The Promise and Threat of Artificial Intelligence in Combating (or Worsening) Employment Discrimination
NAPABA Convention Session #506

Saturday, November 9, 2019
10:45 a.m.-12:00 p.m.

By Jenny R. Yang[1] and Rachel See[2]

Artificial intelligence ("AI") is transforming everything from the way we work, to the way we deliver health care, to how we buy homes, and even how our criminal justice system operates. As these systems rapidly change our world, it is essential that the legal community invests in understanding exactly how AI is changing our society and its implications for inclusion, equity, and social justice. This is especially important in the context of labor and employment law. Increasingly, employers are leveraging AI-powered algorithms to recruit, screen, interview, and rank candidates. AI is often deciding who sees a particular job posting, and a substantial number of job applicants are automatically or summarily rejected by AI-driven screens — in this sense, hiring algorithms act as modern gatekeepers to economic opportunity.

Almost 55 years after the passage of the Civil Rights Act of 1964, workplace inequality persists. For example, research has found that resumes with English-sounding names receive requests for interviews 40% more often than identical resumes with Chinese, Indian, or Pakistani names.[3] AI-powered hiring screens often promise to help employers efficiently identify candidates based on specific criteria while mitigating the bias and subjectivity that may arise with human decision making. Yet, algorithms can also replicate and deepen existing inequities – algorithms do not reveal an objective truth simply because they are based on math. For example, hiring algorithms trained on inaccurate, biased, or unrepresentative data can produce employment outcomes biased along lines of race, sex, or other characteristics protected by antidiscrimination law.

In 2014, Amazon started building an algorithmic recruiting tool to review applicants' resumes and rank top talent for the company.[4] The company hoped to leverage automation in hiring to increase its competitive edge in the same ways that automation had bolstered its warehouse operations and pricing models. Amazon trained the tool to identify patterns based on resumes the company received over a ten-year period. Unfortunately, men had submitted most of those resumes, leading to the algorithm to prefer male candidates and penalize resumes that included phrases like "Society

[1] Jenny R. Yang is a Senior Policy Fellow at the Urban Institute. She served as a Commissioner and Chair of the U.S. Equal Employment Opportunity Commission ("EEOC") from 2013-2018. She thanks Courtney Hinkle, a law student at Georgetown University Law Center, for her assistance with this paper.
[2] Rachel See, serves as Assistant General Counsel for Technology for the EEOC. She manages the Commission's use of technology in its nationwide litigation program.
[3] Jane Burnett, *Strong Job Candidates With Foreign Names Miss Out On Job Interviews, Study Shows*, Ladders (February 24, 2017), https://www.theladders.com/career-advice/study-ethnic-sounding-name-employers-fewer-calls-back.
[4] Jeffrey Dastin, *Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women*, Business News (October 9, 2018), https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G.

for Women Scientists" or graduates of all-women colleges because they contained the word "women." By 2017, the company abandoned that project, acknowledging it could not guarantee the tool would not develop other discriminatory ways of sorting candidates. In another highly publicized incident, a resume screening service being vetted by a potential buyer discovered that its model identified being named "Jared" and playing high school lacrosse as the strongest predictive indicators of success, despite neither qualification having a causal link to job-related performance.[5]

Amazon is certainly not alone in its embrace of AI to help narrow an applicant pool and identify the strongest candidates. Faced with a deluge of online applications, many employers, including Target, Hilton, Cisco, PepsiCo, Ikea, along with major staffing agencies, are turning to these data-driven, predictive tools to inform various stages of their hiring process.[6] In a recent survey by talent software firm CareerBuilder, more than 55 percent of U.S. human resources managers said AI would be a regular part of their work within the next five years.[7] A recent LinkedIn survey of 9,000 hiring managers and recruiters found over half of respondents identified data analytics as "very" or "extremely" important, and nearly one-fifth stated they had "mostly" or "completely adopted" its use in their hiring practices.[8]

Vendors of these algorithmic hiring screens are estimated to be a $500 million industry.[9] However, these products are only a small fraction of the global artificial intelligence software market, which, according to Fortune Business Insights, is expected to experience massive growth in the coming years, with revenues increasing from around $9.5 billion in 2018 to an expected $118.6 billion by 2025.[10] Many employers are motivated by the promise of greater efficiency and cost savings provided by automated decision making. However, companies may rush to automate existing recruitment and hiring processes without a clear understanding of how bias and discrimination may be replicated from past hiring decisions.

Certainly, AI-powered tools, if deployed correctly, offer an opportunity to address workplace equity, and bolster efforts to increase diversity and mitigate subjective bias throughout the hiring process. For example, Catalyte uses AI and predictive analytics to identify people, regardless of background, who have the potential to succeed as software developers.[11] Some employers are

---

[5] Dave Gershgorn, *Companies Are On The Hook If Their Hiring Algorithms Are Biased*, Quartz (October 22, 2018), https://qz.com/1427621/companies-are-on-the-hook-if-their-hiring-algorithms-are-biased.

[6] Miranda Bogen & Aaron Rieke, *Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias*, Upturn, (December 2018), https://www.upturn.org/static/reports/2018/hiring-algorithms/files/Upturn%20--%20Help%20Wanted%20-%20An%20Exploration%20of%20Hiring%20Algorithms,%20Equity%20and%20Bias.pdf

[7] Jeffrey Dastin, *supra* Note 2.

[8] LinkedIn Talent Solutions, *Global Recruiting Trends 2018*, https://business.linkedin.com/content/dam/me/business/en-us/talent-solutions/resources/pdfs/linkedin-global-recruiting-trends-2018-en-us2.pdf

[9] Cathy O'Neil, *Weapons of Math Destruction: How Bid Data Increases Inequality and Threatens Democracy*, at 108 (2016).

[10] Fortune Business Insight, *Artificial Intelligence AI) Market Size, Share and Industry Analysis* (April 2019; last updated Sept. 2019), https://www.fortunebusinessinsights.com/industry-reports/artificial-intelligence-market-100114.

[11] Catalyte – About Us, https://catalyte.io/about-us/ (last visited Sept. 25, 2019).

deploying tools such as TapRecruit, which leverages AI to optimize the language used in job descriptions to attract a more qualified and diverse applicant pool.[12] Using predictive analytics, language detectors can filter out gender-biased wording in job descriptions to attract a more diverse and qualified applicant pool and identify patterns of bias in performance feedback. Anonymous recruitment processes can encourage recruiters to focus on skills, rather than a candidate's first or last name. And tools can alert hiring managers when someone is consistently assigned fewer or less important tasks that may be the result of implicit bias.

Given the potential promise but also risks of unlawful discrimination furthered by AI-powered hiring algorithms, it is critical that employers fully understand the capabilities and drawbacks of AI. In particular, labor and employment attorneys have ethical obligations to comprehend how these technologies are being used and their effects in order to offer adequate client counsel and ensure these tools are compliant with prohibitions on hiring discrimination. This paper outlines (1) professional responsibility rules for lawyers to remain current with emerging technologies in recruitment and hiring; (2) the ways these emerging technologies in recruitment and hiring may introduce bias at each stage of the hiring process; and (3) legal issues for attorneys to understand about AI and the existing legal civil rights framework as well as advice on best practices.

I. Ethics/Professional Responsibility Rules

All attorneys are expected to satisfy their ethical obligation of providing competent representation of their clients. Since 1983, ABA Model Rule 1.1 has required that attorneys "provide competent representation to a client" and must maintain "the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation."[13] Moreover, the ABA Model Code of Professional Responsibility advises a lawyer "should strive to become and remain proficient in [] practice and should accept employment only in matters which [the lawyer] *is* or *intends to become* competent to handle."[14] A lawyer may face disciplinary action where the lawyer knew or should have known that he or she lacked the ability to provide competent service.[15]

These rules have heightened significance given the shifting technological landscape that is directly impacting the practice of law. In 2012, in response to recommendations from the ABA's Ethics 20/20 Commission, the ABA amended comment 8 to Model Rule 1.1, stating, "To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology."[16] To date, 37 states have adopted the duty of technological competence.[17]

---

[12] TapRecruit, https://taprecruit.co/ (lasted visited Sept. 25, 2019).
[13] *See* Model Rules of Prof'l Conduct Preface (Am. Bar Ass'n 1983).
[14] *Id.* at EC 6-1 (emphasis added).
[15] Model Code DR 6-101(a)(1).
[16] Model Rules 1.1 cmt. 8 (emphasis added).
[17] *See* Robert Ambrogi, *Tech Competence*, LawSites, https://www.lawsitesblog.com/tech-competence (last visited Sept.19, 2019).

Despite an affirmative duty to maintain technological competence, legal publications are rife with recent examples of non-technical lawyers running afoul of technological mishaps.[18] Of course, lawyers being slow to adopt or to understand technology is nothing new to the profession; a recent Florida Bar Journal article pointed out that lawyers were slow to adopt the telephone.[19] Nevertheless, most courts have held that technological ignorance is not an excuse with respect to lawyers mishandling electronic discovery issues.[20]

In the realm of electronic discovery ("e-discovery"), a vocal number of judges have been sounding the alarm regarding the lack of technological competence of lawyers appearing in front of them. When discussing crafting of keyword search terms and how keyword searches must be tested to ensure accuracy, Magistrate Judge Andrew Peck of the Southern District of New York proclaimed, "It is time that the Bar — even those lawyers who did not come of age in the computer era — understand this."[21]

Text analytics and predictive analytics technology – also known as AI – has been part and parcel of cutting-edge e-discovery discussions for several years now. The first written opinion approving the use of this technology was in 2012.[22] Since then, there have been numerous cases addressing the use of this technology.[23] Concepts that e-discovery practitioners have been applying in the context of predictive analytics are now being applied to other search methodologies.[24]

While a discussion of what constitutes cutting-edge e-discovery technology in recent years is beyond the scope of this article, the fact remains that attorneys involved in complex litigation –

---

[18] *See, e.g.*, Mike Scarcella, *Jones Day Apologizes for Botched Filing That Revealed Grand Jury Info*, National Law Journal (September 13, 2019), https://www.law.com/nationallawjournal/2019/09/13/jones-day-apologizes-for-botched-filing-that-revealed-grand-jury-info/?slreturn=20190825114206 (discussing letter apologizing submission of improperly redacted brief that revealed grand jury information; issue "involved Microsoft Word and printing to Adobe Acrobat, rather than the redaction software [Jones Day] has in place"). *See also* Dan Packel, *US Law Firm Falls Victim to Alleged Chinese Hacking as Clients Face Threat*, The American Lawyer (February 20, 2019), https://www.law.com/americanlawyer/2019/02/20/law-firms-still-a-great-target-as-us-firm-falls-victim-to-alleged-chinese-hack/.

[19] *See* Mark Britton*, Behind Stables and Saloons: The Legal Profession's Race to the Back of the Technological Pack*, 90 FLA. B.J. 34 (Jan. 2016) ("In 1891, 7,000 businesses in the New York/New Jersey area had telephones. Among those, there were 937 doctors, 363 saloons, 315 stables, and last were 146 lawyers.")

[20] *See James v. National Financial*, C.A. No. 8931-VCL (Del. Ch. Dec. 5, 2014) (sanctioning computer illiterate attorney and admonished that "[p]rofessed technological incompetence is not an excuse for discovery misconduct."). *See also Green v. Blitz U.S.A., Inc.*, 2011 WL 806011 (E.D. Tex. Mar. 1, 2011) (sanctioning responding party who delegated electronic discovery process to an employee who professed he was ""about as computer…illiterate as they get.").

[21] See *William A. Gross Construction Associates, Inc. v. American Manufacturers Mutual Insurance Co.*, 256 F.R.D. 134, 136 (S.D.N.Y. 2009) (Peck, M.J.).

[22] *Monique da Silva Moore, et al., v. Publicis Groupe SA & MSL Group*, No. 11 Civ. 1279 (ALC)(AJP) (S.D.N.Y. Feb. 24, 2012) (Peck, M.J.).

[23] *See, e.g.*, *Rio Tinto PLC v. Vale S.A.*, 14 Civ. 3042, 2015 WL 872294 (S.D.N.Y. March 2, 2015); *In re Broiler Chicken Antitrust Litigation*, 2018 WL 1146371 (N.D. Ill. Jan. 3, 2018).

[24] *See, e.g.*, *City of Rockford v. Mallinckrodt,* 326 F.R.D. 489 (N.D. Ill. 2018) (stating to validate keyword search, null-set sampling at a 95% +/-2% margin of error).

including complex labor and employment litigation – are increasingly being forced to come to grips with the use of predictive analytics and machine learning as part of fulfilling their discovery obligations.

Comment 18 to Comment 8 of ABA Rule 1.1 states, “In the technological realm, competence is seen as having a basic understanding of the implications that new and emerging technologies can have for a lawyer’s day to day practice.”[25] Predictive analytics and machine learning are established technologies that are already part of the day to day practice of lawyers engaged in litigation involving document-intensive e-discovery. Therefore, an attorney’s affirmative ethical obligation to understand the risks and benefits associated with “relevant” technology includes artificial intelligence.

Outside of the duty of competence, the use of AI presents numerous ethical challenges outside of the scope of this discussion. Still, ABA Model Rule 2.1 further provides, “In rendering advice, a lawyer may refer not only to law but to other considerations such as moral, economic, social and political factors, that may be relevant to the client's situation.”[26] Comment 2 to Rule 2.1 provides, “It is proper for a lawyer to refer to relevant moral and ethical considerations in giving advice. Although a lawyer is not a moral advisor as such, moral and ethical considerations impinge upon most legal questions and may decisively influence how the law will be applied.”[27]

Understanding the technical framework and issues raised by the use of AI is a fundamental starting point in providing effective legal and ethical advice. Hot-button issues include Google’s involvement in “Project Maven,” a Department of Defense project intended to use machine learning and AI in order to differentiate people and objects in thousands of hours of drone footage. Over 3,000 Google employees signed a letter urging Google’s CEO to suspend the deal: “We believe that Google should not be in the business of war.”[28] In response to the outcry, Google has withdrawn from Project Maven.[29] Additionally, the company declined to bid on the Pentagon’s Joint Enterprise Defense Infrastructure (JEDI) contract because it “couldn't be assured that it would align with our AI Principles,” which, among other things, stated that Google would not design or deploy AI for weapons, surveillance, or other technology designed principally to injure people.[30] While questions regarding the ethics involving the military use of AI may seem remote from many lawyers’ daily practice, the legal landscape is rapidly developing in other areas that implicate complicated technological and ethical questions.

---

[25] Model Rule 1.1, cmt. 18 to cmt. 8

[26] Model Rule 2.1

[27] Model Rule 2.1 cmt. 2

[28] Scott Shane and Daisuke Wakabayashi, *‘The Business of War’: Google Employees Protest Work for the Pentagon*, N.Y. Times (April 4, 2018), https://www.nytimes.com/2018/04/04/technology/google-letter-ceo-pentagon-project.html

[29] Daisuke Wakabayashi and Scott Shane, *Google Will Not Renew Pentagon Contract That Upset Employees*, N.Y. Times (June 1, 2018), https://www.nytimes.com/2018/06/01/technology/google-pentagon-project-maven.html

[30] Liam Tung, *Google: Here’s Why We’re Pulling Out Of Pentagon’s $10bn JEDI Cloud Race*, ZD Net (Oct. 9, 2018), https://www.zdnet.com/article/google-heres-why-were-pulling-out-of-pentagons-10bn-jedi-cloud-race/

II. How AI Recruiting and Hiring Tools Operate

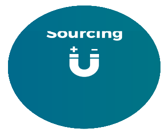

8, https://www.upturn.org/reports/2018/hiring-algorithms/

So how exactly are AI-powered screens changing modern recruiting and hiring practices? Predictive technologies can play very different roles at various stages of the hiring process from determining who sees job advertisements, to screening and predicting an applicant's performance and "fit," to forecasting a candidate's salary requirements.[31]

In the "sourcing" stage, predictive technologies can be used to optimize placement of job advertisements, to market jobs to certain jobseekers, or to identify candidates most likely to be poached from a competitor.[32] These technologies can help artificially shape the universe of potential candidates. In the digital era, online advertising is increasingly the preferred way to reach prospective applicants, and employers may use recruiting algorithms to identify well-qualified individuals and encourage them to apply for a particular job. In particular, recruiters are relying upon social media platforms, taking advantage of the opportunity to recruit passive job seekers already visiting the site for social purposes.[33]

In order to facilitate digital recruitment, social media companies offer tools to assist employers in targeting particular audiences based upon preferred applicants. For example, Facebook mines data from its billions of global users, and infers characteristics, interests, and preferences based on a user's profile and previous behaviors. Facebook then allows recruiters to micro-target certain job posts to a preferred audience using its ad delivery algorithm.[34] A preferred audience can be identified based on available filters, such as location, age, gender, or likes and pages a user has

---

[31] *See* Miranda Bogen, *supra* Note 6.
[32] *Id.*
[33] *See* AdWeek Survey, 92% of Recruiters Use Social Media to Find High-Quality Candidates (Sept. 22, 2015), https://www.adweek.com/digital/survey-96-of-recruiters-use-social-media-to-find-high-quality-candidates/ (reporting 92 percent of recruiters used social media to recruit job applicants); Society for Human Resource Management (SHRM), *SHRM Survey Findings: Using Social Media for Talent Acquisition—Recruiting and Screening* (Jan. 7, 2016), at 11, https://www.shrm.org/hr-today/trends-and-forecasting/research-and-surveys/Documents/SHRM-Social-Media-Recruiting-Screening-2015.pdf (showing research that the opportunity recruit 'passive' candidates who might not otherwise apply is a top reason for recruiters using social media).
[34] Facebook Business – Ads, Audience https://www.facebook.com/business/help/182371508761821 (last visited Sept. 25, 2019).

engaged with. In addition, Facebook offers recruiters the option of reaching beyond an initial targeted audience to "lookalike audiences" that share common qualities, such as demographic information or interests.[35]

While these tools help to focus the universe of audiences seeing the job ads, they risk intentionally excluding individuals based on protected characteristics. For example, a recent study found that Facebook's ad delivery system disproportionately showed lumber industry job ads to white users and taxi driver job ads to Black users.[36] And as one legal scholar noted, "not informing people of a job opportunity is a highly effective barrier" to applying and ultimately, receiving an offer.[37]

Social media companies, as well as employers purchasing these ads, have come under scrutiny for discriminatory advertising. Earlier this year, Facebook settled a class action lawsuit accusing the company of failing to prevent discrimination based on race, age, and gender in employment, housing and credit advertising.[38] As part of the settlement agreement, Facebook agreed to:

> (1) create a separate portal with more limited target options so advertisers cannot target ads based on users' age, gender, race, or categories associate with membership in protected groups, or zip codes or geographic area that is less than 15-mile radius, and cannot considered users' age, gender, or zip code in "lookalike" audiences; (2) implement a system of automated and human review of ads; (3) create a "self-certification" requirement for all advertisers to certify compliance with anti-discrimination laws, and provide education for advertisers on those laws; (4) study the potential for unintended biases in algorithmic modeling on Facebook; and (5) monitor the implementation of the reforms that Facebook is undertaking for three years.[39]

Other civil rights actions have recently been filed challenging this discriminatory advertising practice, illustrating the legal liability that comes with use of these AI-powered screens.[40]

In the "screening" phase, employers are formally reviewing applications, and choosing between unqualified or relatively weak applicants or identifying strong applicants to give closer consideration. Predictive hiring screens can assist in this initial sorting by assessing, scoring, and

---

[35] Facebook Business – Ads, Lookalike Audiences, https://www.facebook.com/business/help/164749007013531 (last visited: Sept. 25, 2019).
[36] Louis Matsakis, *Facebook's Ad System Might Be Hard-Coded For Discrimination*, Wired (April 6, 2019), https://www.wired.com/story/facebooks-ad-system-discrimination/
[37] Pauline Kim & Sharion Scott, *Discrimination in Online Employment Recruiting*, St. Louis Univ. L.J., Volume 63, No. 1, 12 (2019).
[38] Outten & Golden, LLP, *Facebook Agrees to Sweeping Reforms to Curb Discriminatory Ad Targeting Practice* (March 19, 2019), https://www.outtengolden.com/news/facebook-agrees-sweeping-reforms-curb-discriminatory-ad-targeting-practices.
[39] Adam T. Klein and Nantiya Ruan, *Attorney Competence in the Algorithm Age*, 72nd Annual NYU Conference on Labor – AI & Automation: Impact on Work and Workers (June 14, 2019), 1, 13 (citing *Bradley et al. v T-Mobile U.S., Inc. et al.*, No. 5:17-cv-07232-BLF (N.D. Cal. 2018).
[40] *See, e.g., Bradley et al. v T-Mobile U.S., Inc. et al.*, No. 17-cv-07232-BLF (N.D. Cal. 2018) (arguing defendants, as job advertisers, purposefully excluded older people from seeing their Facebook ads); *National Fair Housing Alliance v. Facebook*, No. 1:18-cv-02689 (S.D.N.Y. 2018) (arguing a challenge of discrimination in housing ads based on race, national origin, sex, familiar status, and disability).

ranking applicants for minimum and preferred qualifications or skillsets. These tools range from basic screening questions[41] to review of resumes with machine learning techniques[42] to predictive assessments using online tests[43] or "neuroscience" web games.[44]

While AI-powered screens might have the appearance of objectivity, they can introduce bias at several points in the process, including in how the assessments are trained, built, validated, and monitored for bias. Algorithms are often developed based on data from past hiring decisions and evaluations of top performers at the company. However, using this historic data can reproduce inequity, even when tools explicitly ignore race, gender, age, and other protected bases.[45]

The information that humans have selected for consideration may incorporate prior interpersonal, institutional, and systemic social biases that will screen out certain groups that have been historically underrepresented in a workplace. When the data used to train the algorithm is not diverse, the analytical models may build in barriers to certain groups that have been underrepresented—including candidates who could perform the job as well, or better, but have a very different profile than current top performers.

Additionally, bias and discrimination may be introduced into a system when recruiters and hiring managers misinterpret or place undue weight on results and recommendations generated by AI-powered tools. If algorithms are fed information and data points about a company's top performers, the algorithm will produce a profile and then predict who will be a successful candidate based on their similarity to that profile. The algorithm may be matching characteristics of people rather than factors causally linked with job performance, such as ability or skills. Ultimately, algorithms are designed by people, and people build their own biases and stereotypes into the models. Compounding the risk of these algorithms is the lack of racial and gender diversity in the technology sector, increasing the possibility that bias and patterns of discrimination in the models will go undetected.[46]

Unfortunately, there is little transparency surrounding the methods used or validation results of these algorithmic tools. For example, one vendor offers a prescreening assessment test to measure aptitude, skills, and personality traits for top perform applicants modeled off of responses from top performers currently at the company.[47] Although the vendor has stated the test is validated and does not have a discriminatory effect, the methodology and results remain a secret, and it is unclear how much the algorithms might account for diversity of the current workforce, or historic barriers to advancement at the company.[48]

---

[41] Mya –About Us, https://mya.com/about (last accessed: Sept. 25, 2019).
[42] Ideal – Resume Screening, https://ideal.com/resume-screening/ (last accessed: Sept. 25, 2019).
[43] Koru – About Us, https://www.joinkoru.com/about/ (last accessed: Sept. 25, 2019).
[44] Pymetrics – About Us, https://www.pymetrics.com/about/ (last accessed: Sept. 25, 2019).
[45] *See* Solon Barocas and Andrew D. Selbst, *Big Data's Disparate Impact*, 104 California L. Rev. 671 (2016).
[46] S.M. West et al., *Discriminating Systems: Gender, Race and Power in AI*, AI Now Institute (April 2019).
[47] Miranda Bogen, *supra* Note 6, at 29-32.
[48] *Id.*

In addition, at the interviewing stage, some employers are using automated voice scans or video interviews that can measure an applicant's facial expression and eye contact, tone or word choice.[49] Many employers hope these tools can limit subjective decision making at this stage, where emphasis is often placed on soft-skills and overall "fit." However, these technologies have sparked public concern about bias, as the facial analysis systems often fail to accurately read faces of women of color, and voice scans struggle to understand region and nonnative accents.[50] It is also unclear how these types of tests are causally related to job performance.[51]

III. Legal Landscape & Future Best Practices

For policymakers, regulators and interested stakeholders across a broad range of industry sectors, it can be difficult to comprehend the expanding ways AI is rapidly being deployed. In the hiring context, there is surprisingly little research on how AI-powered screens are actually being utilized and their effect. Many vendors offering these tools are reluctant to make information publicly available, citing the need to protect proprietary information on their algorithms. According to one study, vendors' websites generally do not make clear what validity models are used, what methodologies are deployed, how the validation data is selected, or how these procedures might vary depending on the particular client.[52] Greater transparency and understanding is essential as predictive models can create concerns about replicating discrimination and bias in hiring and could violate longstanding civil rights laws prohibiting discrimination on the basis of protected characteristics.

Title VII of the Civil Rights Act of 1964 prohibits employers from discriminating on the basis of race, color, religion, sex, and national origin. The statute prohibits discrimination based on: (1) disparate treatment; and (2) disparate impact. Under "disparate treatment," Title VII prohibits *intentional* discrimination on protected categories, such as a covered employer testing the reading ability of African American applicants or employees but not testing the reading ability of their white counterparts.[53] Under "disparate impact," Title VII prohibits employers from using neutral tests or selection procedures that have the *effect* of disproportionately excluding persons based on protected categories where the tests or selection procedures are not "job-related and consistent with business necessity."[54] Employment screens can also violate the Americans with Disabilities Act ("ADA"), and the Age Discrimination in Employment Act ("ADEA").

Legal questions surrounding the use of algorithmic screens in hiring build on decades of guidance, litigation and research on the appropriate use of selection devices. For example, there has been much litigation over the use of cognitive tests that assess applicants on areas such as

[49] *Id.* at 36-37.
[50] *Id.*
[51] *Id.*
[52] Manish Raghavan & Solon Barocas, et al., *Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices*, at 8 (June 21, 2019), https://arxiv.org/abs/1906.09208#.
[53] Equal Employment Opportunity Commission – Employment Tests and Selection Procedures, https://www.eeoc.gov/policy/docs/factemployment_procedures.html (last accessed: Sept. 25, 2019).
[54] *Id.*

reasoning, memory, speed and accuracy, or math and reading comprehension, where they disproportionately exclude African American and Latino candidates. Physical ability tests that measure strength and stamina as well as the ability to perform a particular task have also led to litigation where there is a disparate impact on women or certain ethnic groups in workplaces such as police and fire departments, or truck driving and warehouse jobs. These screens can also raise age and disability concerns as well.

Employers have increasingly utilized personality tests and integrity tests for pre-employment screening, which seek to assess an applicant's traits or temperament such as dependability, cooperativeness, conscientiousness, or aim to predict the likelihood that a person will engage in certain conduct such as theft or safety-related behaviors. Personality tests have generally not had a disparate impact based on race, gender or ethnicity, but they have raised concerns under the ADA and ADEA. This can include concerns of bias in determining traits important for the job or whether the testing is a prohibited pre-offer medical examination if the screens provide evidence that would lead to identifying a mental disorder or impairment. Even where a test is given after a job offer is made it can raise ADA concerns if the test screens out people with disabilities.

The Uniform Guidelines on Employee Selection Procedures ("Uniform Guidelines"), adopted by the EEOC and other federal agencies in 1978, provide employers with a set of principles to determine whether a test or selection procedure is non-discriminatory.[55] The guidance provides a framework for assessing the appropriate use of assessments that hinges on "validation:" the disputed selection method must be shown to be sufficiently related to or predictive of job performance. The Uniform Guidelines offer three criteria to evaluate the validity of a given test: criterion, content, and construct. Criterion-related validity provides evidence that the selection procedure or test is predictive or significantly correlated with job performance. Content validity provides the selection procedure or test includes similar tasks representative of the job. Construct validity provides the selection procedure or test identifies some identifiable characteristic deemed relevant to job performance.

Algorithmic selection models raise critical issues concerning how to apply this existing Title VII legal framework. These practices also raise technological challenges, and issues around transparency. In 2015, as the EEOC celebrated its 50th anniversary, it set out to better understand how work was changing and what government regulators needed to understand to adapt. The Commission held several meetings exploring how technology is changing who has an opportunity to work, looking at social media and online recruiting and big data and employment screens.[56]

[55] Equal Employment Opportunity Commission, Civil Service Commission, et al. *Uniform Guidelines On Employee Selection Procedures* Fed. R, 43(166):38290-38315 (1978).

[56] *See, e.g.,* Equal Employment Opportunity Commission, Big Data in the Workplace, Hearing Transcript Oct. 13, 2016, https://www.eeoc.gov/eeoc/meetings/10-13-16/transcript.cfm; EEOC at 50: Progress and Continuing Challenges in Eradicating Employment Discrimination, July 1, 2015, https://www.eeoc.gov/eeoc/meetings/7-1-15/index.cfm ; Social Media in the Workplace Examining Implications for Equal Employment Opportunity Law, March 12, 2014, https://www.eeoc.gov/eeoc/meetings/3-12-14/index.cfm.

What became clear was that the government did not have a window into what was actually happening with the rapid advances in predictive selection devices. The agency saw a need to engage with the tech community, employers and civil rights organizations to better understand what kinds of guidance could be most helpful in providing a framework for the use of algorithmic hiring that advances equity. The Commission created a task force that sought to answer these questions and explore concerns.

In November 2017, the EEOC hired its first Chief Data Officer and in May 2018 reorganized the agency's Office of Research, Information & Planning into the new Office of Enterprise Data and Analytics ("OEDA") to develop an enterprise-wide data analytics system. In December 2018, the EEOC held a series of "Data Dialogues" to introduce stakeholders to the agency's new OEDA.[57] Issues concerning AI-driven hiring screens are increasingly of interest to Congress, and last year, Senators Kamala Harris, Patty Murray, and Elizabeth Warren wrote to the EEOC requesting a report on the legality and potential issues with the use of facial analysis in pre-employment assessments.[58]

Important questions remain over how existing legal and policy frameworks should be applied or adapted to the use of algorithms in hiring. Algorithms, by their design, often predict applicant characteristics that correlate with job performance without considering the existence of a causal link between them. Many screening models attempt to demonstrate that algorithms are job related by assessing the personal characteristics associated with their best performing employees as evidence of criterion-related validity. But, often these models fail to provide a robust analysis of how the characteristics being predicted are required for job performance. Correlation by itself can be a misleading metric for assessing validity of algorithm since correlation does not demonstrate causation. Although algorithms have the potential to uncover job-related characteristics with strong predictive power, they could just as easily identify correlations arising from statistical noise or, more troubling, from previously undetected bias in the training data. In addition, algorithmic hiring models that continue to train—and therefore change—after deployment raise complex validation challenges.

Ultimately, employers bear responsibility under Title VII for any discrimination in their hiring practices, even if using a third-party vendor tool. Vague assertions about an algorithm's legal compliance or job-relatedness will not protect an employer from legal responsibility. Thus, employers and their counsel must understand how AI models work, including the validation studies supporting screening models and the research demonstrating the job-relatedness of the model. Employers should also understand what vendors actually do when they claim to have de-biased algorithms, and employers should have a process in place to monitor any discriminatory effect that might exist in or be caused by these tools as used in their own workplaces.

There is a tendency with new technology to want to move fast, to test new ideas and disrupt existing practices, but in the area of employment discrimination, these are highly consequential

---

[57] *See* EEOC Holds Data Dialogue Session with Agency Stakeholders, https://www.eeoc.gov/eeoc/newsroom/release/12-14-18.cfm

[58] Devin Coldewey, *Sen. Harris Tells Federal Agencies To Get Serious About Facial Recognition*, TechCrunch (Sept. 18, 2018), https://techcrunch.com/2018/09/18/sen-harris-tells-federal-agencies-to-get-serious-about-facial-recognition-risks/

decisions impacting people's economic opportunity. The widespread interest in using AI to make hiring decisions has brought to the forefront some important issues, including that many employers do not know what they are trying to achieve or what skills and abilities make someone successful on the job. AI alone will not inherently produce fairer employment outcomes: employers will need to invest in ensuring the criteria for selection is both fair and job-related.

Important questions remain about how best to strike a balance between the desire to innovate and create greater efficiencies in decision-making processes and the impact of these decisions on individuals and communities. To create a future with greater equity, AI-driven hiring screens must be designed to ensure that they are treating all people fairly with a commitment to meaningful transparency and accountability. This requires rigorous validation of screens and adequate documentation of decisions to ensure AI-powered models perform as intended and possess a sufficient business. By coming together across disciplines, technologists, business leaders, civil society, lawyers, and others have an important role to play in creating a comprehensive technical, policy, and legal framework for eliminating discrimination and bias in hiring algorithms. Collectively, these efforts can work to expand opportunity for those who have historically been underrepresented in our workplaces and industries.