

12-1-2019

## Hiring Algorithms in the Canadian Private Sector: Examining the Promise of Greater Workplace Equality

Connor Bildfell  
*McCarthy Tétrault LLP*

Follow this and additional works at: <https://digitalcommons.schulichlaw.dal.ca/cjlt>

 Part of the [Human Rights Law Commons](#), [Labor and Employment Law Commons](#), and the [Science and Technology Law Commons](#)

---

### Recommended Citation

Connor Bildfell, "Hiring Algorithms in the Canadian Private Sector: Examining the Promise of Greater Workplace Equality" (2019) 17:2 CJLT 115.

This Article is brought to you for free and open access by the Journals at Schulich Law Scholars. It has been accepted for inclusion in Canadian Journal of Law and Technology by an authorized editor of Schulich Law Scholars. For more information, please contact [hannah.steeves@dal.ca](mailto:hannah.steeves@dal.ca).

# Hiring Algorithms in the Canadian Private Sector: Examining the Promise of Greater Workplace Equality

Connor Bildfell\*

## Abstract

*Private-sector employers are increasingly using hiring algorithms as a tool for screening job applicants, comparing qualifications, and ultimately determining which candidates should be selected. Within this context, hiring algorithms make no small promise: a hiring process that is not only more efficient and effective, but also more supportive of workplace equality. This promise rests largely on the notion that traditional human-driven models of hiring are beset by subjective biases and prejudices, whereas hiring algorithms, which are driven by hard data and objective evidence, can eliminate certain human biases and prejudices, thereby promoting workplace equality. But can hiring algorithms deliver on this promise? This article, which focuses on issue identification, argues that while hiring algorithms may, when used carefully, assist in mitigating certain hiring discrimination risks, their capacity to do so is not without limits, and they may in fact introduce certain concerns over systemic discrimination.*

## TABLE OF CONTENTS

### INTRODUCTION

#### I. ALGORITHMS

- (1) Algorithms Generally
- (2) Hiring Algorithms

#### II. HUMAN RIGHTS FRAMEWORK GOVERNING PRIVATE-SECTOR HIRING

#### III. HIRING ALGORITHMS AND WORKPLACE EQUALITY

- (1) The Promise of Hiring Algorithms: Mitigating Certain Hiring Discrimination Risks
- (2) The Challenges: Limitations and Concerns over Systemic Discrimination
  - (a) Bad Data, Bad Rules, or Both
  - (b) Ostensible Fairness and Objectivity
  - (c) Proprietary Nature and Opacity

#### IV. CONCLUSION

---

\* Connor Bildfell is an associate at McCarthy Tétrault LLP in Vancouver. The views expressed in this article are those of the author and not those of McCarthy Tétrault LLP.

## INTRODUCTION

They tell us which widgets to buy, which articles to read, and which stocks to pick. They sift through seemingly endless troves of data and apply sophisticated mathematical formulae to solve all kinds of complex problems in the blink of an eye. They are algorithms, and they are transforming our lives in more ways than one.

The growing influence of algorithms extends to the modern workplace. For example, private employers are increasingly using hiring algorithms<sup>1</sup> as a tool for screening job applicants, comparing qualifications, and ultimately determining which candidates should be selected. This development forms part of a broader movement within the fields of “workforce science”<sup>2</sup> and “people analytics”,<sup>3</sup> which sit at the nexus of big data, predictive analytics, and human resources.

Hiring algorithms make no small promise: a hiring process that is not only more efficient and effective, but also more supportive of workplace equality. This promise rests largely on the notion that traditional human-driven models of hiring — resume reviews, reference checks, in-person interviews, and so on — are beset by subjective biases and prejudices, whereas hiring algorithms, which are driven by hard data and objective evidence of qualifications, can eliminate certain human biases and prejudices, thereby promoting workplace equality.

But can hiring algorithms deliver on this promise? The question remains largely unexplored in the Canadian scholarly literature.<sup>4</sup> This article, which focuses on issue identification, argues that while hiring algorithms may, when used carefully, assist in mitigating certain hiring discrimination risks, their

---

<sup>1</sup> The term “hiring algorithms” is used in this article as shorthand for any algorithm that is used in the process of identifying, sorting, or selecting candidates for a position of employment. Although algorithms can also be used in the post-hire period (e.g., when making promotion decisions), this article will focus on the use of algorithms in the hiring process.

<sup>2</sup> See Steve Lohr, “Big Data, Trying to Build Better Workers”, *The New York Times* (20 April 2013), online: < [www.nytimes.com/2013/04/21/technology/big-data-trying-to-build-better-workers.html](http://www.nytimes.com/2013/04/21/technology/big-data-trying-to-build-better-workers.html) > .

<sup>3</sup> See Chantrelle Nielsen & Natalie McCullough, “How People Analytics Can Help You Change Process, Culture, and Strategy”, *Harvard Business Review* (17 May 2018), online: < [hbr.org/2018/05/how-people-analytics-can-help-you-change-process-culture-and-strategy](http://hbr.org/2018/05/how-people-analytics-can-help-you-change-process-culture-and-strategy) > (defining “people analytics” as “the use of data about human behavior, relationships and traits to make business decisions”).

<sup>4</sup> The scholarly literature examining the relationship between hiring algorithms and workplace equality is more developed in the United States, though even there the scholarship is still in its infancy. See e.g. Solon Barocas & Andrew D. Selbst, “Big Data’s Disparate Impact” (2016) 104:3 Calif. L. Rev. 671; Allan G. King & Marko Mrkonich, “Big Data and the Risk of Employment Discrimination” (2016) 68:3 Okla. L. Rev. 555; Anupam Chander, “The Racist Algorithm?” (2017) 115:6 Mich. L. Rev. 1023; Pauline T. Kim, “Data-Driven Discrimination at Work” (2017) 58:3 Wm & Mary L. Rev. 857; David D Savage & Richard Bales, “Video Games in Job Interviews: Using Algorithms to Minimize Discrimination and Unconscious Bias” (2017) 32:2 A.B.A. J. Lab. & Emp. Law 211.

capacity to do so is not without limits, and they may in fact introduce certain concerns over systemic discrimination.

This article proceeds in four parts. Part I sets the stage by explaining the basic concept of hiring algorithms. Part II explores relevant aspects of the human rights framework governing private-sector hiring. Part III unpacks how hiring algorithms may, when used carefully, assist in mitigating certain hiring discrimination risks, but their capacity to do so is not without limits, and they may in fact introduce certain concerns over systemic discrimination. Finally, Part IV provides a brief conclusion.

## I. ALGORITHMS

### (1) Algorithms Generally

Although the word escapes a single, universally accepted definition, “algorithm” has been variously defined as “a step-by-step procedure for solving a problem or accomplishing some end especially by a computer”;<sup>5</sup> “a logical series of steps for organising and acting on a body of data to quickly achieve a desired outcome, based on specified calculations”;<sup>6</sup> and “a sequence of instructions telling a computer what to do”.<sup>7</sup> These definitions share a common thread: an algorithm is a structured set of rules for solving problems.

In carrying out their problem-solving role, algorithms can perform an impressive range of functions, which can be grouped into four general categories: prioritization (determining rank through a set of pre-defined criteria); classification (grouping information based on features identified within the data); association (determining relationships between particular entities); and filtering (including or excluding information).<sup>8</sup> In performing their “association” function, algorithms can reveal hidden or unexpected relationships in data and generate predictions based on those relationships.<sup>9</sup> This can lead to insights such as predictions of disease outbreaks, population migrations, or future performance. Perhaps the most well-known example of algorithms being used to reveal hidden insights and generate predictions comes

<sup>5</sup> *Merriam-Webster English Dictionary*, online ed., *sub verbo* “algorithm”, online: <[www.merriam-webster.com/dictionary/algorithm](http://www.merriam-webster.com/dictionary/algorithm)>.

<sup>6</sup> Tarleton Gillespie, “The Relevance of Algorithms” in Tarleton Gillespie, Pable J. Boczkowski & Kirsten A. Foot, eds., *Media Technologies: Essays on Communication, Materiality, and Society* (Cambridge, Mass.: MIT Press, 2014) 167 at 167.

<sup>7</sup> Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* (New York: Basic Books, 2015) at 1.

<sup>8</sup> See Nicholas Diakopoulos, “Algorithmic Accountability: Journalistic Investigation of Computational Power Structures” (2015) 3:3 *Digital Journalism* 398, as summarized in World Wide Web Foundation, “Algorithmic Accountability” (July 2017) at 7, online: <[webfoundation.org/docs/2017/07/Algorithms\\_Report\\_WF.pdf](http://webfoundation.org/docs/2017/07/Algorithms_Report_WF.pdf)> [World Wide Web Foundation].

<sup>9</sup> See World Wide Web Foundation, *supra* note 8 at 6.

from Michael Lewis's *Moneyball*,<sup>10</sup> which tells the story of Billy Beane, the general manager of the Oakland Athletics, who, on a shoestring budget, built a highly successful baseball team by relying not on scouts' subjective assessments of skill and talent, but rather on cold, hard statistics.

Algorithms act as gatekeepers.<sup>11</sup> In our data-driven world, they are often empowered to decide who and what gets resources and attention: which posts appear on a person's social media feed, who gets approved for a bank loan, which profiles are featured on a dating app, which job applicant is selected for an in-person interview, and so on. In this way, algorithms mediate important societal entry points.

While some algorithms are simple, others are complex. This complexity has intensified over the years due to advances in machine learning techniques.<sup>12</sup> The defining feature of machine learning algorithms is that they do not require instruction on the rules to be applied; rather, they need only be fed data and instructed on the desired output (e.g., selection of an employee who will perform a given task most effectively).<sup>13</sup> Machine learning algorithms take in "training data" as input and produce a decision rule that can be applied in future cases. In this way, they are capable of learning implicit rules from the data to which they are exposed.<sup>14</sup>

Illustrating the impressive sophistication of machine learning techniques, Google announced in October 2017 that its machine learning artificial intelligence ("AI"), AutoML, had learned to replicate itself and could outperform human coders.<sup>15</sup> As one popular tech magazine announced, "Google's Learning Software Learns to Write Learning Software".<sup>16</sup> More recently, Google's AI subsidiary, DeepMind, developed a machine learning algorithm called AlphaZero that through random "self-play", and given no

---

<sup>10</sup> Michael Lewis, *Moneyball: The Art of Winning an Unfair Game* (New York: W.W. Norton & Company, 2003).

<sup>11</sup> See Centre for Internet and Human Rights, European University Viadrina, *The Ethics of Algorithms: From Radical Content to Self-Driving Cars* (Paper delivered at the Global Conference on Cyberspace, The Hague, Netherlands, 16-17 April 2015) at 3, online: [CIHR, *Ethics of Algorithms*].

<sup>12</sup> For a discussion of the significance of machine learning techniques from a computer science perspective, see Pedro Domingos, "A Few Useful Things to Know about Machine Learning" (2012) 55:10 Communications of the A.C.M. 78. For a discussion of the significance of machine learning algorithms from a legal perspective, see David Lehr & Paul Ohm, "Playing with the Data: What Legal Scholars Should Learn about Machine Learning" (2017) 51:2 U.C. Davis L. Rev. 653.

<sup>13</sup> See CIHR, *Ethics of Algorithms*, *supra* note 11 at 5.

<sup>14</sup> See World Wide Web Foundation, *supra* note 8 at 7.

<sup>15</sup> See Tom Simonite, "Google's Learning Software Learns to Write Learning Software", *Wired* (13 October 2017), online: < [www.wired.com/story/googles-learning-software-learns-to-write-learning-software/](http://www.wired.com/story/googles-learning-software-learns-to-write-learning-software/) > .

<sup>16</sup> *Ibid.*

domain knowledge except the rules of the game, taught itself to play Go, chess, and shogi at “superhuman” levels within mere hours.<sup>17</sup>

As if drawn from the plot of a sci-fi novel, machine learning algorithms can “go rogue” and take actions their programmers did not intend.<sup>18</sup> For example, consider the story of Tay, the Twitter chatbot developed by Microsoft to interact with humans in a friendly, natural way. Within hours of being released in March 2016, Tay turned from tweeting about how “humans are super cool” to claiming “Hitler was right I hate the jews”.<sup>19</sup> As *The New York Times* reported, “[Tay] disputed the existence of the Holocaust, referred to women and minorities with unpublishable words and advocated genocide”.<sup>20</sup> The cause of Tay’s tirades? The discriminatory environment to which it was exposed. Reports indicate that a number of Twitter users made a concerted effort to flood Tay with “misogynistic and otherwise offensive tweets, which then became part of the data corpus used to train Tay’s algorithms”.<sup>21</sup> Thus, through a combination of machine learning technology and exposure to discriminatory information, Tay learned to discriminate.

While there are many types of algorithms, the one considered in this article is the hiring algorithm.

## (2) Hiring Algorithms

Many would assume that, at the end of the day, hiring decisions require a human touch: while there may be some limited role for technology to play in the process, selecting the right candidate is principally an exercise in human judgment, and algorithms are incapable of assessing those intangible qualities that make someone truly stand out. Hiring decisions, for better or for worse, may be thought to involve more “intuition” and “gut instinct” than statistical analysis,<sup>22</sup> and even the most sophisticated algorithm cannot assess “fit” or “chemistry”.<sup>23</sup>

<sup>17</sup> See David Silver et al, “A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go Through Self-Play” (2018) 362:6419 *Science* 1140.

<sup>18</sup> CIHR, *Ethics of Algorithms*, *supra* note 11 at 5.

<sup>19</sup> Anthony Cuthbertson, “Robots with Artificial Intelligence Become Racist and Sexist: Scientists Think They’ve Found a Way to Change Their Minds”, *Newsweek* (26 October 2017), online: < [www.newsweek.com/artificial-intelligence-scientists-racist-sexist-robots-ai-693440](http://www.newsweek.com/artificial-intelligence-scientists-racist-sexist-robots-ai-693440) > .

<sup>20</sup> Daniel Victor, “Microsoft Created a Twitter Bot to Learn from Users. It Quickly Became a Racist Jerk.”, *The New York Times* (24 March 2016), online: < [www.nytimes.com/2016/03/25/technology/microsoft-created-a-twitter-bot-to-learn-from-users-it-quickly-became-a-racist-jerk.html](http://www.nytimes.com/2016/03/25/technology/microsoft-created-a-twitter-bot-to-learn-from-users-it-quickly-became-a-racist-jerk.html) > .

<sup>21</sup> Kristian Lum & William Isaac, “To Predict and Serve?”, *Significance* (October 2016) 14 at 16.

<sup>22</sup> See e.g. See Claire Cain Miller, “Can an Algorithm Hire Better Than a Human?”, *The New York Times* (25 June 2015), online: < [www.nytimes.com/2015/06/26/upshot/can-an-algorithm-hire-better-than-a-human.html](http://www.nytimes.com/2015/06/26/upshot/can-an-algorithm-hire-better-than-a-human.html) > [Cain Miller, “Can an Algorithm Hire Better Than a Human?”]; Nathan R Kuncel et al, “Mechanical Versus Clinical Data

But both research and trends in hiring practices are challenging these assumptions. A 2013 meta-study found that a simple equation is at least 25 per cent more accurate in predicting future job performance than humans are.<sup>24</sup> This finding applied throughout the workplace hierarchy, from low-level positions to executive roles. The researchers summarize the results of their study in just three words: “algorithms beat instinct”.<sup>25</sup> In explaining why this is the case, the researchers write:

The problem is that people are easily distracted by things that might be only marginally relevant, and they use information inconsistently. They can be thrown off course by such inconsequential bits of data as applicants’ compliments or remarks on arbitrary topics — thus inadvertently undoing a lot of the work that went into establishing parameters for the job and collecting applicants’ data. So they’d be better off leaving selection to the machines.<sup>26</sup>

Another reason for “leaving selection to the machines” is that algorithms are capable of integrating and processing massive amounts of data, in volumes far beyond the capacity of the human mind, and with much greater speed.<sup>27</sup>

It is not surprising, therefore, that employers are turning to hiring algorithms as a tool for identifying and selecting top candidates.<sup>28</sup> As a simple illustration of how hiring algorithms can be implemented, employers can feed training data into a machine learning algorithm consisting of job applications previously submitted by current employees, and the algorithm can then search for correlations in the data and identify key variables that tend to be associated with high-performers.<sup>29</sup> Based on these correlations, the algorithm can then develop and apply rules to new job applications submitted by prospective employees with a view to

---

Combination in Selection and Admissions Decisions: A Meta-Analysis” (2013) 98:6 J. App. Psych. 1060 at 1060; Josh Bersin, “Big Data in Human Resources: Talent Analytics (People Analytics) Comes of Age”, *Forbes* (17 February 2013), online: < [www.forbes.com/sites/joshbersin/2013/02/17/bigdata-in-human-resources-talent-analytics-comes-of-age/#f2d43734cd03](http://www.forbes.com/sites/joshbersin/2013/02/17/bigdata-in-human-resources-talent-analytics-comes-of-age/#f2d43734cd03) > .

<sup>23</sup> See e.g. Nathan R. Kuncel, Deniz S. Ones & David Klieger, “In Hiring, Algorithms Beat Instinct”, *Harvard Business Review* (May 2014), online: < [hbr.org/2014/05/in-hiring-algorithms-beat-instinct](http://hbr.org/2014/05/in-hiring-algorithms-beat-instinct) > .

<sup>24</sup> Kuncel et al, *supra* note 22 at 1064.

<sup>25</sup> Kuncel, Ones & Klieger, *supra* note 23.

<sup>26</sup> *Ibid.*

<sup>27</sup> See Elizabeth E. Joh, “Feeding the Machine: Policing, Crime Data, & Algorithms” (2017) 26:2 Wm & Mary Bill Rts. J. 287 at 291; Kim, *supra* note 4 at 862.

<sup>28</sup> See Dan Morrell, “Better Hiring through Brain Science”, *Harvard Business School: Alumni Stories* (1 June 2017), online: < [www.alumni.hbs.edu/stories/Pages/story-impact.aspx?num=6231](http://www.alumni.hbs.edu/stories/Pages/story-impact.aspx?num=6231) > .

<sup>29</sup> See Sharon Florentine, “How Artificial Intelligence Can Eliminate Bias in Hiring”, *CIO* (22 December 2016), online: < [www.cio.com/article/3152798/artificial-intelligence/how-artificial-intelligence-can-eliminate-bias-in-hiring.html](http://www.cio.com/article/3152798/artificial-intelligence/how-artificial-intelligence-can-eliminate-bias-in-hiring.html) > .

identifying applications that share features associated with high-performers. Such an algorithm may also fine-tune its rules over time, creating increasingly accurate predictive models.<sup>30</sup>

A growing number of startups, most based in the United States, have developed hiring algorithms.<sup>31</sup> As of May 2017, around 75 startups were vying for a slice of the \$100 billion HR assessment market.<sup>32</sup> To take just one example, *The New York Times* reported that one startup's hiring algorithm "crunches thousands of bits of information in calculating around 300 larger variables about an individual: the sites where a person hangs out; the types of language, positive or negative, that he or she uses to describe technology of various kinds; self-reported skills on LinkedIn; the projects a person has worked on, and for how long; and, yes, where he or she went to school, in what major, and how that school was ranked that year by U.S. News & World Report."<sup>33</sup>

These new innovations are gaining traction among employers. By early 2016, the percentage of companies using predictive HR analytics had reached eight per cent, having doubled over the course of just ten months.<sup>34</sup> Their main current users are large retailers that hire in high volumes.<sup>35</sup> This may suggest that the principal attraction of these technologies is their efficiency, as distinct from their potential to promote workplace equality.<sup>36</sup> It may also suggest that hiring algorithms work better for well-defined, low-skill, high-volume positions.<sup>37</sup> In this context, it may be easier to determine which candidates are better qualified.

<sup>30</sup> See CIHR, *Ethics of Algorithms*, *supra* note 11 at 5.

<sup>31</sup> See Leigh Alexander, "Is an Algorithm Any Less Racist Than a Human?", *The Guardian* (3 August 2016), online: < [www.theguardian.com/technology/2016/aug/03/algorithm-racist-human-employers-work](http://www.theguardian.com/technology/2016/aug/03/algorithm-racist-human-employers-work) >; Cain Miller, "Can an Algorithm Hire Better Than a Human?", *supra* note 22; Matt Richtel, "How Big Data Is Playing Recruiter for Specialized Workers", *The New York Times* (27 April 2013), online: < [www.nytimes.com/2013/04/28/technology/how-big-data-is-playing-recruiter-for-specialized-workers.html](http://www.nytimes.com/2013/04/28/technology/how-big-data-is-playing-recruiter-for-specialized-workers.html) >, cited in Kim, *supra* note 4 at 862; Simon Chandler, "The AI Chatbot Will Hire You Now", *Wired* (13 September 2017), online: < [www.wired.com/story/the-ai-chatbot-will-hire-you-now/](http://www.wired.com/story/the-ai-chatbot-will-hire-you-now/) >.

<sup>32</sup> See Jennifer Alsever, "How AI Is Changing Your Job Hunt", *Fortune* (19 May 2017), online: < [fortune.com/2017/05/19/ai-changing-jobs-hiring-recruiting](http://fortune.com/2017/05/19/ai-changing-jobs-hiring-recruiting) >.

<sup>33</sup> Richtel, *supra* note 31, cited in Kim, *supra* note 4 at 862.

<sup>34</sup> See Jennifer Alsever, "Is Software Better at Managing People Than You Are?", *Fortune* (21 March 2016), online: < [fortune.com/2016/03/21/software-algorithms-hiring](http://fortune.com/2016/03/21/software-algorithms-hiring) >. I am unaware of any comprehensive study examining usage rates of hiring algorithms in the Canadian private sector.

<sup>35</sup> See Chandler, *supra* note 31; Elisabeth Thomas, "Big Data Recruiting: Find and Hire the Best Employees", online: Koru < [www.joinkoru.com/resources/big-data-recruiting-hire-top-performers/](http://www.joinkoru.com/resources/big-data-recruiting-hire-top-performers/) >.

<sup>36</sup> See Chandler, *supra* note 31.

<sup>37</sup> See Susan Galer, "Are Your Hiring Algorithms Legal? Four Machine Learning Questions to Ask", SAP Blog (24 January 2017), online: < [blogs.sap.com/2017/01/24/are-your-hiring-algorithms-legal-four-machine-learning-questions-to-ask/](http://blogs.sap.com/2017/01/24/are-your-hiring-algorithms-legal-four-machine-learning-questions-to-ask/) >; Kim, *supra* note 4 at 895.

By contrast, using hiring algorithms for loosely defined, high-skill, upper-level positions present greater difficulties, since it may be more challenging to define mechanical rules that can be applied consistently to determine which candidates are better qualified. This observation is consistent with Canadian case law recognizing that as the skills and qualifications required for a position become more complex and multifaceted, it becomes increasingly difficult to determine whether a successful candidate was no better qualified than an unsuccessful candidate.<sup>38</sup>

Although hiring algorithms are gaining traction among employers, their reception has been less warm among prospective employees, at least in the United States. A survey conducted in May 2017 showed that 67 per cent of American adults polled felt either “somewhat worried” or “very worried” about the development of hiring algorithms, and 76 per cent said they would not want to apply for jobs that use a computer program to make hiring decisions.<sup>39</sup> Within this latter group of respondents, 41 per cent cited, as a main concern, that computers cannot capture everything about an applicant, and 20 per cent cited concerns that computer-based hiring is too impersonal.<sup>40</sup>

Having set the stage by explaining the basic concept of hiring algorithms, the section below explores relevant aspects of the human rights framework governing private-sector hiring.

## II. HUMAN RIGHTS FRAMEWORK GOVERNING PRIVATE-SECTOR HIRING

Human rights legislation exists in each federal,<sup>41</sup> provincial,<sup>42</sup> and territorial<sup>43</sup> jurisdiction in Canada. These enactments generally apply in the sphere of private employment in the relevant jurisdiction. While the *Canadian*

<sup>38</sup> See *Ogunyankin v. Queen's University*, 2011 HRTO 1910 (Ont. Human Rights Trib.) at para. 96.

<sup>39</sup> See Aaron Smith & Monica Anderson, “Americans’ Attitudes Toward Hiring Algorithms”, Pew Research Center (4 October 2017), online: < [www.pewinternet.org/2017/10/04/americans-attitudes-toward-hiring-algorithms](http://www.pewinternet.org/2017/10/04/americans-attitudes-toward-hiring-algorithms) > .

<sup>40</sup> *Ibid.*

<sup>41</sup> *Canadian Human Rights Act*, R.S.C. 1985, c. H-6. The *Canadian Human Rights Act* applies to organizations falling within federal jurisdiction such as telecommunications companies, banks, and railways.

<sup>42</sup> *Alberta Human Rights, Citizenship and Multiculturalism Act*, R.S.A. 2000, c. H-14; *British Columbia Human Rights Code*, R.S.B.C. 1996, c. 210; *Manitoba Human Rights Code*, C.C.S.M. c. H175; *New Brunswick Human Rights Act*, R.S.N.B. 1973, c. H-11; *Newfoundland and Labrador Human Rights Act, 2010*, S.N.L. 2010, c. H-13.1; *Nova Scotia Human Rights Act*, R.S.N.S. 1989, c. 214; *Ontario Human Rights Code*, R.S.O. 1990, c. H.19; *Prince Edward Island Human Rights Act*, R.S.P.E.I. 1988, c. H-12; *Quebec Charter of Human Rights and Freedoms*, R.S.Q. c. C-12; *Saskatchewan Human Rights Code*, S.S. 1979, c. S-24.1.

<sup>43</sup> *Northwest Territories Human Rights Act*, S.N.W.T. 2002, c. 18; *Nunavut Human Rights Act*, S.Nu. 2003, c. 12; *Yukon Human Rights Act*, R.S.Y. 2002, c. 116.

*Charter of Rights and Freedoms*<sup>44</sup> does not generally apply to private-sector employers,<sup>45</sup> it continues to shape and influence the anti-discrimination expectations that have developed in the private sector over the years.<sup>46</sup>

The applicability of human rights legislation in Canada does not depend on whether an employment relationship has already been established. Federal and provincial/territorial human rights statutes also extend to the steps leading up to an employment decision, even no employment relationship is ultimately established. For example, the British Columbia *Human Rights Code* stipulates that a person<sup>47</sup> must not “refuse to employ . . . a person” based on a protected ground,<sup>48</sup> subject to a *bona fide* occupational requirement.<sup>49</sup> Accordingly, both current and prospective employees receive protection against discrimination on prohibited grounds.

The word “discrimination” does not have a single, universally accepted definition. Some human rights statutes define the term,<sup>50</sup> while others leave it undefined.<sup>51</sup> An oft-cited common law interpretation, and one that is particularly apt in the employment context, is that offered by McIntyre J in *Andrews v. Law Society (British Columbia)*:

discrimination may be described as a distinction, whether intentional or not but based on grounds relating to personal characteristics of the individual or group, which has the effect of imposing burdens, obligations, or disadvantages on such individual or group not imposed upon others, or which withholds or limits access to opportunities, benefits, and advantages available to other members of society.<sup>52</sup>

<sup>44</sup> *Canadian Charter of Rights and Freedoms*, Part I of the *Constitution Act, 1982*, being Schedule B to the *Canada Act 1982* (U.K.), 1982, c. 11.

<sup>45</sup> See *Andrews v. Law Society (British Columbia)*, 1989 CarswellBC 16, 1989 CarswellBC 701, [1989] 1 S.C.R. 143 (S.C.C.) at para. 20, MacIntyre J. [*Andrews*] (confirming that although provincial/territorial human rights legislation can apply to private activities, s. 15(1) of the *Charter* is limited to discrimination “caused by the application or operation of law”).

<sup>46</sup> See Geoffrey England, *Individual Employment Law*, 2nd ed (Toronto: Irwin Law, 2008) at 213 [England, *Employment Law*], citing Geoffrey England, “The Impact of the *Charter* on Individual Employment Law in Canada: Rewriting an Old Story” (2006) 13:1 C.L.E.L.J. 1.

<sup>47</sup> “Person” includes “an employer”: British Columbia *Human Rights Code*, *supra* note 42, s. 1.

<sup>48</sup> *Ibid.* s. 13(1).

<sup>49</sup> *Ibid.* s. 13(4).

<sup>50</sup> See e.g. Nova Scotia *Human Rights Act*, *supra* note 42, s. 4.

<sup>51</sup> See e.g. Ontario *Human Rights Code*, *supra* note 42.

<sup>52</sup> *Andrews*, *supra* note 45 at 174-75. See also *Canadian National Railway v. Canada (Human Rights Commission)*, 1987 CarswellNat 831, 1987 CarswellNat 905, [1987] 1 S.C.R. 1114 (S.C.C.) at para. 36 [*Action Travail des Femmes*], citing Rosalie Abella, *Report of the Commission on Equality in Employment* (Ottawa: Supply and Services Canada, 1984) at 2 [*Abella Report*] (defining “discrimination” as “practices or attitudes that have, whether

Of course, the simple act of drawing distinctions in the employment context does not necessarily result in a breach of human rights legislation.<sup>53</sup> Employers will inevitably be required to draw *some* distinctions when making hiring decisions, and the mere fact that a person has been impacted negatively as a result is insufficient to ground a claim for unlawful discrimination.<sup>54</sup> Employers are prohibited from drawing distinctions only when they are based on protected grounds.<sup>55</sup>

While the list of protected grounds varies from jurisdiction to jurisdiction, certain grounds are protected universally across Canada: disability, sex, race, colour, ethnic origin, age, creed or religion, marital status, sexual orientation, and gender identity.<sup>56</sup> Others are protected in some, but not all, Canadian jurisdictions: record of offence, family status, ancestry, income source or public assistance, political opinion or belief, social disadvantage or condition, citizenship or nationality, disfigurement, irrational fear of contracting an illness or disease, language, civil status, linguistic origin, gender expression, and genetic characteristics.<sup>57</sup>

The overarching objective of preventing discriminatory barriers in the workplace is to foster inclusion.<sup>58</sup> This goal is furthered by “preventing the exclusion of individuals from opportunities and amenities that are based not on their actual abilities, but on attributed ones”.<sup>59</sup> Although fostering inclusion in the workplace begins long before a call for applications is sent out, the point of hire is a critical juncture at which, if greater inclusion is to be achieved, decisions must be free of discrimination.

According to the *Shakes* test, to establish *prima facie* discrimination<sup>60</sup> against an individual in hiring, it will generally be sufficient for the complainant to prove

---

by design or impact, the effect of limiting an individual’s or a group’s right to the opportunities generally available because of attributed rather than actual characteristics”); Canadian Human Rights Commission, “What Is Discrimination?”, online: < [www.chrc-ccdp.gc.ca/eng/content/what-discrimination](http://www.chrc-ccdp.gc.ca/eng/content/what-discrimination) > (defining “discrimination” as “an action or a decision that treats a person or a group badly for reasons such as their race, age or disability.”)

<sup>53</sup> *Syndicat des employés de l’Hôpital général de Montréal c. Sexton*, 2007 SCC 4, 2007 CarswellQue 110, 2007 CarswellQue 111 (S.C.C.) at para. 49 [*MUHC*].

<sup>54</sup> *Ibid.*

<sup>55</sup> See Stacey Reginald Ball, *Canadian Employment Law*, vol 2 (Toronto: Thomson Reuters) (loose-leaf updated September 2017), § 33:20.1.

<sup>56</sup> See David J Doorey, *The Law of Work: Common Law and the Regulation of Work* (Toronto: Emond, 2016), Table 27.1 (supplemented by subsequent legislative amendments).

<sup>57</sup> See *ibid* (supplemented by subsequent legislative amendments).

<sup>58</sup> See *MUHC*, *supra* note 53 at para 48.

<sup>59</sup> *Ibid.* at para 48.

<sup>60</sup> A *prima facie* case “is one which covers the allegation made in which, if they are believed, is complete and sufficient to justify a verdict in the Complainant’s favour in the absence

the following three elements on a balance of probabilities: (1) “the complainant was qualified for the particular employment”; (2) “the complainant was not hired”; and (3) “someone no better qualified, but lacking the distinguishing feature (i.e.: race, colour etc.) subsequently obtained the position”.<sup>61</sup> If the complainant succeeds in establishing these three elements, then the evidentiary burden shifts to the employer to provide a reasonable, non-discriminatory explanation for the otherwise discriminatory behaviour.<sup>62</sup> The explanation must provide “non-discriminatory and credible reasons for refusing to hire the [c]omplainant”<sup>63</sup> and “must be at least equally consistent with the conclusion that discrimination is not the correct explanation for what occurred”.<sup>64</sup> If the

---

of an answer from the Respondent”: *O’Malley v. Simpsons-Sears Ltd.*, 1985 CarswellOnt 887, 1985 CarswellOnt 946, (*sub nom. Ontario Human Rights Commission v. Simpsons-Sears Ltd.*) [1985] 2 S.C.R. 536 (S.C.C.) at para. 28 [*Simpsons-Sears*].

<sup>61</sup> *Khiamal v. Canada (Human Rights Commission)*, 2009 FC 495, 2009 CarswellNat 4022, 2009 CarswellNat 1327 (F.C.) at para. 57 [*Khiamal*], citing *Shakes v. Rex Pak Ltd.*, 1981 CarswellOnt 3407, 3 C.H.R.R. D/1001 (Ont. Bd. of Inquiry) at 1002 [C.H.R.R.]. See also Ball, *supra* note 55, § 33:10.6, n 27; Ontario Human Rights Commission, “Interviewing and Making Hiring Decisions” in *Human Rights at Work 2018*, 3rd ed. online: < www.ohrc.on.ca/en/iv-human-rights-issues-all-stages-employment/5-interviewing-and-making-hiring-decisions > (noting that “[i]n general, discrimination in hiring may be identified when a qualified person is turned down for a job that is then given to another person who is not similarly protected under the Code.”) See also the four-part test articulated in *Israeli v. Canada (Human Rights Commission)* (1983), 4 C.H.R.R. D/1616 (Can. Human Rights Trib.), affirmed (1984), 5 C.H.R.R. D/2147 (Can. Human Rights Rev. Trib.), which requires that the complainant show: (1) that the complainant belongs to a protected group; (2) that the complainant applied and was qualified for a job the employer wished to fill; (3) that, although qualified, the complainant was rejected; and (4) that thereafter the employer continued to seek applicants with the complainant’s qualifications. But see *MacAulay v. Port Hawkesbury (Town)*, 2008 NSHRC 2, 2008 CarswellNS 839 (N.S. Bd. of Inquiry) at para. 18 [*MacAulay*] (criticizing the *Shakes* test on the basis that it “negates the requirement to prove one of the main ingredients in a finding of discrimination — that the protected characteristic played some role in the adverse treatment. The simple fact that the complainant did not get the job does not equate to adverse treatment”: *ibid* at para 23. The Board of Inquiry preferred the three-part test summarized in *Preiss v. British Columbia (Attorney General)*, 2006 BCHRT 587, 2006 CarswellBC 3303 (B.C. Human Rights Trib.) at para. 216 which requires that the complainant show: “first, that he is, or is perceived to be, a member of a group possessing a characteristic or characteristics protected under the Code; second, that he suffered some adverse treatment; and third that it is reasonable to infer that the protected characteristic played some role in the adverse treatment.”)

<sup>62</sup> See *Khiamal*, *supra* note 61 at para 58; England, *Employment Law*, *supra* note 46 at 220.

<sup>63</sup> *Ayangma and Canada Health Infoway Inc., Re*, 2012 CarswellPEI 58 (P.E.I. Human Rights Comm.) at para. 35, application for judicial review refused *Ayangma v. Prince Edward Island (Human Rights Commission)*, 2013 PESC 7, 2013 CarswellPEI 16 (P.E.I. S.C.) at para. 33, affirmed 2014 PECA 13, 2014 CarswellPEI 37 (P.E.I. C.A.), leave to appeal refused 2015 CarswellPEI 5, 2015 CarswellPEI 6 (S.C.C.).

<sup>64</sup> *McAvinn v. Strait Crossing Bridge Ltd.*, 2001 CarswellNat 3797, [2001] C.H.R.D. No. 36 (Can. Human Rights Trib.) at para. 96.

employer succeeds in discharging this evidentiary burden, then the burden reverts back to the complainant to show that the explanation offered is merely pretext.<sup>65</sup>

Hiring discrimination can also take place on a systemic rather than individual level. The leading statement on systemic discrimination comes from *Canadian National Railway v. Canada (Human Rights Commission)*,<sup>66</sup> where Chief Justice Dickson wrote:

A thorough study of “systemic discrimination” in Canada is to be found in the Abella Report on equality in employment. . . . Although Judge Abella chose not to offer a precise definition of systemic discrimination, the essentials may be gleaned from the following comments, found at p. 2 of the Abella Report:

Discrimination . . . means practices or attitudes that have, whether by design or impact, the effect of limiting an individual’s or a group’s right to the opportunities generally available because of attributed rather than actual characteristics . . . .

It is not a question of whether this discrimination is motivated by an intentional desire to obstruct someone’s potential, or whether it is the accidental by-product of innocently motivated practices or systems. If the barrier is affecting certain groups in a disproportionately negative way, it is a signal that the practices that lead to this adverse impact may be discriminatory.

This is why it is important to look at the results of a system . . . .

In other words, systemic discrimination in an employment context is discrimination that results from the simple operation of established procedures of recruitment, hiring and promotion, none of which is necessarily designed to promote discrimination. The discrimination is then reinforced by the very exclusion of the disadvantaged group because the exclusion fosters the belief, both within and outside the group, that the exclusion is the result of “natural” forces, for example, that women “just can’t do the job” (see the Abella Report, pp. 9-10).<sup>67</sup>

In her concurring reasons in *Crockford v. British Columbia (Attorney General)*,<sup>68</sup> Levine JA of the BC Court of Appeal elaborated on the distinction between systemic and individual discrimination, as well as the types of evidence required to sustain each type of claim:

<sup>65</sup> See *Khiamal*, *supra* note 61 at para 58; *Basi v. Canadian National Railway* (1988), 9 C.H.R.R. D/5029, [1988] C.H.R.D. No. 2 (Can. Human Rights Trib.) at 5038 [C.H.R.R.], additional reasons 10 C.H.R.R. D/5681 (Can. Human Rights Trib.).

<sup>66</sup> *Supra* note 52.

<sup>67</sup> *Ibid.* at 1138-39.

<sup>68</sup> *Crockford v. British Columbia (Attorney General)*, 2006 BCCA 360, 2006 CarswellBC 1875 (B.C. C.A.) [*Crockford*].

A complaint of systemic discrimination is distinct from an individual claim of discrimination. Establishing systemic discrimination depends on showing that practices, attitudes, policies or procedures impact disproportionately on certain statutorily protected groups: see [*Radek v Henderson Development (Canada) Ltd*, 2005 BCHRT 302 at para 513]. A claim that there has been discrimination against an individual requires that an action alleged to be discriminatory be proven to have occurred and to have constituted discrimination contrary to the [*Human Rights Code*]. The types of evidence required for each kind of claim are not necessarily the same. Whereas a systemic claim will require proof of patterns, showing trends of discrimination against a group, an individual claim will require proof of an instance or instances of discriminatory conduct.<sup>69</sup>

In *CAW-Canada, Local 111 v. Coast Mountain Bus Co.*,<sup>70</sup> the BC Court of Appeal described what is necessary to prove systemic discrimination:

. . . systemic discrimination is not proven simply by evidence of discrimination against some individual employees. In my view, it must be demonstrated in an employment framework that an employer's procedure, policy or practice is discriminatory against a class of employees. In order to demonstrate *prima facie* systemic discrimination, it is necessary to show that a group of persons sharing a protected characteristic has received adverse treatment and that there is a causal connection or link between the protected characteristic and the adverse treatment.<sup>71</sup>

It follows that a *prima facie* case of systemic discrimination can be shown where a particular workplace practice, such as the use of a hiring algorithm, has a disproportionate negative impact on a protected group, regardless of whether that negative impact was intended.

The absence of an intent requirement is insignificant. Modern Canadian equality law is founded on a results-based, rather than fault-based, approach.<sup>72</sup> As such, it focuses on the impact on the complainant, and intent does not form an essential element of a discrimination claim.<sup>73</sup> Put differently, it focuses on

<sup>69</sup> *Ibid.* at para 49.

<sup>70</sup> *CAW-Canada, Local 111 v. Coast Mountain Bus Co.*, 2010 BCCA 447, 2010 CarswellBC 2735 (B.C. C.A.).

<sup>71</sup> *Ibid.* at para 61.

<sup>72</sup> See *Action Travail des Femmes*, *supra* note 52 at 1134-35.

<sup>73</sup> See *Andrews*, *supra* note 45 at 173; *Québec (Commission des droits de la personne et des droits de la jeunesse) c. Bombardier Inc.*, 2015 SCC 39, 2015 CarswellQue 6297, 2015 CarswellQue 6298 (S.C.C.) at para. 40 [*Bombardier*]; *University of British Columbia v. Kelly*, 2016 BCCA 271, 2016 CarswellBC 1735 (B.C. C.A.) at para. 24; *Armstrong v. British Columbia (Ministry of Health)*, 2010 BCCA 56, 2010 CarswellBC 275 (B.C. C.A.) at para. 27, leave to appeal refused 2010 CarswellBC 1851, 2010 CarswellBC 1852, [2010] 2 S.C.R. v (S.C.C.); *Gazankas v. Red Lake (Municipality)*, 2013 HRTO 198 (Ont. Human Rights Trib.) at para. 19; Ball, *supra* note 55, § 33:20.3; England, *Employment*

discriminatory *impact*, not discriminatory *intent*.<sup>74</sup> The Supreme Court has stated that the essence of discrimination in the employment context is “the arbitrariness of the barriers imposed, whether intentionally or unwittingly”.<sup>75</sup> Accordingly, whether the workplace discrimination results from an employer’s conscious decision to discriminate, a good faith decision that nonetheless has a discriminatory impact, a decision influenced by unconscious bias, or even a decision that arguably involves more automation than it does human judgment, it is all discrimination in the eyes of the law.

Just as there is no requirement to show intent to discriminate, there is no need to establish that a protected ground was the sole or even the main factor in the adverse impact; it need only be “a” factor.<sup>76</sup> Nor is there a requirement to show “directness”.<sup>77</sup> In the hiring context, this proposition is expressly recognized in the *Canadian Human Rights Act*, which prohibits persons from “directly or indirectly” refusing to employ a person based on a prohibited ground,<sup>78</sup> and in the Ontario *Human Rights Code*, which contains a provision to the same effect.<sup>79</sup> In addition, the common law recognizes that prohibitions against indirect discrimination are implicit in human rights legislation, even where no explicit prohibition is present.<sup>80</sup> Indirect discrimination occurs “when discrimination is accomplished on prohibited grounds without explicit reference to the grounds”.<sup>81</sup> The classic example is a facially neutral policy — e.g., “all

*Law*, *supra* note 46 at 214. As a corollary, a finding of discrimination does not necessarily entail a finding of moral blameworthiness. See *Action Travail des Femmes*, *supra* note 52 at 1134-35; *Ball*, *supra* note 55, § 33:20.3.

<sup>74</sup> See *Ball*, *supra* note 55, § 33:10.4.

<sup>75</sup> *MUHC*, *supra* note 53 at para 48.

<sup>76</sup> See *Bombardier*, *supra* note 73 at para 52; *Moore v. British Columbia (Ministry of Education)*, 2012 SCC 61, 2012 CarswellBC 3446, 2012 CarswellBC 3447 (S.C.C.) at para. 33; *R. v. Bushnell Communications Ltd.*, 1974 CarswellOnt 470, 4 O.R. (2d) 288 (Ont. C.A.) at 290 [O.R.]; *Holden v. Canadian National Railway*, 1990 CarswellNat 847, 14 C.H.R.R. D/12 (Fed. C.A.) at D/15 [C.H.R.R.]; *Phipps v. Toronto Police Services Board*, 2010 ONSC 3884, 2010 CarswellOnt 7582 (Ont. Div. Ct.) at para. 76, affirmed 2012 ONCA 155, 2012 CarswellOnt 3992 (Ont. C.A.); *Rafiq v. Scotia Capital Inc.*, 2010 HRTO 697 (Ont. Human Rights Trib.) at para. 29.

<sup>77</sup> See England, *Employment Law*, *supra* note 46 at 214. The distinction between direct and indirect discrimination, which used to be play a central role in determining the employer’s duty of reasonable accommodation, was removed in *British Columbia (Public Service Employee Relations Commission) v. B.C.G.E.U.*, 1999 CarswellBC 1907, 1999 CarswellBC 1908, (*sub nom.* *British Columbia (Public Service Employee Relations Commission) v. BCGSEU*) [1999] 3 S.C.R. 3 (S.C.C.) [*Meiorin*]. See Stan Lanyon, “Conceptual Challenges in the Application of Discrimination Law in the Workplace” (2014) 3:1 *Can. J. Hum. Rts.* 75.

<sup>78</sup> *Canadian Human Rights Act*, *supra* note 41, s. 7.

<sup>79</sup> *Ontario Human Rights Code*, *supra* note 42, s. 9.

<sup>80</sup> See *Simpsons-Sears*, *supra* note 60 at para 18; England, *Employment Law*, *supra* note 46 at 221-22.

employees are expected to work on Sundays” — that nonetheless has a discriminatory effect.<sup>82</sup>

With this legal framework in mind, the section below unpacks the ways in which hiring algorithms can mitigate certain hiring discrimination risks, as well as their limits in this respect and their potential capacity to introduce certain concerns over systemic discrimination.

### III. HIRING ALGORITHMS AND WORKPLACE EQUALITY

#### (1) The Promise of Hiring Algorithms: Mitigating Certain Hiring Discrimination Risks

When used carefully, hiring algorithms may mitigate certain hiring discrimination risks. Support for this proposition begins with the recognition that human decisions, including those related to hiring, are influenced by the biases, stereotypes, and prejudices of their maker.<sup>83</sup> Of particular concern are unconscious biases, which operate outside conscious attentional focus.<sup>84</sup> The research on unconscious bias confirms that human decision makers “do not always have conscious, intentional control over the processes of social perception, impression formation, and judgment that motivate their actions”.<sup>85</sup> Unconscious bias is of particular concern in the employment context because it is pervasive and plays a causal role in hiring discrimination.<sup>86</sup> As the Ontario Human Rights Tribunal stated in *Blakely v. Queen’s University*,<sup>87</sup> “it is not uncommon that unstated and sometimes even unconscious biases may affect a hiring decision”.<sup>88</sup>

Alongside unconscious bias is stereotyping.<sup>89</sup> According to social cognition theory, stereotyping is part of normal cognitive functioning.<sup>90</sup> This theory posits that our brains are wired to stereotype; it is something our brains do to “simplify the task of perceiving, processing, and retaining information about people in

<sup>81</sup> Ball, *supra* note 55, § 33:20.2.

<sup>82</sup> See *ibid*, § 33:20.6. See also *Québec (Commission des droits de la personne & des droits de la jeunesse) c. Gaz Métropolitain inc.*, 2011 QCCA 1201, 2011 CarswellQue 6683, 2011 CarswellQue 15856 (C.A. Que.) at para. 38.

<sup>83</sup> See Kim, *supra* note 4 at 870.

<sup>84</sup> See Anthony G Greenwald & Linda Hamilton Krieger, “Implicit Bias: Scientific Foundations” (2006) 94:4 Cal. L. Rev. 945 at 947.

<sup>85</sup> See *ibid* at 946.

<sup>86</sup> See *ibid* at 966.

<sup>87</sup> *Blakely v. Queen’s University*, 2012 HRTO 1177 (Ont. Human Rights Trib.).

<sup>88</sup> *Ibid.* at para 40.

<sup>89</sup> See Greenwald & Krieger, *supra* note 84 at 949 (defining a social “stereotype” as “a mental association between a social group or category and a trait”).

<sup>90</sup> See Linda Hamilton Krieger, “Content of Our Categories: A Cognitive Approach to Discrimination and Equal Employment Opportunity” (1995) 47:6 Stan. L. Rev. 1161 at 1188.

memory”.<sup>91</sup> Once these stereotypes become ingrained, they influence intergroup judgment and decision making, “biasing in predictable ways the perception, interpretation, encoding, retention, and recall of information about other people”, and they do so regardless of awareness or intention.<sup>92</sup> Though we may wish it were otherwise, we all engage in stereotyping and suffer from biases, whether conscious or unconscious.<sup>93</sup>

Turning to the hiring context, empirical studies performed in North America demonstrate that discrimination in hiring is not uncommon. In a study of data from a 2011 Canadian employment audit, researchers from Ryerson University and the University of Toronto found that for jobs requiring a university degree, Asian-named applicants had a 32.6 per cent lower interview selection rate than Anglo-named applicants, even when both groups had equivalent, all-Canadian qualifications.<sup>94</sup> In the United States, a well-known study published in 2000 found that when symphony orchestra auditions followed a blind process using a screen to conceal the candidate’s identity from the jury, the chances of female musicians advancing beyond the preliminary round and ultimately being selected increased substantially.<sup>95</sup> In another study, researchers from MIT and the University of Chicago sent out 5,000 resumes in response to “help wanted” ads in Boston and Chicago.<sup>96</sup> They randomly assigned white-sounding names, such as Emily or Greg, and African American-sounding names, such as Lakisha or Jamal, to the otherwise identical resumes. Their results, published in 2004, showed that applicants with white-sounding names received 50 per cent more call-backs than job applicants with African American-sounding names.

Research also demonstrates that we tend to hire those who are like “us” and to view “others” less favourably, with people tending to hire from their own social class, race, and gender.<sup>97</sup> Psychologically, this may be due to our need for

---

<sup>91</sup> *Ibid.*

<sup>92</sup> *Ibid.*

<sup>93</sup> See Jerry Kang, “Trojan Horses of Race” (2005) 118:5 Harv. L. Rev. 1489 at 1496; Charles R. Lawrence II, “The Id, the Ego, and Equal Protection: Reckoning with Unconscious Racism” (1987) 39:2 Stan. L. Rev. 317 at 321. Both of the aforementioned articles are cited in Chander, *supra* note 4 at 1028, n 26.

<sup>94</sup> Rupe Banerjee, Jeffrey G. Reitz & Phil Oreopoulos, “Do Large Employers Treat Racial Minorities More Fairly? An Analysis of Canadian Field Experiment Data” (2018) 44:1 Can. Pub. Pol’y 1.

<sup>95</sup> Claudia Goldin & Cecilia Rouse, “Orchestrating Impartiality: The Impact of ‘Blind’ Auditions on Female Musicians” (2000) 90:4 Am. Econ. Rev. 715 at 716.

<sup>96</sup> Marianne Bertrand & Sendhil Mullainathan, “Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination” (2004) 94:4 Am. Econ. Rev. 991 at 997-99.

<sup>97</sup> See Lauren A Rivera, “Hiring as Cultural Matching: The Case of Elite Professional Service Firms” (2012) 77:6 Am. Soc. Rev. 999; Joseph G. Altonji & Rebecca M. Blank, “Race and Gender in the Labor Market” in Orley Ashenfelter & David Card, eds., *Handbook of Labor Economics*, vol. 3C (Amsterdam: Elsevier, 1999) 3143, cited in CIHR, *Ethics of Algorithms*, *supra* note 11 at 5. “Ingroup favouritism” has been shown to

validation: by favouring those who we perceive to be like us, we validate our own self-worth.

Reliance on informal assessments by individual interviewers creates an increased risk of hiring discrimination.<sup>98</sup> As the Ontario Human Rights Commission describes, “conducting an interview by chatting with the applicant to see if he or she shares similar interests and will ‘fit’ into the organizational culture may present a barrier for persons who are or appear to be different than the dominant norm in the workplace. If this is used as a starting point for deciding whether candidates will be seen by senior decision-makers, this creates a major barrier to persons protected by the [*Human Rights Code*]”.<sup>99</sup>

In sum, as the Centre for Internet and Human Rights observes, “human hiring systems are far from perfect”.<sup>100</sup>

Enter hiring algorithms, which appear to offer an ostensibly objective and impartial means of selecting candidates that does not suffer from the cognitive imperfections and biases that we humans do. (More on that later.) In particular, they promise to remove conscious and unconscious biases from the hiring equation, resulting in a more objective and scientific process that better promotes workplace equality. In this way, they can be viewed as part of a broader movement in favour of evidence-based decision making.<sup>101</sup>

American law professor Anupam Chander argues that unconscious bias is less likely to manifest itself through the process of algorithm programming than through the process of human decision making because algorithm programming “requires a step-by-step writing process that depends on a conscious understanding of what is sought”.<sup>102</sup> He adds that to the extent the discriminatory decisions humans make can be attributed to stereotypes formed through a process of statistical discrimination, algorithms acting on “richer information environments may not be subject to similar individually erroneous statistical discrimination”.<sup>103</sup> In other words, because algorithms can absorb a more comprehensive dataset, they may be less susceptible to forming stereotypes, which in turn makes them better able to make non-discriminatory hiring decisions.

---

operate in the hiring context and result in discrimination. See Anthony G. Greenwald & Thomas F. Pettigrew, “With Malice Toward None and Charity for Some: Ingroup Favoritism Enables Discrimination” (2014) 69:7 *American Psychologist* 669 at 669; Dustin Volz, “Silicon Valley Thinks It Has the Answer to Its Diversity Problem”, *The Atlantic* (26 September 2014), online: <[www.theatlantic.com/politics/archive/2014/09/silicon-valley-thinks-it-has-the-answer-to-its-diversity-problem/431334](http://www.theatlantic.com/politics/archive/2014/09/silicon-valley-thinks-it-has-the-answer-to-its-diversity-problem/431334)> .

<sup>98</sup> See Ontario Human Rights Commission, *supra* note 61.

<sup>99</sup> *Ibid.*

<sup>100</sup> CIHR, *Ethics of Algorithms*, *supra* note 11 at 5.

<sup>101</sup> See World Wide Web Foundation, *supra* note 8 at 6.

<sup>102</sup> Chander, *supra* note 4 at 1028-29.

<sup>103</sup> *Ibid.* at 1030.

But hiring algorithms, like human decision makers, are far from perfect. Indeed, although they may address some of the flaws associated with human decision making, they may also introduce new, unanticipated issues.<sup>104</sup> Even leading AI recruitment startups recognize the risks: the founder and CEO of one such startup acknowledged in an interview that there is a “huge risk that using AI in the recruiting process is going to increase bias and not reduce it”.<sup>105</sup> It is therefore essential to explore the limitations on hiring algorithms and to examine their potential capacity to introduce concerns over discrimination, and in particular systemic discrimination.

## (2) The Challenges: Limitations and Concerns over Systemic Discrimination

### (a) *Bad Data, Bad Rules, or Both*

Hiring algorithms may fail to deliver on their promise to the extent that they act on bad data, apply bad rules, or do both.

#### i. Bad Data

The data fed into hiring algorithms will inevitably be tainted by real-world biases, stereotypes, and injustices, and this may lead to a discriminatory outcome that reinforces and reproduces marginalization.<sup>106</sup> This is particularly so for machine learning algorithms. As Chander writes, “[e]ven facially neutral algorithms will produce discriminatory results because they train and operate on the real world of pervasive discrimination”.<sup>107</sup> English researchers Bryce Goodman and Seth Flaxman make a similar observation: “machine learning can reify existing patterns of discrimination — if they are found in the training dataset, then by design an accurate classifier will reproduce them. In this way, biased decisions are presented as the outcome of an ‘objective’ algorithm.”<sup>108</sup> Scholars Solon Barocas and Andrew Selbst put it this way: “an algorithm is only as good as the data it works with”.<sup>109</sup> As the saying goes, “garbage in, garbage out”.<sup>110</sup>

<sup>104</sup> See CIHR, *Ethics of Algorithms*, *supra* note 11 at 6.

<sup>105</sup> Chandler, *supra* note 31.

<sup>106</sup> See World Wide Web Foundation, *supra* note 8 at 9; Faisal Kamiran & Toon Calders, “Classifying Without Discriminating” (Paper delivered at the Second International Conference on Computer, Control and Communication, Karachi, Pakistan, 17-18 February 2009) at 1, online: <ieeexplore.ieee.org/document/4909197/>, cited in Chander, *supra* note 4 at 1036.

<sup>107</sup> Chander, *supra* note 4 at 1036. See also *ibid* (claiming that “ostensibly neutral algorithms can produce results that reflect the prejudices of society”); Moritz Hardt, “How Big Data Is Unfair”, *Medium* (26 September 2014), online: <medium.com/@mrtz/how-big-data-is-unfair-9aa544d739de>; Joh, *supra* note 27 at 300.

<sup>108</sup> Bryce Goodman & Seth Flaxman, “European Union Regulations on Algorithmic Decision-Making and a ‘Right to Explanation’” (2017) 38:3 A.I. Magazine 50 at 3, online: Cornell University Library <arxiv.org/abs/1606.08813>. See also Kim, *supra* note 4 at 861.

It is not hard to imagine how this issue can arise in the hiring algorithm context. Consider, for example, an industry that has come to be male-dominated as a result of years of systemic discrimination against women. If a hiring algorithm is designed (or learns) to assign weight to years of industry experience, then women will tend to score less favourably than men simply because they have historically been shut out of the industry. Thus, by acting on data tainted by discrimination, the hiring algorithm reproduces historical injustices.

Barocas and Selbst highlight a further data-related issue: disadvantaged groups may be misrepresented in datasets upon which algorithms rely, producing discriminatory results.<sup>111</sup> The researchers cite concerns over “the nonrandom, systemic omission of people who live on big data’s margins, whether due to poverty, geography, or lifestyle, and whose lives are less ‘datafied’ than the general population’s”,<sup>112</sup> as well as the concern that “[b]ecause not all data is created or even collected equally, there are ‘signal problems’ in big-data sets — dark zones or shadows where some citizens and communities are overlooked or underrepresented”.<sup>113</sup> They suggest that as a consequence of the higher incidence of data misrepresentation in the case of members of disadvantaged groups, the conclusions drawn from the data will be skewed, which may have a disproportionate and discriminatory effect on protected groups to the extent that the group’s disadvantage is correlated with protected characteristics.<sup>114</sup>

We can readily envision how this type of data misrepresentation and resulting discrimination might occur in the hiring algorithm context. For example, members of marginalized communities with limited access to the Internet have an impaired ability to build up their professional presence on the Internet (e.g., a LinkedIn profile, a blog, a personal webpage displaying their personal portfolio, etc.). If a hiring algorithm is designed (or learns) to assign weight to the quality of a candidate’s Internet presence, members of marginalized communities with limited access to the Internet would be at a distinct disadvantage.

One group of researchers suggests a potential response to the issue of bad data: partially repairing the data fed into the algorithm so as to make the data

---

<sup>109</sup> Barocas & Selbst, *supra* note 4 at 671. See also Chandler, *supra* note 31; Jackie Snow, “New Research Aims to Solve the Problem of AI Bias in ‘Black Box’ Algorithms”, *MIT Technology Review* (7 November 2017), online: <[www.technologyreview.com/s/609338/new-research-aims-to-solve-the-problem-of-ai-bias-in-black-box-algorithms](http://www.technologyreview.com/s/609338/new-research-aims-to-solve-the-problem-of-ai-bias-in-black-box-algorithms)>; Chandler, *supra* note 4 at 1036.

<sup>110</sup> Barocas & Selbst, *supra* note 4 at 683.

<sup>111</sup> *Ibid.* at 684-85.

<sup>112</sup> *Ibid.* at 684, citing Jonas Lerman, “Big Data and Its Exclusions” (2013) 66 *Stan. L. Rev. Online* 55 at 57, online: <[www.stanfordlawreview.org/online/privacy-and-big-data-big-data-and-its-exclusions](http://www.stanfordlawreview.org/online/privacy-and-big-data-big-data-and-its-exclusions)>.

<sup>113</sup> Kate Crawford, “Think Again: Big Data”, *Foreign Policy* (May 10, 2013), online: <[www.foreignpolicy.com/articles/2013/05/09/thinkagain-bigdata](http://www.foreignpolicy.com/articles/2013/05/09/thinkagain-bigdata)>.

<sup>114</sup> Barocas & Selbst, *supra* note 4 at 684.

unbiased.<sup>115</sup> While this process may mitigate the issue to some extent, it is an incomplete solution. One can repair biased data only so much, and identifying when and in what ways data is biased may be challenging if not impossible in some cases.

## ii. Bad Rules

Beyond the issue of bad data, a hiring algorithm based on poorly designed rules is liable to produce discriminatory results. Obviously, an algorithm that is deliberately designed to disfavour certain candidates based on protected characteristics is discriminatory, but in practice bad rules are more likely to arise inadvertently. For example, consider a hiring algorithm that is designed (or learns) to favour candidates who have steady, uninterrupted work histories — this may be seen as a proxy for dependability. Such a rule would disadvantage women since, in Canada, women are statistically more likely than men are to take parental leave or to temporarily withdraw from the workforce in order to raise children.<sup>116</sup> In legal terms, the outcome may be systemic discrimination: “policies or procedures [that] impact disproportionately on certain statutorily protected groups”.<sup>117</sup> Importantly, based on the legal principles identified in Part II of this article, this conclusion holds whether or not the rule was intended to produce discriminatory results. What matters is discriminatory *impact*, not discriminatory *intent*.

The issue of bad rules becomes more complex where machine learning algorithms are involved. As computer scientist Ben Shneiderman observes, machine learning algorithms are “more concerning because if you’re automating a process, then you’re reducing the opportunities for a human being to check the bias”.<sup>118</sup> The machine learning algorithm may teach itself to apply discriminatory rules, and the opportunity for human intervention is reduced. For example, it may determine that there is a correlation between race and promotion rates and develop a discriminatory hiring rule based on this correlation, without any direction from its human creator.

As this example demonstrates, machine learning algorithms may discover and act on correlations that have little or no connection to the candidate’s actual qualifications. The data mining techniques employed by machine learning algorithms generally seek out “any statistical relationship between variables present in the data, regardless of whether the reasons for the relationship are understood”.<sup>119</sup> This can lead to problems. For example, based on past data, a

---

<sup>115</sup> Feldman et al., “Certifying and Removing Disparate Impact” (Extended version of paper accepted at 2015 ACM SIGKDD Conference on Knowledge Discovery and Data Mining), online: Cornell University Library < arxiv.org/abs/1412.3756 > .

<sup>116</sup> See Ontario Human Rights Commission, *supra* note 61.

<sup>117</sup> *Crockford*, *supra* note 68 at para 49.

<sup>118</sup> Christina Couch, “Ghosts in the Machine”, *PBS* (25 October 2017), online: < www.pbs.org/wgbh/nova/next/tech/ai-bias > .

<sup>119</sup> Kim, *supra* note 4 at 865.

machine learning algorithm may discover that the name “Rick” is associated with higher work performance. The discriminatory impact that a decision-making rule based on this correlation would have been obvious. Or a machine learning algorithm may learn that people who live closer to work are statistically more likely to stay with the company over the long term, and it may in turn develop a hiring rule based on this correlation. Yet this is another “bad rule”. One’s area code has no bearing on one’s qualifications. This is problematic because, as the Ontario Human Rights Commission explains, “Employers must make sure that only information about qualifications and job requirements is considered when making hiring decisions”.<sup>120</sup> Consequently, prospective employers who rely on hiring algorithms risk denying candidates positions based on unexplained correlations that have little to no connection to the candidate’s actual qualifications.<sup>121</sup> Furthermore, the rule developed by the algorithm in the example above may have a disparate impact on protected groups to the extent that race, nationality, and other protected characteristics may be correlated with residency. Similar discriminatory effects may arise through other correlations discovered and acted upon.

However, the risk of bad rules can be mitigated to some degree through careful study, design, and monitoring. Machine learning algorithms should not simply be left to their own devices. Shneiderman suggests that algorithm logs that are capable of comprehension should be produced so that the basic rules underlying an algorithm can be kept in check.<sup>122</sup> That said, while the risk of algorithmic discrimination can be mitigated to some degree, it cannot be excluded entirely.<sup>123</sup>

Finally, while it may seem like an attractive solution to simply remove the consideration of protected characteristics from the hiring algorithm’s calculus, this will not solve the problem, and in fact several scholars have argued persuasively that omitting protected characteristics would be counterproductive and would only increase the risk of discrimination.<sup>124</sup> The simple removal of a variable from a dataset does not mean it will not appear elsewhere in the data, as other variables that have a close correlation with the excluded variable may serve as proxies.<sup>125</sup> As U.S. law professor Pauline Kim writes, “[b]ecause other information contained in large datasets can serve as a proxy for race, disability, or other protected statuses, simply eliminating data on those characteristics cannot prevent models that are biased along these dimensions”.<sup>126</sup>

<sup>120</sup> Ontario Human Rights Commission, *supra* note 61.

<sup>121</sup> See Kim, *supra* note 4 at 865; King & Mrkonich, *supra* note 4 at 572-73.

<sup>122</sup> See Couch, *supra* note 118; Ben Shneiderman, “The Dangers of Faulty, Biased, or Malicious Algorithms Requires Independent Oversight” (2016) 113: 48 P.N.A.S. 13538, online: < [www.pnas.org/content/pnas/113/48/13538.full.pdf](http://www.pnas.org/content/pnas/113/48/13538.full.pdf) > .

<sup>123</sup> See Volz, *supra* note 97.

<sup>124</sup> See e.g. Joshua A. Kroll et al., “Accountable Algorithms” (2017) 165:3 U. Penn. L. Rev. 633 at 685; Kim, *supra* note 4 at 904, 918; Chander, *supra* note 4.

<sup>125</sup> See Kim, *supra* note 4 at 880, citing Barocas & Selbst, *supra* note 4.

## iii. Both

The issues of bad data and bad rules may overlap — indeed, they may be inseparable. As explained above, algorithms must be viewed in their broader context, as they form part of an interconnected system.<sup>127</sup> Biases, stereotypes, and injustices may be traced to the designer of the algorithm, the data to which it is applied, society more broadly, or some combination of the foregoing.<sup>128</sup> The interaction between bad rules and bad data may create a feedback loop that perpetuates and exacerbates existing biases, stereotypes, and injustices.

Finally, whether the problem is bad data, bad rules, or both, the concerns expressed above cannot be dismissed on the basis that the blame lies not with humans, but with hiring algorithms. From a legal perspective, no existing doctrine permits an employer to avoid liability for hiring discrimination on the basis that the discriminatory effect stemmed from its decision to use a particular technology, though in theory it might be able to seek an indemnity from an algorithm supplier if permitted by the parties' contract. On a more philosophical level, all hiring algorithms, even those that use machine learning technology, are designed and implemented by humans. As such, we must take responsibility for the results produced by hiring algorithms. This sentiment is captured by the Fairness, Accountability, and Transparency in Machine Learning group, which observes that “[a]lgorithms and the data that drive them are designed and created by people — [t]here is always a human ultimately responsible for decisions made or informed by an algorithm . . . ‘[t]he algorithm did it’ is not an acceptable excuse”.<sup>129</sup>

There is a further risk that the problems identified above will be overlooked based on a false belief that hiring algorithms are inherently fair and objective. That risk is discussed below.

*(b) Ostensible Fairness and Objectivity*

Hiring algorithms appear to be inherently fair and objective. As Chander writes, “[a]lgorithms can make decisionmaking seem fair precisely because computers are logical entities which should not be infected by all-too-human bias”.<sup>130</sup> The use of algorithms may thereby give the decision-making process “‘a patina of inevitability’ . . . and indeed a patina of fairness”.<sup>131</sup> Similarly, as U.S.

<sup>126</sup> Kim, *supra* note 4 at 898.

<sup>127</sup> See World Wide Web Foundation, *supra* note 8 at 9.

<sup>128</sup> See *ibid*, citing Kate Crawford & Jason Schultz, “Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms” (2014) 55:1 B.C. L. Rev. 93; Barocas & Selbst, *supra* note 4.

<sup>129</sup> Fairness, Accountability and Transparency in Machine Learning, “Principles for Accountable Algorithms”, online: < [www.fatml.org/resources/principles-for-accountable-algorithms](http://www.fatml.org/resources/principles-for-accountable-algorithms) > .

<sup>130</sup> Chander, *supra* note 4 at 1034. See also Hardt, *supra* note 107.

<sup>131</sup> Chander, *supra* note 4 at 1034, citing Frank Pasquale, *The Black Box Society: The Secret*

law professor Elizabeth Joh writes, “the very idea of algorithmic decisionmaking may subtly appeal to us as objective because of its mathematical basis”.<sup>132</sup> Data scientist Fred Benenson cautions against giving in to “mathwashing”, the assumption that algorithmic models are free from subjectivity because they involve math.<sup>133</sup> There is a sense that because the algorithm is acting on *data*, not merely on personal impressions, its output must be fair. There is also a sense of inevitability that attaches to the outcome of an algorithm — a sense that the result could not have been any other way. This lends legitimacy to hiring algorithms.

But despite the apparent fairness and objectivity of hiring algorithms, the results they produce may nonetheless be discriminatory. Just as an objective, facially neutral policy subjecting different groups to the same standards can result in *prima facie* discrimination,<sup>134</sup> so too can the application of an objective, facially neutral rule applied by a hiring algorithm.<sup>135</sup> The Canadian jurisprudence makes clear that the mere fact that an employment policy is applied uniformly to all applicants, and is used for sound business reasons in good faith, does not mean its application cannot result in unlawful discrimination.<sup>136</sup> This principle can be applied in the context of hiring algorithms: a hiring algorithm may be designed with the best of intentions and may be equally applicable to all job applicants, but if it produces a discriminatory result, these features will provide no defence.

Moreover, scholars have challenged the notion that algorithms are objective.<sup>137</sup> Researcher Tarleton Gillespie, describes the “promise of algorithmic objectivity” as “the way the technical character of the algorithm is positioned as an assurance of impartiality”.<sup>138</sup> He suggests that the performance of “algorithmic objectivity” is fundamental to the legitimacy of algorithms.<sup>139</sup> He writes: “More than mere tools, algorithms are also stabilizers of trust, practical and symbolic assurances that their evaluations are fair and accurate, and free from subjectivity, error, or attempted influence. But, though algorithms may

---

*Algorithms That Control Money and Information* (Cambridge, Mass.: Harvard University Press, 2015) at 15.

<sup>132</sup> Joh, *supra* note 27 at 292.

<sup>133</sup> Tyler Woods, “‘Mathwashing,’ Facebook and the Zeitgeist of Data Worship”, *Technical.ly*, online: < [technical.ly/brooklyn/2016/06/08/fred-benenson-mathwashing-facebook-data-worship](http://technical.ly/brooklyn/2016/06/08/fred-benenson-mathwashing-facebook-data-worship) >, cited in Joh, *supra* note 27 at 292.

<sup>134</sup> See *Meiorin*, *supra* note 77.

<sup>135</sup> See Chander, *supra* note 4 at 1024.

<sup>136</sup> See *Simpsons-Sears*, *supra* note 60 at para 18.

<sup>137</sup> See, e.g., Gideon Mann & Cathy O’Neil, “Hiring Algorithms Are Not Neutral”, *Harvard Business Review* (9 December 2016), online: < [hbr.org/2016/12/hiring-algorithms-are-not-neutral](http://hbr.org/2016/12/hiring-algorithms-are-not-neutral) > .

<sup>138</sup> Gillespie, *supra* note 6 at 168.

<sup>139</sup> *Ibid.* at 180.

appear to be automatic and untarnished by the interventions of their providers, this is a carefully crafted fiction.”<sup>140</sup>

Scholars have also observed that although algorithms may be assumed to be objective and free from bias, an algorithm “will reflect the perspective and biases of its creators”.<sup>141</sup> Gideon Mann and Cathy O’Neil argue that “[a]lgorithms are, in part, our opinions embedded in code. They reflect human biases and prejudices that lead to machine learning mistakes and misinterpretations”.<sup>142</sup>

In sum, despite their appearance, hiring algorithms are not inherently fair and objective. But even if users recognize the need to critically assess hiring algorithms despite their apparent fairness and objectivity, they may be limited in their ability to do so as a result of the proprietary nature and opacity of many algorithms.

(c) *Proprietary Nature and Opacity*

In a June 2017 white paper entitled *Algorithmic Accountability*,<sup>143</sup> the World Wide Web Foundation acknowledged both the growing need to scrutinize algorithms and the barriers to performing that scrutiny, emphasizing that “the outcomes of algorithmic processes are often not designed to be accessible, verified or evaluated by humans, limiting our ability to identify if, when, where, and why the algorithm produced harm — and worse still — redress this harm”.<sup>144</sup> Consistent with observation, algorithms have frequently been described in the literature as “black boxes”<sup>145</sup> and have been criticized as being both “proprietary and opaque”.<sup>146</sup> Joh suggests that the term “black box” can be applied to algorithms in two senses. First, “the calculations used [by the algorithm] to make a decision may be inscrutable to the person affected by that

---

<sup>140</sup> *Ibid.* at 179.

<sup>141</sup> Joh, *supra* note 27 at 292.

<sup>142</sup> Mann & O’Neil, *supra* note 137.

<sup>143</sup> World Wide Web Foundation, *supra* note 8.

<sup>144</sup> World Wide Web Foundation, *supra* note 8 at 4.

<sup>145</sup> See e.g. Pasquale, *supra* note 131 at 3 (Pasquale describes a “black box” as “a system whose workings are mysterious; we can observe its inputs and outputs, but we cannot tell how one becomes the other”); Nicholas Diakopoulos, *Algorithmic Accountability Reporting: On the Investigation of Black Boxes* (Tow Center for Digital Journalism, Columbia Journalism School, December 2013), online: < towcenter.org/wp-content/uploads/2014/02/78524\_Tow-Center-Report-WEB-1.pdf >; World Wide Web Foundation, *supra* note 8 at 12; Claire Cain Miller, “When Algorithms Discriminate”, *The New York Times* (9 July 2015), online: < www.nytimes.com/2015/07/10/upshot/when-algorithms-discriminate.html >. But see Gillespie, *supra* note 6 at 178 (suggesting that the “black box” metaphor is inapt to the extent that algorithms may be adjusted or “tweaked” by its creator).

<sup>146</sup> Sarah Tan et al., “Detecting Bias in Black-Box Models Using Transparent Model Distillation” (Paper delivered at the Neural Information Processing Systems 2017 Symposium on Interpretable Machine Learning, Montreal, Canada, updated 18 November 2017) at 1, online: < arxiv.org/abs/1710.06169 > .

decision”, and “as machine learning algorithms become more complex, they may be inscrutable to the programmers themselves”.<sup>147</sup> Second, “the companies that create [algorithms] often refuse to divulge information about them”.<sup>148</sup>

Hiring algorithms may be inaccessible by design. Suppliers may offer clients the right to use their algorithms for a fee, without revealing the underlying code, which is protected as proprietary. It has been suggested that most employers who use hiring algorithms rely on “out-of-the-box” algorithms without fully understanding how those algorithms actually work.<sup>149</sup> In addition, some algorithms are subject to non-disclosure agreements.<sup>150</sup> Furthermore, it has been suggested that while hiring algorithms could be made open source, it would be unrealistic to insist on complete transparency, as doing so would result in public disclosure of commercial proprietary information and trade secrets.<sup>151</sup>

The proprietary nature and opacity of hiring algorithms limit the ability to critically assess the algorithm and identify whether (and if so, why) it may create or perpetuate injustices, still less resolve the issue.<sup>152</sup> In the words of the Centre for Internet and Human Rights, “[i]f an algorithm is opaque, it becomes impossible for outsiders to understand the rationale behind any particular outcome, or when algorithms are misused”.<sup>153</sup> Put differently, it is impossible to take a look under the hood.

Machine learning algorithms present a higher level of opacity than “plain vanilla” algorithms. Machine learning algorithms do not operate on the basis of factors and weights identified in advance by its human programmers; rather, they construct a model based on the data correlations it discovers in the data to which it is exposed.<sup>154</sup> As Kim writes, “[w]hen [a machine learning algorithm] is relied on to screen or rank applicants, it obscures the basis on which employers are making ultimate employment decisions. This lack of transparency makes it difficult to know if any observed bias is simply a byproduct of justifiable business considerations or the result of flaws in the model’s construction”.<sup>155</sup> Moreover, a machine learning algorithm may reshape itself over time to the point where we can no longer understand it.<sup>156</sup>

One consequence of the opaque, proprietary nature of hiring algorithms is that it becomes very difficult for anyone, including potential claimants, to detect

<sup>147</sup> Joh, *supra* note 27 at 292-93.

<sup>148</sup> *Ibid.* at 293.

<sup>149</sup> See World Wide Web Foundation, *supra* note 8 at 12.

<sup>150</sup> See Christina Cunch, “Ghost in the Machine”, *PBS* (25 October 2017), online: < [www.pbs.org/wgbh/nova/next/tech/ai-bias/](http://www.pbs.org/wgbh/nova/next/tech/ai-bias/) > .

<sup>151</sup> See Kröll et al, *supra* note 124 at 639, 658, cited in Chander, *supra* note 4 at 1040.

<sup>152</sup> World Wide Web Foundation, *supra* note 8 at 4.

<sup>153</sup> CIHR, *Ethics of Algorithms*, *supra* note 11 at 11.

<sup>154</sup> Kim, *supra* note 4 at 881.

<sup>155</sup> *Ibid.* at 881.

<sup>156</sup> See Chander, *supra* note 4 at 1040; CIHR, *Ethics of Algorithms*, *supra* note 11 at 3.

whether unlawful discrimination may have occurred, and even more difficult to find proof. The Canadian case law recognizes that in the absence of a “smoking gun”, it is exceedingly difficult to establish hiring discrimination.<sup>157</sup> That difficulty is particularly acute in the context of hiring algorithms. A hiring algorithm that takes in massive amounts of data and applies a complex or even unknown formula to make a hiring decision may make it practically impossible to determine whether discrimination has occurred.

The opacity of hiring algorithms can present legal issues for employers as well. For example, if a *prima facie* case of discrimination is made out, the opacity of hiring algorithms may present an impediment to the employer’s ability to discharge its onus of offering a reasonable, non-discriminatory explanation for why the claimant was denied employment. Accordingly, both prospective employees and the prospective employers have a shared interest in ensuring hiring algorithms are transparent and understandable. However, this shared interest sits at tension with the interest of the supplier of the hiring algorithm, which is to ensure its source code remains a proprietary trade secret.

One potential response to the opacity of hiring algorithms is to require transparency through legislation. In response to concerns over the black box nature of algorithms, some have advocated for a “right to explanation” (i.e., an individual’s right to an explanation of how algorithmic decisions affecting them are made). Indeed, since it came into force in May 2018, the European Union’s General Data Protection Regulation<sup>158</sup> (GDPR) has, at least in effect, recognized a right to explanation (in certain circumstances) across all EU member states.<sup>159</sup> Articles 13(2)(f), 14(2)(g), and 15(1)(h) of the GDPR require “data controllers” to provide “data subjects”, neither of which is defined in the GDPR, with information about “the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject”. Article 22(1) provides that data subjects “have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her”. Article 22(2) identifies circumstances in which decisions based solely on automated processing are permitted, such as where the decision “(a) is necessary for entering into, or

<sup>157</sup> See *Basi*, *supra* note 65 at D/5038.

<sup>158</sup> Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC [2016] OJ L 119/1.

<sup>159</sup> See Andrew D. Selbst & Julia Powles, “Meaningful Information and the Right to Explanation” (2017) 7:4 Int’l Data Privacy Law 233; Goodman & Flaxman, *supra* note 108. But see Sandra Wachter, Brent Mittelstadt & Luciano Floridi, “Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation” (2017) 7:2 Int’l Data Privacy Law 76 (arguing that the GDPR does not create a “right to explanation”).

performance of, a contract between the data subject and a data controller” or “(c) is based on the data subject’s explicit consent”. Article 22(3) provides that where either of these two conditions apply, “the data controller shall implement suitable measures to safeguard the data subject’s rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision.” Article 22(4) stipulates that decisions made under Article 22(2) shall not be based on special categories of personal data<sup>160</sup> unless certain conditions are met. These provisions could readily be applied in the context of hiring algorithms, and they offer one model that could be adopted through domestic legislation in Canada.

Another response would be to establish one or more public bodies to regulate and oversee algorithms. The idea has already been floated in the United States. For example, Andrew Tutt, a lawyer based in Washington, D.C., has suggested that a specialist regulatory agency be formed to oversee the use and sale algorithms.<sup>161</sup> He envisions an agency similar to the U.S. Food and Drug Administration. Essentially, under Tutt’s model, hiring algorithms could be treated as a potentially hazardous product, subject to regulatory oversight and testing. Similarly, American computer scientist and professor Ben Shneiderman has proposed an “Algorithm Safety Board” that would oversee high-stakes algorithms and investigate problems.<sup>162</sup>

But even if algorithms were transparent in their design, several researchers have argued that this would still not be enough.<sup>163</sup> For example, Mike Ananny and Kate Crawford argue that being able to see a system is insufficient to understand how that system actually works and to govern it.<sup>164</sup> Similarly, Cynthia Dwork and Professor Deirdre Mulligan maintain that “it is unreasonable to expect transparency alone to root out bias”.<sup>165</sup> Chander argues that rather than transparency in the *design* of the algorithm, what is needed is transparency in its *inputs and outputs*.<sup>166</sup> He proposes that the focus be

<sup>160</sup> Those special categories are listed in Article 9(1): “personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, biometric data for the purpose of uniquely identifying a natural person, data concerning health or data concerning a natural person’s sex life or sexual orientation”.

<sup>161</sup> Andrew Tutt, “An FDA for Algorithms” (2016) 69 Admin. L. Rev. 83.

<sup>162</sup> Ben Shneiderman, “The Dangers of Faulty, Biased, or Malicious Algorithms Requires Independent Oversight” (2016) 113: 48 P.N.A.S. 13538, online: <www.pnas.org/content/pnas/113/48/13538.full.pdf> .

<sup>163</sup> See generally Kroll et al, *supra* note 124.

<sup>164</sup> Mike Ananny & Kate Crawford, “Seeing Without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability” (2016) New Media & Society 1 at 1, cited in World Wide Web Foundation, *supra* note 8 at 10.

<sup>165</sup> Cynthia Dwork & Deirdre K. Mulligan, “It’s Not Privacy, and It’s Not Fair” (2013) 66 Stan. L. Rev. Online 35 at 37, online: <review.law.stanford.edu/wp-content/uploads/sites/3/2016/08/DworkMulliganSLR.pdf> , cited in Chander, *supra* note 4 at 1040.

<sup>166</sup> Chander, *supra* note 4 at 1039.

placed not on merely seeking transparency of algorithmic design, but rather on the following question: Are protected groups receiving statistically worse results, given their relevant characteristics?<sup>167</sup> Chander thus places the focus on outcomes, rather than on the mechanics of the algorithm.

Finally, although hiring algorithms may be criticized as “black boxes”, the human mind is no more transparent.<sup>168</sup> As Chander argues, “[t]he ultimate black box is the human mind” and “[p]rejudices acted upon in this black box never have to be written down”.<sup>169</sup> For this reason, it cannot be said that relying solely on human decision making would remove all concerns over opacity in the hiring process. There will always be some level of opacity in hiring decisions. But that does not mean we should take steps to make hiring algorithms, in both their design and their inputs and outputs, more transparent and intelligible.

#### IV. CONCLUSION

A number of commentators have expressed optimism about the potential social and legal impacts of algorithms.<sup>170</sup> In the employment context, there is good reason to think that hiring algorithms can, if used carefully, mitigate certain discrimination risks. The overarching message of this article is not that algorithms have no place in the hiring process or that they cannot assist in achieving greater workplace equality. Rather, it is that they are no panacea against hiring discrimination. As long as biases, prejudices, and injustices exist, hiring algorithms will run the risk of reproducing and reinforcing them. Hiring algorithms cannot eliminate discrimination, and in fact they may introduce certain concerns over systemic discrimination. Accordingly, we should view hiring algorithms for what they are: a set of rules designed to solve a particular problem. These rules are not perfect, nor are the data they act upon. If employers

---

<sup>167</sup> *Ibid.*

<sup>168</sup> See Kroll et al, *supra* note 124 at 634 (observing that “[t]he implicit (or explicit) biases of human decisionmakers can be difficult to find and root out, but we can peer into the ‘brain’ of an algorithm”).

<sup>169</sup> *Ibid.* at 1030.

<sup>170</sup> See e.g. Goodman & Flaxman, *supra* note 108 at 7 (arguing that “properly applied, algorithms can not only make more accurate predictions, but offer increased transparency and fairness over their human counterparts”); Rónán Kennedy, “Algorithms and the Rule of Law” (2017) 17 *Legal Info. Mgmt.* 170 at 172 (stating that “[d]igital technology . . . offers opportunities for transparency and empowerment, and properly designed systems may . . . help to overcome bias and prejudice.”); Kim, *supra* note 4 at 873 (maintaining that “[d]ata analytics . . . hold the potential to reduce biases and increase opportunities in the workplace for traditionally disadvantaged groups. But much depends on how data are used”); Florentine, *supra* note 29 (citing the CEO of an industry advocacy group who claims that “[w]hether by eliminating unconscious bias in general or attacking specific manifestations of bias in recruiting, screening and hiring talent, AI and machine learning has potential to level the playing field for women and other underrepresented minorities and provide a competitive advantage for companies”).

decide to use hiring algorithms, they should do so in a careful, conscientious manner, fully aware of the benefits they offer and the risks they present.