



Legal and Ethical Challenges for HR in Machine Learning

R. H. Hamilton¹ · H. Kristl Davison²

Accepted: 16 May 2021 / Published online: 31 May 2021

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

The technology of machine learning, a type of artificial intelligence, will enable organizations to analyze their use and deployment of human resources (HR) in new ways that ultimately will allow them to manage more effectively, but it will also present challenges for HR managers who are unprepared. In this paper we discuss some of the legal and ethical concerns in the HR context that accompany machine learning. Legal concerns include possible violations of both US employment discrimination laws and the provisions of the European General Data Protection Regulation, while ethical concerns for HR revolve around employee desires for privacy and justice. We assess that some data analysis activities that are legal nonetheless might not be appropriate in some cases and might be demotivating to employees, resulting in lowered performance or even counterproductive behaviors if HR mishandles the context. We conclude by offering guidelines for HR managers to assess the appropriateness of machine learning projects.

Keywords Machine learning · Employment discrimination · General Data Protection Regulation · Privacy · Organizational justice

Introduction: Data Possibilities and Dilemmas

In 2018 and 2019, Facebook faced significant scrutiny from many fronts. Multiple scandals plagued the social media giant, resulting in a plummet in the company's reputation. Part of the reason for Facebook's challenge was inherent in its business model: collecting data from individuals and selling it to third parties (LaForgia et al., 2018). The greater the specificity about user preferences that Facebook was able to convey to its clients, the more clients were willing to pay Facebook. As the specificity (and personal nature) of the data that Facebook collected became widely known, Facebook encountered criticism from

✉ R. H. Hamilton
bhamilton@bus.olemiss.edu

H. Kristl Davison
davisonhk@appstate.edu

¹ Department of Management, School of Business, The University of Mississippi, University, MS 38677, USA

² Walker College of Business, Appalachian State University, 4087 Peacock Hall, Boone, NC 28608, USA

legislative committees in the US and UK (Blumenthal, 2018; LaForgia et al., 2018; Reuters, 2019).

To the average human resources (HR) manager, Facebook's troubles may seem far away. However, companies will soon be inundated with new possibilities for data collection and direct observation and analysis of employee actions. Unprepared HR senior managers may face a series of legal and ethical dilemmas that, while not on the scale of the Facebook debacle, could complicate employee engagement and potentially even create employee backlash and counterproductive work behaviors (Sackett, 2002). On the other hand, approached appropriately, artificial intelligence techniques such as *machine learning* could help companies more effectively manage resources in a way that benefits both employees and the firm.

In this paper, we first review how machine learning will usher in new possibilities to assist HR. Next, we discuss some of the legal concerns that HR managers must address up front in the machine learning process, and we also review two key ethical concerns in the HR context. We conclude with practical implications and a series of proposed guidelines for senior HR managers.

The Ascendancy of Machine Learning

When the phrase “artificial intelligence” is used in the pop culture context, it typically refers to a future age in which autonomous vehicles and robots could take over large segments of society's tasks (e.g., as depicted in the movie derived from Isaac Asimov's science fiction tale, *I, Robot*). While in the medium-term robots and autonomous vehicles may indeed prove to be game-changers with accompanying opportunities and challenges, at present the most significant concern for HR managers from artificial intelligence technologies comes from the derivative known as machine learning. The history of machine learning dates to the earliest days of computing, when pioneer Alan Turing proposed the question “can computers learn?” (Turing, 1950). The principle behind machine learning is to allow software to continuously process and analyze data to determine if there are correlations that are not obvious or apparent, such that the computer “learns” more about that dataset through pattern recognition (Algorithmia, 2020; Chakravorty, 2016). Typically, the machine learning software takes a segment or subset of the dataset to “train” itself on the potential correlations, sometimes using prior instructions. The software continually processes that data until it “learns” or discovers the key relationships among variables through pattern recognition and develops algorithms that are, in turn, compared to the rest of the dataset to confirm associations among variables and specific predictors. In essence, the software is performing a kind of cross-validation on the dataset using the associations derived from the subset. Thus, machine learning can help the decision-making process by elucidating the potential relationships between variables.

An example of machine learning in HR would be identifying which kinds of employee training programs have been most effective in maintaining high product quality. The machine learning process could compare data from employee training sessions, data from production sources, data from quality management reports as well as customer comments on quality. The goal of machine learning in the HR context is to create *predictive analytics* that can help HR managers to make better decisions in managing human capital (Levenson, 2018; Minbaeva, 2018; Waters et al., 2018). In this context,

HR machine learning projects can be seen as a subset of the general category of HR analytics that uses data to make decisions (Levenson, 2018; Minbaeva, 2018; McIver et al., 2018; Waters et al., 2018).

What has accelerated the use and applicability of machine learning technologies in the last decade is the increasing availability and variety of data, since machine learning technologies require large datasets to process information. Most traditional HR data sets (e.g., spreadsheet-based datasets of headcount, payroll, etc.) in and of themselves are not appropriate for typical machine learning projects, although they may have value when in combination with other data sources (discussed below). But given the expense of acquiring and analyzing these other sources of data, machine learning projects are not likely to be engaged in by smaller organizations. In the HR context, perhaps the most significant usage of new data sources and machine learning has been in the screening of applicants through social media sources (Angrave et al., 2016; Marler & Boudreau, 2017). Many firms use Linked-In, Twitter, and other sources to find new employees, but machine learning accelerates that process by finding those who might be “ideal” applicants through pattern recognition. For example, DeepSense is an artificial intelligence service that develops personality profiles of potential employees based on social media profiles (Thibodeaux, 2017).

Moreover, emerging data sources will dramatically increase the possibilities for machine learning and likely will have significant application to the HR context. Two of these sources are video analytics and the “Internet of Things.” Digitalized video has been available for many years, with cameras tracking many aspects of life, but the capacity to analyze video for specific patterns has accelerated in the last few years. A well-known consumer example is the use of facial recognition as the default security screen in the Apple iPhone X. But vendors including Hewlett Packard, IBM, and Google now sell specialized software that can analyze video from social media and YouTube pages (Hewlett Packard Enterprises, 2016; IBM, 2015; Pande, 2017). While most of this software is currently focused on marketing applications to determine customer acceptance of specific products and services, it is not difficult to extrapolate a usage of video analytics for training and development purposes. For example, the CEO of Starbucks has acknowledged that the firm is actively using machine learning for menu and service decisions (CNBC, 2019). In turn, Starbucks could use video to determine specific instances where employees are “serving customers with a smile” and determine how to replicate those behaviors in other locations. Moreover, HR could assess through video analytics whether specific incentives or training programs are associated with particularly exceptional service, or whether other factors—such as turnover of key employees (Call et al., 2015a; Nyberg & Ployhart, 2013)—make exceptional service a rarity in certain locations.

The Internet of Things will also dramatically increase the availability of data. The Internet of Things concept refers to the fact that many products, such as automobiles, refrigerators, and washing machines, now include sensors that transmit data about their status, and soon these sensors will be able to communicate with each other (Krotov, 2017; Porter & Heppelmann, 2014; Saarikko et al., 2017). A classic example of sensors transmitting their status is Radio Frequency Identification (RFID) chips that enable package delivery tracking. While early versions of these systems have been available in the automobile industry for years, as the technology infrastructure improves to be able to handle vastly increased data transmission (including full implementation of “5G”) from all of these embedded sensors collectively, IT managers will be able to network these sensors together with other information systems to not only determine the status of packages but also give operational and strategic input to decision-makers. For example, if sensor data from a significant number of products indicates product failures, then these failures can be traced back to the

particular factories that created the products, and even to the individual employees who produced the defective ones.

In turn, the Internet of Things significantly increases HR capabilities to analyze the effectiveness of a firm's human resources (Hamilton & Sodeman, 2020). For example, machine-learning software could analyze social media posts to search for patterns of dissatisfaction with company products in certain geographic regions. Other software could then evaluate sensor feedback to determine if there are operational concerns with those products (Saarikko et al., 2017), such as problems in the particular plants that make products sold in a geographic region. In turn, turnover data could be analyzed to determine if turnover (e.g., of "star" employees) in those plants (Call et al., 2015a, b) contributed to the plants' manufacturing of lower quality products, resulting in poor recommendations in social media. Alternatively, HR could determine if there had been training lapses, labor issues, manager changes, or other potential causes of lower quality products in those plants. Thus, the use of machine learning in conjunction with technologies such as the Internet of Things has the potential to connect and combine disparate data sources related to marketing, operations, and HR functions (as well as other functional areas within the firm value chain) to discover, resolve, predict, and prevent organizational problems (see Table 1).

Some hurdles remain in order for HR to use these new data sources effectively, including comprehension by HR managers of the firm value chain and key differentiating operational capabilities (such as quality manufacturing or exceptional customer service); the capacity to appropriately research and conduct relevant statistical analyses; the ability of senior HR managers to understand firm data sources and data architecture; and the ability of HR to develop strong relationships with firm stakeholders, such as IT, marketing, or production, especially since many important data sources will be beyond the traditional purview or control of HR (Hamilton & Sodeman, 2020; McIver et al., 2018). Some of these concerns are beyond the scope of this paper but have been discussed previously (Angrave et al., 2016; Hamilton & Sodeman, 2020; Levenson, 2018; Marler & Boudreau, 2017; Minbaeva, 2018; Tomczak et al., 2018).

But less discussed is the fact that any of these potential workplace improvements to be gained through machine learning could easily be engulfed by significant legal and ethical concerns that can arise through the use of this technology. For example, there have been a

Table 1 Potential data sources for machine learning in HR

Data source	Examples of data obtained from the source
Video	Employee facial expressions Customer facial expressions Employee physical behaviors/movements
Product sensors	Product location Product defects
Social media	Customer complaints Customer personality Customer use/misuse of products Employee attitudes
Operations	Equipment failures Change in suppliers
HR	Turnover Training Manager changes
Marketing/Sales	Customer satisfaction surveys Sales dollars

number of concerns raised with using video for facial recognition technologies, especially given that these technologies tend to be considerably less accurate for people of color, and as a result law enforcement has been criticized for using machine learning to apprehend suspects. (Dethlefsen, 2019; Stewart, 2019). Use of facial recognition technology for work-related purposes could face comparable issues (e.g., concerns about being able to accurately recognize “service with a smile” across different racial/ethnic groups). Another concern is that the Internet of Things may be vulnerable to cyberattack and resulting possible manipulation of the data (Stone, 2020), which might call into question employment-related decisions made on the basis of that data. Thus, effectively using machine learning for HR purposes hinges on both the accuracy and appropriateness of the data sources as well as the potential uses of the data.

In this paper, our focus is on some of the challenges that HR managers will face in the machine learning era. First, we discuss below some of the key legal criteria that could impact machine learning projects. We then discuss two of the major ethical concerns that will likely result from the implementation of machine learning, followed by practical suggestions.

HR Legal Issues Impacting Machine Learning

Although machine learning has the potential to improve organizational efficiency and effectiveness, the importance of ensuring the appropriate treatment of employees during the machine learning process cannot be emphasized enough. The most significant legal concern would be the inadvertent violation of employment discrimination laws, such as the US equal employment opportunity (EEO) laws (e.g., Title VII of the Civil Rights Act of 1964, Americans with Disabilities Act, Age Discrimination in Employment Act), the Canadian Human Rights Act, and European Union regulations such as the Employment Equality Framework Directive (see Table 2).

Obviously, if machine learning algorithms predicted that firm performance would improve if all minorities of a certain race/ethnicity were fired, we hope that the vast

Table 2 Examples of potential legal and ethical concerns

Examples of Legal Concerns	Impact
US Employment discrimination laws (Title VII, ADA, ADEA)	Adverse impact; perpetuating historic discriminatory patterns
Fair Credit Reporting Act (FCRA)	Data sets and resulting analytics not in compliance
Health Insurance Portability AND Accountability Act (HIPAA)	Data sets and resulting analytics not in compliance
General Data Privacy Regulation (GDPR)	Employee consent and privacy by design required; some analyses (such as employee surveillance) may not be feasible; “right to be forgotten” could impact analyses
Examples of Ethical Concerns	Impact
Privacy	Feelings of “intrusion” if disclosure not sought or analytics violate personal space; possible counter-productive behaviors
Organizational justice	Inappropriate processes or unexpected outcomes (even if accurate) could generate employee feelings of injustice

majority of HR managers would reject the conclusion as spurious as well as illegal. Such an example of *disparate treatment* on the part of the software would be obviously wrong to HR. However, *disparate (adverse) impact* discrimination issues are likely to be less apparent and may serve to exacerbate previous discriminatory patterns. For example, suppose that a machine learning algorithm determines that breaks in previous employment records are related to absenteeism or turnover. This might seem to be a reasonable basis on which to reject an applicant, but it is likely to have an adverse impact on women, given that women tend to have more fragmented careers (e.g., for childrearing, but also for other issues; see Zimmerman & Clark, 2016).

Beyond adverse impact, machine learning could perpetuate past discrimination in more subtle ways, even if protected class status were not explicitly included in the analysis. For example, suppose that a particular group of employees are identified through machine learning as “stars” based on their historical and current performance. Machine learning might be used to create a profile of the experiences, background, and other characteristics that are common among these stars. The organization then would use this profile to determine which other employees should receive training in order to develop them into future stars. But if subtle forms of discrimination had been tolerated by the organizational culture in the past (e.g., certain races or ethnicities had been excluded from training by prejudicial supervisors), the machine learning might inadvertently perpetuate such discrimination. In this case, the machine might “learn” that only members of certain races or ethnicities became stars in the past, and it would be technically correct (that is, the identification of current stars is *not* spurious but accurate), but in trying to identify future stars, machine learning might create an algorithm that includes this implicit prejudice. In other words, human biases led to the identification of organizational stars who had a certain pattern of characteristics; the machine learning would attempt to replicate this pattern, and in doing so would replicate these implicit biases. In this case, the machine learning algorithm would result in adverse impact, based on historical disparate treatment.

Such bias can be inadvertently built into the algorithms without explicitly including protected class status. To continue with the above example, HR or payroll databases will likely include the employee’s zip code, yet in many cities different zip codes will have different distributions of races and ethnicities. As part of the machine learning analysis of past star data, the software learns to associate certain zip codes with development into stars and embeds zip code as a predictor of stardom, perpetuating discrimination. A similar scenario could occur for age discrimination. For example, suppose that supervisors have been reluctant to provide training to older employees, assuming that “you can’t teach an old dog new tricks” or that training is wasted on older individuals who will be close to retirement. The algorithms that are created could reflect that prejudice, for example, where machine learning finds a significant (negative) relationship between organizational tenure and “stardom” because no employees over a certain age have been given the opportunity to become stars.

In a real-world example, in 2018 it was publicly reported (Reuters, 2018) that in 2015 Amazon’s attempt at using machine learning to screen applicants failed because the machine-learning process screened out female applicants. The machine learning software learned to reject resumes that included items such as “women’s chess club captain” and focused on applications that included such words as “captured” or “executed” (words more likely to be on male resumes), simply because those were the words on previous applications of successful hires, since Amazon had hired mostly males in the past. The attempt to have machine-learning processes be a primary screening mechanism without human intervention was reportedly abandoned (Reuters, 2018).

Even more insidious would be the use of machine learning to formalize a kind of disparate treatment. Barocas and Selbst (2016) describe the possibility of “masking” by decision makers, in which they intentionally bias data collection or analysis, so as to ensure that the algorithms produced by the machine learning are unfavorable to certain protected class subgroups. That is, the decision-makers have introduced adverse impact into the machine learning process but have done so in an intentional manner that is difficult to detect. In this case, masking is a pretext for disparate treatment (Barocas & Selbst, 2016). For example, imagine that a manager desires to discriminate against individuals with disabilities. Obviously, introducing a code for disability into the analysis would be a red flag for discrimination. The manager instead ensures that the cost of equipment for employees is part of the data analyzed, a proxy that would probably not be a red flag for organizational leaders. However, accommodations for individuals with disabilities might necessitate their having higher equipment costs. In turn, the machine learning would determine that the profit returned per employee would be less for those with higher equipment costs, and those with high equipment costs should be terminated. In this scenario, the supervisor has engaged in intentional discrimination against disabled individuals but masked it with a proxy that appears reasonable to decision-makers.

Beyond the US EEO laws noted above, other laws will be relevant to machine learning issues. In the US these include the Fair Credit Reporting Act (FCRA) and Health Insurance Portability and Accountability Act (HIPAA). These laws were not originally intended to apply to machine learning or to workforce analysis processes in general. However, for several years researchers have suggested that FCRA applies to organizations using the Internet to assist the screening process (Black et al., 2015; Davison et al., 2012). In 2016, the Federal Trade Commission ruled that credit information cannot be used in a way that would create adverse impact and also determined that material information must be provided to those whose data are analyzed (Federal Trade Commission, 2016). Since some HR data sets are originally sourced from third parties, it is conceivable that credit information could be mixed in with the data (Federal Trade Commission, 2016). Similarly, any employee data that might include health information (which could include a significant amount of HR data related to benefits or absences) is likely to be governed by HIPAA. Health insurance data, even if generated internally, cannot legally be used in the US for any non-health-related purpose. This implies that to be legal, data must be stripped of identifying health or credit information before it is analyzed (Cohen & Mello, 2018). At the very least, this fact would complicate using some datasets for evaluation of workforce performance.

In the EU, the General Data Protection Regulation (GDPR) covers any data that is processed in Europe (input, queried, transmitted, distributed, or stored), even if non-European firms or non-European employees are involved (European Commission, 2018). For years, in organizations based in Europe or in cases where employment applicants are from European countries, consent from data subjects has normally been required before personal data is analyzed (Davison et al., 2012). Under GDPR, the data subject is required to have access to that data at any time, which could complicate machine-learning processes. For example, the GDPR’s “right to be forgotten” (European Commission, 2018, p. 11) could impact machine learning in locations where there has been high turnover and exiting employees ask for their data to be removed. If there are disputes between employer and employee, the GDPR has a specific bias toward individual rights: “an individual can ask to have the processing of their personal data restricted while it is determined whether or not your legitimate interest overrides their interest” (European Commission, 2018, p. 12). Moreover, as machine learning technologies mature and create automated decisions (regarding, for example, bonuses or promotions), the GDPR expressly gives the right to challenge those

decisions (European Commission, 2018, p. 12). Under the GDPR, data scientists and data architects are required to have more stringent preparation for the analysis as GDPR requires “data protection by design and by default” such that no analysis can be conducted that might infringe on an employee’s data rights (European Commission, 2018, p. 14). It is not completely clear what impact GDPR would have regarding product-tracking sensor data that can be traced back to employees who make products but that does not actually track the employees’ actions. However, we believe that GDPR would be likely to apply to such situations.¹

Given the various laws relevant to machine learning in the HR context, it is clear that while machine learning offers many useful possibilities for organizations to enhance productivity, such endeavors cannot be left solely to data scientists, who may be unfamiliar with the various laws that are relevant to personnel decisions. Moreover, data scientists may be unfamiliar with principles of sound HR practice. For example, in terms of hiring, the Uniform Guidelines on Employee Selection Procedures (UGESP; Equal Employment Opportunity Commission, 1978) address issues of adverse impact and validation procedures; although the UGESP do not have the force of law, they are given great deference by the courts. Senior HR managers must be in on the ground floor of the data selection process and must guide conversations to determine what data to include from a theoretical perspective. HR must also offer guidance to data scientists to be sensitive to analyses that would aggravate possible discriminatory patterns or confirm implicit biases (such as in the earlier Amazon example).

Ethical Concerns Impacting Machine Learning for HR

In the previous section we addressed legal issues surrounding machine learning in the HR context. In this section we address two significant ethical concerns. Even if a firm’s practice is legal, ethical concerns in the treatment of employees remain important to consider as HR is often seen as a guardian in ethics matters (Armstrong, 2016).

Privacy

Perhaps the most significant ethical concern regarding machine learning is a potential loss of privacy. Over a hundred years ago, US Supreme Court Justice Louis Brandeis defined privacy as the desire to be “left alone,” for “solitude and privacy have become more essential to the individual; but modern enterprise and invention have, through invasions of his [*sic*] privacy subjected him to mental pain and distress” (Warren & Brandeis, 1890, p. 196).

Far more than Justice Brandeis could have imagined, machine learning allows for a potentially greater loss than ever of personal privacy through a great quantity of personal data and preferences being gathered, distributed, correlated, reviewed, and analyzed among multiple data sources. Data available from many public websites can be correlated with proprietary data, which in turn can be analyzed to predict outcomes such as customer preferences. Thus, machine learning may find out more about us than we want organizations

¹ We should also note that the impact on GDPR from the withdrawal of the United Kingdom from the EU is not currently fully known, especially in the employment context.

to know, especially in the US context. An infamous public example involved the retailer Target sending information about diapers to a teen after the machine learning concluded (apparently accurately!) from her purchase patterns that she was pregnant – supposedly before her own father knew (Duhigg, 2012).

We have placed our discussion of privacy in this section regarding ethical concerns since legal enforcement of privacy issues varies considerably. As discussed above, the GDPR is essentially a response by the EU to privacy concerns. However, in the US, no such broad right to privacy applies to employees. US public sector employees have some privacy protections under the Fourth Amendment to the Constitution. In California, a right to privacy is enshrined within the state constitution (California Constitution, Article 1, Sect. 1), and more than 10 US states have laws against employers requiring access to employees' personal Internet or social media accounts (National Conference of State Legislatures, 2019). However, outside of these and a few other exceptions, private sector employees in the US have relatively limited legal privacy protections in the workplace. Even in the context of a few of the recent state laws on privacy, such as in the California Consumer Privacy Act of 2018, the Virginia Consumer Data Protection Act and the Maryland Personal Information Protection Act, the statutes are more focused on consumer protection rather than on employee protection, thus differing from GDPR (Dobias, 2021; European Commission, 2018; Frosh, 2021; Office of the Attorney General, 2021). Moreover, common law privacy protections (under tort law) may not apply to many machine learning initiatives. Generally, elements of a privacy violation claim involve an intentional intrusion into private affairs (Walsh, 2019). But under US law, private firms own their own data, not employees, and typical workplace actions are under the purview and guidance of the employer. Thus, unless specific discrimination laws are violated, an algorithm is not likely to legally be considered an "intruder." Therefore, especially in the US, various invasive practices to collect data for machine learning purposes are likely to be legal in most situations.

Nonetheless, even if many machine learning practices are legal from a privacy perspective, HR managers should consider whether these practices might be perceived by employees as violations of privacy and unethical, and thus demotivating. For example, an MIT researcher as well as several consulting firms have used sensors in employee badges to "track" the behaviors of employees and then compare these behaviors with performance data (Peck, 2013), in order to determine what kinds of individual actions, activities, and collaborations between employees result in higher creativity and efficiency. Companies such as Humanyze, Wal-Mart, Microsoft, and Amazon have developed tracking systems that collect audio, geolocation, accelerometer, and other data from employees throughout their workday (Heath, 2016; Shell, 2018; Sheng, 2019). These systems can identify the most productive actions and knowledge-sharing activities by individuals and teammates, including identifying informal workgroups that have members across multiple locations (Shell, 2018). Tracking data can be used for real-time assessment of employee performance, and when combined with other data sources, can be helpful for linking the job, product quality, and training. But from a privacy perspective, does the use of tracking systems violate employees' desire to keep some of their actions private or, as Justice Brandeis might put it, do these systems inflict mental pain? For example, most employees might consider it to be an invasion of privacy for the employer to track how often they visited the restroom or water cooler during the day (and such an invasion of privacy might ultimately lead to violation of the Americans with Disabilities Act by impacting individuals with certain medical conditions). For that matter, should an employer be given access to every conversation between employees? Privacy might also be jeopardized if, for example, an employee is tracked going to meet with an ombudsperson or union representative.

Ultimately, tracking devices are likely to be perceived negatively by employees, who will see these as evidence of management's distrust. Indeed, in some Amazon warehouses, these pervasive monitoring systems that fed machine learning-driven decision-making were perceived by employees as draconian and wearying (Spitznagel, 2019).

Moreover, as many HR functions move to mobile phone applications (i.e., "apps") such as WorkBright, Namely, Halogen, TalentSpace, and Zenefits (Stringfellow, 2019), the potentially persistent nature of some of these apps (especially with regard to the phone's GPS location-based services) has been called into question. For example, in 2019, using software from a location-based data services company, a New York Times investigative team was able to identify and track within minutes a Microsoft software engineer as he traveled to Amazon headquarters on several occasions to interview for, and eventually accept, a position with Amazon (Thompson & Warzel, 2019). While the manager was reportedly not entirely surprised about the ability of the apps to track him in this way, he also questioned whether the tradeoff between the app services provided and the loss of privacy was worth it (Thompson & Warzel, 2019).

Will similar data collection strategies in other organizations, from any data source, be considered "invasive surveillance," and ultimately a de-motivator, negating any positive impact from machine learning? A 2019 Accenture poll found that 64% of employees are concerned about a possible eroding of their privacy (Sheng, 2019), and many could perceive these systems as creating a pervasive "Big Brother" environment. We would recommend that HR managers closely monitor the research processes of machine learning in the organization, being careful to caution against any data collection procedures that would seem to invade employees' privacy. Data collection that is perceived as particularly invasive may result in counterproductive behaviors (Sackett, 2002) from some employees; some of these behaviors could include knowledge-hiding or knowledge-hoarding (Serenko & Bontis, 2016), which would be particularly difficult to detect. HR managers should also take seriously the recommendations of Guzzo et al. (2015), who suggest that informed consent and privacy protections should be built into the data collection process, as well as Illingworth (2015), who suggests that the process should be based on "privacy through accountability" (2015, p. 569) in terms of allowing users (i.e., employees) to have control over how their data is utilized. Additionally, requiring informed consent from employees for machine learning projects will improve the climate of inclusiveness and transparency in the organization, which is related to the ethical concern of organizational justice.

Organizational Justice

John Rawls suggests that justice is the "first virtue of social organizations" (Rawls, 1971, p. 3). Similarly, Loi et al. (2012) contend that organizational justice is a fundamental value and virtue in any organization, with justice being a base for the assumptions of how humans should be treated. The concept of organizational justice has been prominently discussed in the HR, ethics, and organizational behavior literature (Armstrong, 2016; Cohen-Charash & Spector, 2001; Colquitt, 2001; Cropanzano & Stein, 2009; Konovsky, 2000; Loi et al., 2012; Rawls, 1971), with a focus typically on three primary forms of justice—distributive, procedural, and interactional. Distributive justice deals with the fairness of the outcomes of organizational decisions and has been linked to Aristotle's law of proportions (i.e., that fairness can be judged by comparing ratios of one's contributions and another's contributions to their respective rewards; see Adams, 1963; Cropanzano & Stein, 2009).

In contrast, procedural justice refers to the fairness of the processes by which decisions are made and can be seen as a form of Kantian deontological ethics, insofar as the actions of decision makers should be judged based not on the outcomes of the process, but on the rightness or wrongness of the actions, or internal standards and convictions (see Cropanzano & Stein, 2009). Justice in HR processes such as selection, performance appraisal, and compensation may help give employees “voice” as stakeholders, which in turn gives them feelings of value and control over their situations (Cropanzano & Stein, 2009; Konovsky, 2000; Loi et al., 2012). Finally, interactional justice, the fairness of personal treatment, includes two sub-categories: interpersonal justice (treating others with respect and sensitivity) and informational justice (communicating explanations for decisions; Bies & Moag, 1986; Colquitt, 2001). Although interactional justice is typically viewed as separate from procedural justice, some scholars have treated it as part of procedural justice, and it does share procedural justice’s deontological focus on the ethicality of the actions or behaviors of decision-makers, rather than on the outcomes.

The use of machine learning in organizations has the potential to be unethical insofar as it might violate any of these categories of justice. First, since machine learning involves software that makes decisions about data, presumably without human intervention, it is possible for that software to violate procedural justice principles. Fair procedures include aspects such as consistent application of rules, accurate use of information, providing the opportunity to be heard, and safeguarding against biases (Greenberg, 2011). One of the advantages of using machine learning algorithms *appears* to be the consistent application of rules, insofar as the presumably unbiased computer determines the result, not a potentially biased human. However, this assumes that the inputs are accurate and unbiased, for as the old computer adage goes, “garbage in, garbage out.” Indeed, as McAbee et al. (2017, p. 284) have noted: “*Someone* must decide what data to collect...and...*someone* must then analyze and interpret the data” (italics in original). Thus, a human is ultimately responsible for the fairness of the process, and this is where we see the role of HR to be crucial. Even if the data input into the analytics is technically accurate, it could be deficient (i.e., not presenting the whole picture of performance, such as sales dollar figures that ignore customer retention), or contaminated (i.e., affected by other factors beyond the employee’s control, such as sales territory). To protect procedural justice, HR should particularly safeguard against biases and ensure that accurate employee information is used. HR should also provide for an opportunity for employees to be heard (e.g., an appeals process for any employment decisions made based on algorithms). Decisions directly made by algorithms without an allowance for employee “voice” (Cropanzano & Stein, 2009; Konovsky, 2000; Loi et al., 2012) are very likely to be perceived as unfair and intruding on employee rights while eliminating employees’ sense of control.

However, even if an appropriate process is used, it is also possible for considerable employee backlash to occur regarding unfavorable distributive justice. That is, even if technically accurate, the results of machine learning analytics are likely to change perceptions about who are good performers, since current appraisals are often based on subjective judgments, while analytics may be based on hard data from sensors or other real-time measures. Some employees may therefore believe they are being unfairly treated through analytics (e.g., they might assume the data to be deficient or contaminated), if they perceive that their contributions are greater than the algorithms conclude. But sometimes the outcomes will not be perceived to be fair even when the data and conclusions are accurate, as, for example, when the algorithms correctly determine that employees who were previously considered strong performers are actually not the

strongest performers. When there is employee pushback regarding less desirable outcomes as a result of accurate conclusions from machine learning, HR managers should demonstrate care and compassion.

Related to that, the two forms of interactional justice, interpersonal and informational justice, will be relevant. Interpersonal justice would involve actively displaying concern for individuals and providing a rationale face-to-face for any decisions derived from machine learning (rather than simply defaulting to “well, since the computer says so, we can’t do anything about this!” or notifying employees about a negative machine learning decision via an impersonal medium such as e-mail). Informational justice would include informing employees up front about the machine learning evaluation process, including which data sources might be used for the evaluation, as well as clearly discussing why any changes are being implemented as a result of the machine learning process. Thus, ensuring informational justice in machine learning should involve transparency, which is often lacking in machine learning, leading to its being perceived as a “black box.” (We discuss the implications of the “black box” nature of machine learning in greater detail in the next section).

Taking all of these aspects of justice together, HR’s role to ensure fairness should also involve procedures for seeking employee input into the analytic process, communicating with employees about the use of analytics in decision-making, ensuring the decisions made based on the analytics are appropriate and just, and doing all of this in a respectful and sensitive manner. Given the legal and ethical concerns we have outlined, we discuss below some practical implications and specific guidance for HR managers as they navigate this emerging field of machine learning.

Practical Implications

HR has important responsibilities in terms of reducing firms’ exposure to litigation, as well as the ethical responsibilities of the profession for promoting fairness and justice (e.g., SHRM, 2014). Given these responsibilities, it is crucial that HR take a lead in terms of addressing the legal and ethical issues inherent in machine learning when applied to employment contexts. (See Table 3 for an example of a process). In particular, an important theme of this paper is that while complex analyses of variables may be performed (i.e., are possible thanks to machine learning), it does not mean they are legally or ethically appropriate.

Minimizing Legal Exposure

As discussed above, there are multiple ways in which machine learning can potentially violate US as well as international laws. In order to avoid running afoul of these laws, it is imperative for HR to not only develop skills in understanding data analytics but also to partner with data scientists from the very earliest stages of machine learning projects. Essentially, HR needs to know what variables are included in the datasets, and if possible, what variables used in the algorithms are being built via machine learning. Below we discuss how this partnership can help prevent violations of the various laws discussed earlier.

First, the EEO laws prohibit discrimination in terms of both disparate treatment and disparate impact. Unless intentionally used for enhancing diversity and inclusion initiatives,

Table 3 Process for legal & ethical use of Machine Learning (ML) in HRM

Stages in the Process	Steps to follow / Decisions to make
Stage 1: Determine Purpose of Using ML for HRM Purposes	<ul style="list-style-type: none"> ● Work with organizational leadership to determine the relevant outcomes (turnover, performance, etc.) to be predicted by ML ● Ensure these outcomes are of strategic importance to the organization
Stage 2: Determine the Proposed Data Collection Process	<ul style="list-style-type: none"> ● Partner with data scientists within and/or outside the organization <ul style="list-style-type: none"> ○ Identify possible sources of data (internal databases, new data collection initiatives such as employee tracking, etc.) ● Examine the data sources and data collection initiatives for legal and ethical concerns <ul style="list-style-type: none"> ○ Consult a demographically diverse team to review these sources/initiatives ○ Consider bringing in outside experts to examine potential for biases ○ Identify specific variables to exclude <ul style="list-style-type: none"> ■ Protected class variables ■ Credit or health-related data ■ Variables that might result in disparate impact ■ Variables that are potentially invasions of privacy ■ Variables that potentially create justice issues ● Review with the data scientists and organizational leadership what variables must be omitted from analyses and why
Stage 3: Communicate with Employees About the Proposed Process	<ul style="list-style-type: none"> ● Organizational leaders, HR representatives, data scientists, and members of review teams should be involved in the communications <ul style="list-style-type: none"> ○ Explain the rationales for using ML in terms of how it will help them and the organization ○ Explain the process that was followed to ensure legality and protect employees' privacy ○ Explain other safeguards to protect the security of their data and opportunities for appeals ● Where possible, informed consent should be obtained from individual employees, especially for sensitive data

Table 3 (continued)

Stages in the Process	Steps to follow / Decisions to make
Stage 4: Review the Results after the Algorithm Has Been Developed	<ul style="list-style-type: none"> ● Examine the algorithms for possible bias across protected class subgroups ● Examine any recommended employment decisions for discrimination (e.g., adverse impact) across protected class subgroups <ul style="list-style-type: none"> ○ Examine impact on other relevant groups that may not have protected class status (e.g., veterans, younger employees, employees with children) ● If bias or adverse impact is found: <ul style="list-style-type: none"> ○ Attempt to identify the source of bias or discrimination and remove the problematic variable(s) if possible <ul style="list-style-type: none"> ■ Re-evaluate the modified algorithm for predictive validity and disparate impact ○ If it is not possible to identify the source of bias or discrimination, reconsider using the algorithm in light of evolving legal guidance ● Provide an appeals process for other employees who are affected by the algorithm's recommendations <ul style="list-style-type: none"> ○ Compare other sources of information (e.g., supervisor and coworker input) that may contradict the algorithm's recommendations ○ Ensure that the process is perceived to maintain justice

HR professionals should insist that no explicit protected class variables (e.g., race, sex, age, etc.) are included in the datasets, even if they were to be predictive of relevant outcomes. Beyond these, there are other variables that might lead to disparate impact (e.g., zip codes, marital status, number of dependents), and it would be wise to explicitly exclude some of the more commonly problematic variables. Variables pertaining to credit or health, which might be in certain HR databases, also should be removed for FRCA and HIPAA considerations. In addition to the involvement of senior HR managers, employees from various job categories, demographic groups, and legal representatives, as well as union representatives when applicable, could be consulted to evaluate the variables that are included in the datasets. Outside experts (e.g., in the field of AI and ethics or diversity and inclusion) might also be consulted at this stage, to assess the possible impact of the algorithms. We realize that given the large number of potential variables being combined in the algorithms, such a detailed level of attention to every variable might not be considered feasible by organizations, but we believe it is important to take such steps to ensure legal compliance.

The end results (e.g., recommendations as to employment-related actions) of the machine learning algorithms also should be examined for adverse impact across various protected class subgroups. We would also suggest that adverse impact be examined for other combinations of protected classes (e.g., race by sex combinations, such as white males compared to black males, or women with children compared to men with children); even if such comparisons are not typically part of disparate impact analyses, these could be indicative of discrimination (e.g., sex-plus discrimination).

Moreover, when disparate impact is found, it would be advisable to identify the particular variable or combinations of variables that are causing it and remove those from the dataset, then re-evaluate both the prediction and disparate impact using the modified algorithm.

However, we recognize that it may not be possible to identify the offending variables in many cases. Although there is a trend toward developing machine learning algorithms that are more transparent (i.e., explainable machine learning; Defense Advanced Research Projects Agency, 2020) in terms of the outcomes of the machine learning process, in many cases HR may be faced with the problem of machine learning being a “black box.” In other words, the algorithm may provide an employment decision about an applicant or employee, but the logic path may not be apparent to the user (Card, 2017). This is certainly troubling from a legal as well as an ethical perspective, as the employer cannot readily point to a particular reason why the decision was made.

Related to this, while discussion of the technical properties of machine learning software are beyond the scope of this paper, we should also caution that a particular type of machine learning, termed “unsupervised learning,” may be especially susceptible to the concerns we have outlined. In unsupervised learning, the software discovers the variables that are being used to train the algorithm, rather than a priori by data scientists, so variables that promote adverse impact might inadvertently become part of the resulting algorithm. Particular caution should be used when the machine learning software is using so-called “deep learning” processes which layer the data analyses (IBM, 2020; Mathworks, 2020a, b). We urge HR managers at the very minimum to be aware of and monitor machine learning projects that use unsupervised or deep learning structures.²

The problem of the black box becomes particularly concerning when disparate impact is present. It is important to note that under US EEO law, a plaintiff who challenges a selection practice for disparate impact is required to identify the particular practice that is causing the disparate impact. However, in machine learning (and especially in deep learning) it may not be possible to identify any one variable that leads to the disparate impact, but instead there can be a convoluted collection of variables combined in various linear and nonlinear ways that results in disparate impact. For example, in the Amazon case cited earlier, Amazon reportedly did not include variables designating sex, but the algorithm interpreted that when activities typically engaged in by women were part of a resume, those candidates were excluded (Reuters, 2018). Although an employer may be able to assert the validity of the algorithm as a defense to adverse impact, the legal defensibility of using algorithms in HR is an evolving area. These points clearly argue for the importance of only including variables that have a reasonable degree of job relevance in order to comply with the spirit of the law.

Finally, the algorithms should be examined for other forms of bias, such as differential validity and differential prediction (see Berry, 2015, for a convenient summary of various forms of test bias, which may also apply to machine learning). Differential validity occurs when a predictor is less valid for one subgroup relative to another, whereas differential prediction occurs when the equations (e.g., the algorithms) are different for different subgroups. It is possible that an algorithm developed on training data may be biased (e.g., have differential validity) when applied to a new dataset, particularly when one subgroup is underrepresented in the training data (e.g., as in the Amazon case discussed earlier).

In sum, when considering the lack of transparency in the black box of machine learning, care must be taken by HR on the front end and back end of the process to ensure legal defensibility and fairness. In other words, HR should oversee the process in a way that

² Since the field of machine learning is rapidly progressing, and given that a technical discussion of machine learning issues is beyond the scope of this paper, we refer the reader to IBM (2020), Mathworks (2020a, b), and RSIP Vision (2021) for a basic overview of these concerns.

ensures that (1) training datasets are as representative of various protected classes as possible, (2) only job relevant, non-discriminatory variables are input into the algorithms, and (3) algorithms do not have bias or discriminatory effects in the outcome.

Upholding Ethical Principles

To address the ethical issues inherent in machine learning, HR must also forge a strong relationship with the employees it serves to protect. With regard to the ethical concerns we are exploring in this paper, this relationship should be built on the principles of privacy and justice, with a particular focus on ensuring that employees are provided with protections from harm as well as justifications for the use of their data.

As mentioned earlier, in the EU, the GDPR provides privacy rights to individuals with respect to their data. Given the global nature of many organizations, some organizations may already have to comply with the GDPR's privacy regulations. But, even if an organization is not under the GDPR's purview, it would be advantageous to provide employees with the same kinds of privacy rights, so they will perceive that they are part of an ethical organization. A likely starting point for HR would be to examine the data collection processes to determine whether the processes might be considered invasive. Given that there are a variety of criteria or situations where different employees might feel their privacy is invaded, it would not be sufficient for an HR professional to unilaterally determine where privacy is violated (e.g., "I don't mind wearing a tracking device, so it should be fine with others"). One possible vehicle for determining whether a machine learning process violates privacy would be to convene a privacy panel consisting of representatives from various employee groups (including those who would be affected by the practice), as well as legal advisors and union representatives when the machine learning process involves especially controversial areas. The panel should also be diverse in demographics, as different groups may have different ideas of what is acceptable and what is not. Such a panel might also consider the databases that will be included in the analysis and determine whether the use of those might constitute unethical invasions of privacy (even if their use is legal).

Additionally, HR professionals should follow the various principles of justice described earlier, particularly in focusing on procedural and interactional justice. Using panels of employees as described above to review data collection practices and database choices as well as to review individual variables in the algorithms would be helpful for safeguarding against biases and ensure that accurate information is used, which would be important for procedural justice. Moreover, including representatives from the various employee groups will provide a degree of voice in the process for those affected by the algorithms. To ensure both procedural and distributive justice, an appeals process will be essential. For an employee who faces a negative employment decision based on an algorithm, a fair hearing is essential, ideally by a disinterested party or parties. (Note that HR would *not* typically be perceived as a disinterested party in this situation.) A lack of an appeals process that is perceived to be fair could result in significant employee dissatisfaction, given that a negative employment outcome generated by machine learning processes is likely to be seen as highly arbitrary and quite literally "inhuman." That is, a procedural justice perspective would likely require human oversight of machine learning.

Perhaps most important, however, is the need for HR to communicate with employees in an interpersonally sensitive manner and to provide information at all stages of the process. First, the rationale for using machine learning should be clearly explained to employees in terms of how it will benefit them and the organization. The reasons for using certain

databases or variables should be addressed, as well as safeguards implemented to protect their privacy and the appeals process. These messages may come from HR but also should be reinforced by employee representatives from the panels suggested earlier. Where possible, informed consent should be obtained from individual employees, especially for sensitive data. All of these steps should serve to insert a human dimension into what could be perceived as an often-dehumanizing process. As much as possible, HR managers should ensure that machine learning projects are used as an instrument of improvement not only for the organization but also for the employees, in a non-judgmental way. For example, if machine learning projects are used as a method of “continuous improvement” and creating opportunities for advancement, promotion, and higher remuneration, they will be much more likely to be accepted than if they are used as a method to cut bonuses or downsize the organization.

Ensuring Appropriate Data Management

Undergirding all of this discussion is the basic requirement for HR managers to be familiar with data sources used for employee-targeted machine learning projects. Given how machine learning projects are implemented, it is no longer feasible for HR to simply rely on operations managers, IT personnel, or software vendors to be in charge of evaluating the appropriateness of data. HR managers must have a basic awareness of how given data sets might be legally or ethically inappropriate.

For example, with regard to video data, there is a large body of evidence that machine learning applications tend to discriminate against people of color (Dethlefsen, 2019; Stewart, 2019). That is partly because the machine learning training data often contain far more “white” faces, such that faces of color are inappropriately excluded or singled out. Until software is created that overcomes this limitation, or people of color are included more in a more representative manner in the training datasets, it is not a good idea for projects involving video analytics to have an end-goal of evaluating individual performance, especially where there are negative consequences. As previously mentioned, machine learning projects focused on encouraging positive actions for emulation (e.g., “service with a smile”) which do not have negative consequences for individuals’ employment, are the types of projects that are likely most appropriate in the current context.

Similarly, while the Internet of Things does not have as obvious a discriminatory component as video data, it is also true that collection of sensor data might be particularly sensitive to issues of missing data (e.g., when the sensor can’t transmit back to headquarters), which would skew the results. Additionally, this kind of data might be susceptible to cyber-attack with resulting manipulation, changing, or corruption of the data. Other concerns with the Internet of Things may arise as it becomes more prevalent. From that perspective, the Internet of Things is a good example of how HR managers will have to keep up with implications for further new machine learning data sources.

Conclusion – Guidelines for Using Machine Learning for HR

A central theme of the legal and ethical issues with respect to machine learning addressed here is that there must be *transparency* and *care* by the employer and *consent* by those employed (Guzzo et al., 2015; Illingworth, 2015). A machine learning project is more likely to be legally compliant and perceived as ethically appropriate if employees are

included as stakeholders in the process. Employers (including HR) must communicate with employees about the purpose and means of data collection; ensure to the extent possible that the data collected are accurate, securely maintained, and used in a non-discriminatory manner; and obtain employee consent prior to data collection. Moreover, employees may be more likely to consent to the use of their data if employers clarify how machine learning will improve how the firm is functioning and help diagnose performance problems. Here the employer should make explicit the connection from the externally measured data (e.g., output, product, or customer response data) to the employees' behaviors on the job, to individual and firm performance, and ultimately to enhanced employee learning, compensation, and other rewards (Tomczak et al., 2018). Such lines-of-sight can enable employees to gain a better perspective on the value of their work and how machine learning facilitates their development and success. We also contend that using data sources to track the actual output of employees would be less controversial and more likely to be accepted by most employees than tracking every employee movement. Finally, machine learning outcomes can assist in firm change processes insofar as they constitute an evidence-based practice that can enable a reasoned approach to change (McIver et al., 2018). Thus, if legal and ethical concerns and employee rights are addressed, rather than producing a dystopian dread, machine learning can instead serve as a catalyst toward a more just and harmonious workplace.

Declarations

Human and Animal Study This article does not contain any studies with human participants or animals performed by any of the authors.

Conflict of Interest There are no conflicts of interest to disclose.

References

- Adams, J. S. (1963). Toward an understanding of inequity. *Journal of Abnormal and Social Psychology*, *47*, 422–436.
- Algorithmia. (2020). *How machine learning works*. <https://algorithmia.com/blog/how-machine-learning-works>. Accessed 21 Mar 2021.
- Angrave, D., Charlwood, A., Kirkpatrick, I., Lawrence, M., & Stuart, M. (2016). HR and analytics: Why HR is set to fail the big data challenge. *Human Resource Management Journal*, *26*, 1–11.
- Armstrong, M. (2016). *Handbook of strategic human resource management* (6th ed.). Kogan Page.
- Barocas, S., & Selbst, A. (2016). Big data's disparate impact. *California Law Review*, *104*, 671–732.
- Berry, C. M. (2015). Differential validity and differential prediction of cognitive ability tests: Understanding test bias in the employment context. *Annual Review of Organizational Psychology & Organizational Behavior*, *2*, 435–463.
- Bies, R. J., & Moag, J. S. (1986). Interactional justice: Communication criteria of fairness. In R. J. Lewicki, B. H. Sheppard, & M. H. Bazerman (Eds.), *Research on negotiation in organizations* (pp. 43–55). JAI Press.
- Black, S., Stone, D., & Johnson, A. (2015). Use of social networking websites on applicant's privacy. *Employee Responsibilities and Rights Journal*, *27*, 115–159.
- Blumenthal, E. (2018). Facebook's latest privacy scandal: What we know about the company's handling of user data. *USA Today*. December 19. <https://www.usatoday.com/story/tech/2018/12/19/facebooks-latest-privacy-scandal-what-we-know-now/2361257002/>. Accessed 16 Nov 2019.
- Call, M., Nyberg, A., Ployhart, R., & Weekley, J. (2015a). The dynamic nature of collective turnover and unit performance: The impact of time, quality, and replacements. *Academy of Management Journal*, *58*, 1208–1232.

- Call, M., Nyberg, A., & Thatcher, S. (2015b). Stargazing: An integrative conceptual review, theoretical reconciliation, and extension for star employee research. *Journal of Applied Psychology*, *100*, 623–640.
- Card, D. (2017). *The “black box” metaphor in machine learning*. <https://towardsdatascience.com/the-black-box-metaphor-in-machine-learning-4e57a3a1d2b0>. Accessed 21 Apr 2021.
- Chakravorty, T. (2016). *How machine learning works: An overview*. <https://thenewstack.io/how-machine-learning-works-an-overview/>. Accessed 16 Nov 2019.
- CNBC. (2019). *Starbucks CEO talks artificial intelligence*. <https://www.cnn.com/video/2019/01/25/starbucks-ceo-talks-about-how-artificial-intelligence-is-informing-businessstrategy.html?&qsearchterm=Starbucks%20CEO>. Accessed 25 Jan 2019.
- Cohen, I., & Mello, M. (2018). HIPAA and protecting health information in the 21st century. *JAMA*. <https://doi.org/10.1001/jama.2018.5630>. Published online May 24, 2018.
- Cohen-Charash, Y., & Spector, P. E. (2001). The role of justice in organizations: A meta-analysis. *Organizational Behavior and Human Decision Processes*, *86*, 278–321.
- Colquitt, J. (2001). On the dimensionality of organizational justice: A construct validation of a measure. *Journal of Applied Psychology*, *86*, 386–400.
- Cropanzano, R., & Stein, J. H. (2009). Organizational justice and behavioral ethics: Promises and prospects. *Business Ethics Quarterly*, *19*, 193–233.
- Davison, H., Maraist, C., Hamilton, R., & Bing, M. (2012). To screen or not to screen: Using the Internet for selection decisions. *Employee Responsibilities and Rights Journal*, *24*, 1–21.
- Defense Advanced Research Projects Agency. (2020). *Explainable artificial intelligence*. <https://www.darpa.mil/program/explainable-artificial-intelligence>. Accessed 21 Apr 2021.
- Dethlefsen, J. (2019). *The ethics of machine learning and discrimination*. <https://vce.usc.edu/volume-3-issue-2/the-ethics-of-machine-learning-and-discrimination/>. Accessed 16 Apr 2021.
- Dobias, N. (2021). *Virginia’s Consumer Data Protection Act: What you need to know*. <https://www.pactsafe.com/blog/virginias-consumer-data-protection-act-what-you-need-to-know>. Accessed 14 Apr 2021.
- Duhigg, C. (2012). *How companies learn your secrets*. <http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html>. Accessed 1 Oct 2017.
- Equal Employment Opportunity Commission. (1978). Uniform guidelines on employee selection procedures. *Federal Register*, *43*, 38290–39315.
- European Commission. (2018). *The GDPR: New opportunities, new obligations*. https://ec.europa.eu/commission/sites/beta-political/files/data-protection-factsheet-sme-obligations_en.pdf. Accessed 20 June 2018.
- Federal Trade Commission. (2016). *Big data: A tool for inclusion or exclusion? Understanding the issues*. <https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf>. Accessed 20 June 2018.
- Frosh, B. (2021). *Guidelines for businesses to comply with the Maryland Personal Information Protection Act*. <https://www.marylandattorneygeneral.gov/Pages/IdentityTheft/businessGL.aspx>. Accessed 14 Apr 2021.
- Greenberg, J. (2011). Organizational justice: The dynamics of fairness in the workplace. In S. Zedeck (Ed.), *APA handbook of industrial/organizational psychology*, *3* (pp. 271–328). APA.
- Guzzo, R., Fink, A., King, E., Tonidandel, S., & Landis, R. (2015). Big data recommendations for Industrial-Organizational Psychology. *Industrial and Organizational Psychology*, *8*, 491–508.
- Hamilton, R., & Sodeman, W. (2020). The questions we ask: Opportunities and challenges for using big data analytics to strategically manage human capital resources. *Business Horizons*, *63*, 85–95.
- Heath, T. (2016). This employee ID badge monitors and listens to you at work — except in the bathroom. *Washington Post*. <https://www.washingtonpost.com/news/business/wp/2016/09/07/this-employee-badge-knows-not-only-where-you-are-but-whether-you-are-talking-to-your-co-workers>. Accessed 24 Mar 2018.
- Hewlett Packard Enterprises. (2016). *Real-world applications of video analytics*. <https://www.youtube.com/watch?v=72-R090y0pY>. Accessed 26 Dec 2018.
- IBM. (2015). *IBM intelligent video analytics overview*. <https://www.youtube.com/watch?v=fUKpGLk9MI8>. Accessed 26 Dec 2018.
- IBM. (2020). *What is machine learning?* <https://www.ibm.com/cloud/learn/machine-learning>. Accessed 21 Mar 2021.
- Illingworth, A. (2015). Big data in I-O Psychology: Privacy considerations and discriminatory algorithms. *Industrial and Organizational Psychology*, *8*, 567–575.
- Konovsky, M. (2000). Understanding procedural justice and its impact on business organizations. *Journal of Management*, *26*, 489–511.

- Krotov, V. (2017). The Internet of Things and new business opportunities. *Business Horizons*, 60, 831–841.
- LaForgia, M., Confessore, N., & Dance, G. (2018). Facebook rebuked for failing to disclose data-sharing deals. *New York Times*, December 18. <https://www.nytimes.com/2018/12/19/technology/facebook-data-privacy-criticism.html>. Accessed 16 Nov 2019.
- Levenson, A. (2018). Using workforce analytics to improve strategy execution. *Human Resource Management*, 57, 685–700.
- Loi, R., Lam, L., & Chan, K. (2012). Coping with job insecurity: The role of procedural justice, ethical leadership and power distance orientation. *Journal of Business Ethics*, 108, 361–372.
- Marler, J., & Boudreau, J. (2017). An evidence-based review of HR analytics. *The International Journal of Human Resource Management*, 28, 3–26.
- Mathworks. (2020a). *What is deep learning? 3 things you need to know*. <https://www.mathworks.com/discovery/deep-learning.html>. Accessed 21 Mar 2021.
- Mathworks. (2020b). *What is machine learning? 3 things you need to know*. <https://www.mathworks.com/discovery/machine-learning.html>. Accessed 21 Mar 2021
- McAbee, S., Landis, R., & Burke, M. (2017). Inductive reasoning: The promise of big data. *Human Resource Management Review*, 27, 277–290.
- McIver, D., Lengnick-Hall, M., & Lengnick-Hall, C. (2018). A strategic approach to workforce analytics. *Business Horizons*, 61, 397–407.
- Minbaeva, D. (2018). Building credible human capital analytics for organizational competitive advantage. *Human Resource Management*, 57, 701–713.
- National Conference of State Legislatures. (2019). *State social media privacy laws*. <http://www.ncsl.org/research/telecommunications-and-information-technology/state-laws-prohibiting-access-to-social-media-username-and-passwords.aspx>. Accessed 19 Nov 2019.
- Nyberg, A., & Ployhart, R. (2013). Context-emergent turnover: A theory of collective turnover. *Academy of Management Review*, 38, 109–131.
- Office of the Attorney General. (2021). *California Consumer Privacy Act (CCPA)*. <https://www.oag.ca.gov/privacy/ccpa>. Accessed 14 Apr 2021.
- Pande, A. (2017). *Future of video analytics with cloud, big data and machine learning*. <https://www.youtube.com/watch?v=RK5Uyeg0DUE>. Accessed 26 Dec 2018.
- Peck, D. (2013). *They're watching you at work*. <https://www.theatlantic.com/magazine/archive/2013/12/theyre-watching-you-at-work/354681>. Accessed 1 Oct 2017.
- Porter, M., & Heppelmann, J. (2014). How smart, connected products are transforming competition. *Harvard Business Review*, 92, 64–88.
- Rawls, J. (1971). *A theory of justice*. Belknap Press of Harvard University Press.
- Reuters. (2018). *Amazon ditched AI recruiting tools that favored men for technical jobs*. <https://www.theguardian.com/technology/2018/oct/10/amazon-hiring-ai-gender-bias-recruiting-engine>. Accessed 10 Jan 2019.
- Reuters. (2019). *Facebook's Zuckerberg grilled in U.S. Congress on digital currency, privacy, and U.S. elections*. <https://www.reuters.com/article/us-facebook-congress/facebooks-zuckerberg-grilled-in-u-s-congress-on-digital-currency-privacy-elections-idUSKBN1X2167>. Accessed 16 Nov 2019.
- RSIP Vision. (2021). *Exploring deep learning and CNNs*. <https://www.rsipvision.com/exploring-deep-learning/>. Accessed 10 Mar 2021.
- Sackett, P. (2002). The structure of counterproductive work behaviors: Dimensionality and relationships with facets of job performance. *International Journal of Selection and Assessment*, 10, 5–11.
- Saarikko, T., Westergren, U., & Blomquist, T. (2017). The Internet of Things: Are you ready for what's coming? *Business Horizons*, 60, 667–676.
- Serenko, A., & Bontis, N. (2016). Understanding counterproductive knowledge behavior: Antecedents and consequences of intra-organizational knowledge hiding. *Journal of Knowledge Management*, 20, 1199–1224.
- Shell, E. R. (2018). *The employer-surveillance state*. <https://www.theatlantic.com/business/archive/2018/10/employee-surveillance/568159>. Accessed 10 Apr 2019.
- Sheng, E. (2019). *Employee privacy in the US is at stake as corporate surveillance technology monitors every move*. <https://www.cnbc.com/2019/04/15/employee-privacy-is-at-stake-as-surveillance-tech-monitors-workers.html>. Accessed 27 Apr 2019.
- SHRM. (2014). *Code of ethics*. <https://www.shrm.org/about-shrm/pages/code-of-ethics.aspx>. Accessed 16 Feb 2021.
- Spitznagel, E. (2019). *Inside the hellish workday of an Amazon warehouse employee*. <https://nypost.com/2019/07/13/inside-the-hellish-workday-of-an-amazon-warehouse-employee/> Accessed 10 Mar 2021.

- Stewart, M. (2019). *Handing discriminatory biases in data for machine learning*. <https://towardsdatascience.com/machine-learning-and-discrimination-2ed1a8b01038>. Accessed 16 Apr 2021.
- Stone, M. (2020). *IoT security explained*. <https://cybersecurity.att.com/blogs/security-essentials/internet-of-things-security-explained>. Accessed 21 Mar 2021.
- Stringfellow, A. (2019). *The 52 best mobile apps for human resource (HR) managers*. <https://www.wonolo.com/blog/best-hr-apps/>. Accessed 20 Dec 2019.
- Thibodeaux, W. (2017). *This artificial intelligence can predict how you'll behave at work based on social media*. <https://www.inc.com/wanda-thibodeaux/this-artificial-intelligence-can-use-social-media-to-tell-hiring-managers-about-your-personality.html> Accessed 16 Nov 2019.
- Thompson, S., & Warzel, C. (2019). *Twelve million phones, one data set, no privacy*. <https://www.nytimes.com/interactive/2019/12/19/opinion/location-tracking-cell-phone.html>. Accessed 20 Dec 2019.
- Tomczak, D., Lanson, L., & Aguinis, H. (2018). Evidence-based recommendations for employee performance monitoring. *Business Horizons*, *61*, 251–259.
- Turing, A. (1950). Computing machinery and intelligence. *Mind*, *50*(236), 433–460.
- Walsh, D. J. (2019). *Employment law for human resource practice* (6th ed.). Cengage.
- Warren, S., & Brandeis, L. (1890). The right to privacy. *Harvard Law Review*, *4*, 193–220.
- Waters, S., Streets, V., McFarlane, L., & Johnson-Murray, R. (2018). *The practical guide to HR analytics*. Society for Human Resource Management.
- Zimmerman, L., & Clark, M. (2016). Opting-out and opting-in: A review and agenda for future research. *The Career Development International*, *21*, 603–633.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.