernible from appearance.

4. Cues for auditability

The matrix can serve as a tool for developing new avenues for technical work on the auditability of algorithms. By helping to identify the concepts that the hiring ADS claim to measure and rank, the matrix can serve as a basis for more solid validity and reliability assessments. Validity generally refers to the extent that a statistical tool measures what it is supposed to measure.

The matrix unpacks how the hiring ADS constructs what it is supposed to measure and rank, for example experience, skill, ability, and personality. Social research into the scientific and narrative roots of these frameworks can then help assess what we can take as a baseline for validity which we can then assess in an empirical study or audit. In the context of hiring ADS, we may do so by examining stability, i.e. assess if the construction (measurement and ranking) of any given framework is stable across systems, and a number of other input factors that can be changed (such as file type), in a stable way. If the concept assumes that any of these individual features are somewhat stable across lifetime - such as ability or personality - then any hiring ADS claiming to “objectively” measure those features, should equally do so in a stable way. An empirical study, or audit, of these hiring ADS using the same sample of individuals can reveal instability in prediction across trials.

5. Cues for Regulatory Audit Mandates

With the matrix as a starting point, we can better appreciate the complexity of auditing or analyzing hiring ADS and how this relates to regulatory audit mandates. Traditional audit methodologies may ask about the Goal, Data, and Function of the ADS, but generally do not address issues of cross-model contamination (e.g. Funnel Stage), Assumptions, and Epistemological Roots.

Therefore, a successful hiring ADS audit requires an interdisciplinary group, following from work on algorithmic impact assessment that calls for multidisciplinary perspectives on identifying algorithmic harms (Metcalf et al. 2021). In addition to legal and data science teams that are currently engaging in algorithmic audit work, we suggest the addition of social scientists, psychologists, and historians of science and technology to critically evaluate assumptions and epistemologies, and inform the audit process as a whole.

Auditing has also become a bit of a catch-all phrase, and there is value to parsing out different types of audits based on purpose and audience. In fields such as healthcare and finance, where audits are the norm, audit functions can be divided into two audiences: internal and external. Internal auditors are employed by the company, and external auditors can be a regulatory agency or a third-party group. Third-party groups can be a private firm specializing in audits or a potential client that may use the hiring ADS and wants to conduct their own audit. Further, the private firm may be compensated by a potential hiring ADS client OR by the company itself.

Whether the audit body is internal or external has significant impacts on: accessibility to models and data, chronology (i.e., when an audit is conducted in the development of this model), ability to assess cross-model contamination, and incentives.

Internal audit bodies have less external credibility but better access. In general, internal audit bodies serve to ensure the system is compliant with existing laws and addresses reputational risks. Internal audit groups may have access to data, models, IP, and key employees, which can mean engaging in the earliest stages of development, working closely with developers, data scientists, and project leads at milestones, and mitigating harm before there is adverse impact. These individuals are incentivized to ensure the company performs well, which can put into question the viability of fundamentally addressing issues of false assumptions and flawed epistemologies. Internal audit can range dramatically as ADS have no norms or laws dictating audit work. It is rarely the case that internal audits serve as external validation due to conflicts of interest. Generally, the role of internal audit is to ensure the company avoids legal or reputational backlash once the product is launched.

External audit bodies have more credibility but less access. Even if the audit is paid for by the company, external auditors have pressure to provide quality audit services to retain a good reputation. Regulatory bodies also publish their audit frameworks to allow internal audit teams to ensure adherence, but this also allows for public accountability. However, these auditors are not often granted unconstrained access to data, models, IP, or employees, and, in the current regulatory vacuum, companies have a heavy hand in creating these constraints. External audit bodies, however, are better able to critically analyze fundamentals such as assumptions and epistemologies as well as cross-model contamination (see Moss et al. 2021).

6. Directions for Future Work

In the matrix, we outline a series of questions that need to be answered to understand how hiring ADSs are intended to work and how they actually operate in practice. Several of those questions require ethnographic investigation into the contextual uses and understanding of these hiring ADSs.

Some aspects of this ethnographic investigation are already well-established, particularly for conducting ethnographic interviews (Spradley 1979), undertaking workplace ethnography (Neyland 2008, Ladner 2014), and merging ethnographic fieldwork with archival research (Merry 2002). But methodological innovation is called for in tailoring ethnographic interviews, archival research, and fieldwork to robust accountability processes. Additional work is also needed beyond the investigation of specific hiring ADSs to better understand how hiring managers, workers, applicants, and others within an organization interact with each other and with hiring tools. Who are the hiring managers using the hiring ADS? When are hiring ADSs used in the pipeline and by whom?

Not all hiring ADSs are used the same way by hiring managers, some might take hiring ADSs as suggestions for their own decision-making, others might implement a hiring ADS's outputs directly, and the ways a hiring ADS functions within a hiring managers workflow ought to be inspected as part of any audit or impact assessment. Building out an understanding of how hiring ADS are used in the workplace is a job for ethnography. Gaining ethnographic insight into how job applicants experience and ascribe meaning to hiring ADS is equally important. Future work can and must focus on this side of the social practice of hiring. This ethnographic insight, then, will be essential not just for designing impactful audits, but also for - eventually - creating less invasive and discriminatory hiring technologies, for example in collaboration with worker organizations and unions.

7. Conclusion

In this paper, we have suggested a systematic approach for developing socio-technical assessment for hiring ADS. We have developed a matrix that can serve as a research tool for identifying the concepts that the hiring ADS claim to measure and rank, as well as the assumptions rooted in pseudoscientific essentialized understandings of human nature and capability that they may be based on. We have argued that the matrix can serve as a basis for more solid validity and reliability assessments, as well as a basis for critically investigating emerging auditing standards and practices that fail to address issues around essentialized categories and behaviors.

ACKNOWLEDGMENTS

The authors wish to thank the NYU Center for Responsible AI, Tandon School of Engineering, New York University and Tübingen AI Center, University of Tübingen.

REFERENCES

- Aguera y Arcas, Blaise, Alexander Todorov, and Margaret Mitchell. 2017. "Physiognomy's New Clothes." Medium. May 20, 2017.
<https://medium.com/@blaisea/physiognomys-new-clothes-f2d4b59fdd6a>.
- Ajunwa, Ifeoma. 2021. "The Auditing Imperative for Automated Hiring." Harv. J.L. & Tech.
- Bandy, Jack. 2021. "Problematic Machine Behavior: A Systematic Literature Review of Algorithm Audits." ArXiv:2102.04256 [Cs], February.
<http://arxiv.org/abs/2102.04256>.
- Bogen, Miranda, and Aaron Reike. 2018. "Help Wanted: An Exploration of Hiring Algorithms, Equity and Bias." Washington, DC: Upturn.
- Brown, Shea, Jovana Davidovic, and Ali Hasan. 2021. "The Algorithm Audit: Scoring the Algorithms That Score Us." Big Data & Society 8 (1): 205395172098386.
<https://doi.org/10.1177/2053951720983865>.
- Burrell, Jenna. 2016. "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms." Big Data & Society 3 (1): 205395171562251.
<https://doi.org/10.1177/2053951715622512>.

- Costa, P. 1986. "Personality Stability and Its Implications for Clinical Psychology." *Clinical Psychology Review* 6 (5): 407–23. [https://doi.org/10.1016/0272-7358\(86\)90029-2](https://doi.org/10.1016/0272-7358(86)90029-2).
- Cumbo, Laurie A. 2020. The New York City Council - File #: Int 1894-2020.
- Durupinar, F, N Pelechano, J M Allbeck, Ugur Gudukbay, and N I Badler. 2011. "How the Ocean Personality Model Affects the Perception of Crowds." *IEEE Computer Graphics and Applications* 31 (3): 22–31. <https://doi.org/10.1109/MCG.2009.105>.
- John, Oliver P., Richard W. Robins, and Lawrence A. Pervin, eds. 2008. *Handbook of Personality: Theory and Research*. 3rd ed. New York: Guilford Press.
- Koshiyama, Adriano, Emre Kazim, Philip Treleaven, Pete Rai, Lukasz Szpruch, Giles Pavey, Ghazi Ahamat, et al. 2021. "Towards Algorithm Auditing: A Survey on Managing Legal, Ethical and Technological Risks of AI, ML and Associated Algorithms." *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3778998>.
- Ladner, Sam. 2014. *Practical Ethnography: A Guide to Doing Ethnography in the Private Sector*. Walnut Creek, CA: Left Coast Press.
- Levendowski, Amanda, Trademarks as Surveillance Transparency (October 31, 2019). 36 Berkeley Tech. L. J. (2021 Forthcoming), Available at SSRN: <https://ssrn.com/abstract=3544195>
- Lussier, Kira. 2018. "Personality, Incorporated: Psychological Capital in American Management, 1960-1995." Toronto, ON: University of Toronto.
- Merry, Sally Engle. 2002. "Ethnography in the Archives." In *Practicing Ethnography in Law*, edited by June Starr and Mark Goodale, 128–42. New York: Palgrave Macmillan US. <https://doi.org/10.1007/978-1-137-06573-5>.
- Metcalf, Jacob, Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, and Madeleine Clare Elish. 2021. "Algorithmic Impact Assessments and Accountability: The Co-Construction of Impacts." In *Proceedings of the ACM Conference on Fairness, Accountability and Transparency*. Toronto, ON: ACM. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3736261.
- Moss, Emanuel, Elizabeth A. Watkins, Ranjit Singh, and Jacob Metcalf. 2021. "Assembling Accountability through Algorithmic Impact Assessment." New York: Data & Society Research Institute.
- Neyland, Daniel. 2008. *Organizational Ethnography*. Los Angeles: Sage.
- Pasquale, Frank. 2015. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge: Harvard University Press.
- Sanchez-Monedero, Javier, Lina Dencik, and Lilian Edwards. 2020. "What Does It Mean to Solve the Problem of Discrimination in Hiring? Social, Technical and Legal Perspectives from the UK on Automated Hiring Systems." *ArXiv:1910.06144 [Cs]*, January. <http://arxiv.org/abs/1910.06144>.
- Semel, Beth. 2020. "The Body Audible: From Vocal Biomarkers to a Phrenology of the Throat." *Somatosphere* (blog). September 21, 2020. <http://somatosphere.net/2020/the-body-audible.html/>.

- Sloane, Mona. 2021. "The Algorithmic Auditing Trap." Medium OneZero. March 17, 2021. <https://onezero.medium.com/the-algorithmic-auditing-trap-9a6f2d4d461d>.
- Spradley, James P. 1979. *The Ethnographic Interview*. New York: Holt, Rinehart and Winston.
- Stark, Luke, and Jesse Hoey. 2020. "The Ethics of Emotion in Artificial Intelligence Systems." OSF Preprints, 12. <https://doi.org/10.31219/osf.io/9ad4u>.
- van den Broek, Elmira, Anastasia Sergeeva, and Marleen Huysmann. 2019. "Hiring Algorithms: An Ethnography of Fairness in Practice." In *Proceedings of the 40th International Conference on Information Systems*, 9. Munich, Germany.
- Wilson, Christo, Avijit Ghosh, Shan Jiang, Alan Mislove, Lewis Baker, Janelle Szary, Kelly Trindel, and Frida Polli. 2021. "Building and Auditing Fair Algorithms: A Case Study in Candidate Screening," 12.