“The addition of a new class of worker, driven by AI, promises to challenge the path to greater inclusion by having the potential to exponentially increase disruption not just in organizations but in society, government, and our everyday lives.”

AI × I = AI²

The OD Imperative to Add Inclusion to the Algorithms of Artificial Intelligence

By Frederick A. Miller, Judith H. Katz, & Roger Gans

Since its beginnings, one of the functions of OD has been to create organizations that enable people to do their best work and to create workplaces built on principles of democracy and participation. A major element of creating such workplaces is identifying and ameliorating discriminatory practices and cultures in organizations. Artificial intelligence (AI) is in the process of complicating and confounding that function in ways we may not have seen coming. There are growing concerns about human (and other) biases being built into the machine-learning algorithms that are increasingly impacting our organizations, their processes, and our lives. But just as AI has the potential to reify and magnify the effects of human bias, it also offers unprecedented opportunity to build inclusive practices into the fundamental practices and processes of organizations.

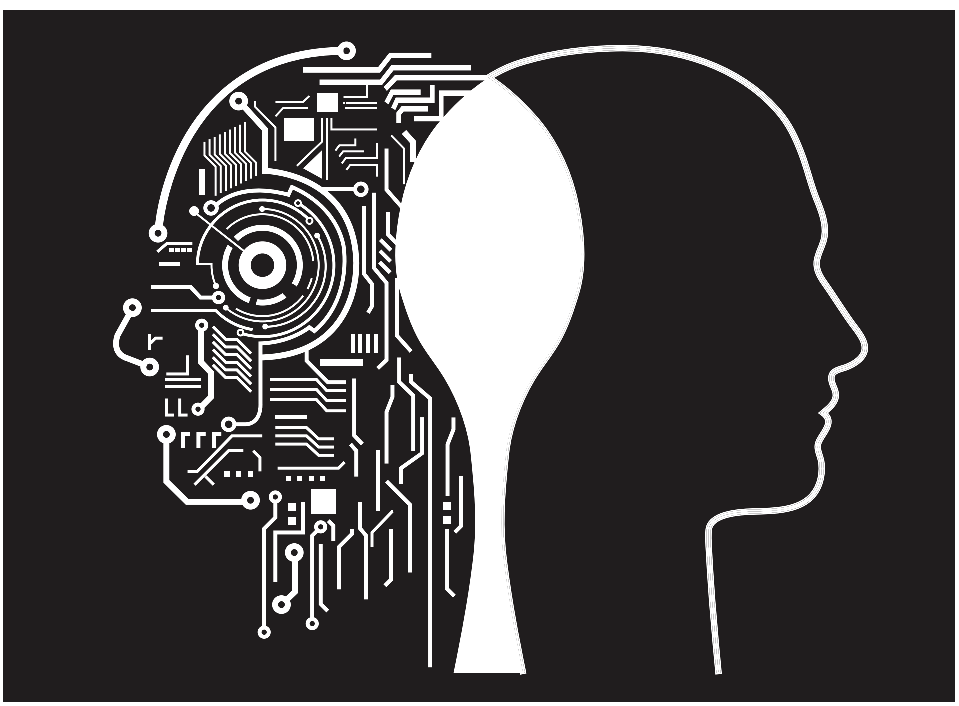

As the following will show, it is clear that responsible AI developers must find ways to incorporate awareness of the potential for bias and the value of inclusion into the algorithms that guide machine learning processes. But our experience suggests that if the developers of AI systems hope to eliminate discrimination and build inclusion into their software, they first will need to do those things with their own culture. In addition, those who are working in organizations need to be mindful of the potential for bias in such processes and software so they can insure the processes being implemented are not contributing to biases that may already exist within the workplace. In this article we discuss some of the dangers and opportunities presented by AI, and the implications for organizations, the people of those organizations, and OD practitioners tasked with assisting them to survive and thrive.

A New Class of Worker Brings Danger and Opportunity

Organizations have been, and continue to be, disrupted and transformed by the addition of women, people of color, people from different countries and ethnic origins, people with different sexual orientation and gender identities, and differently-abled people into the workforce and workplace. Organizations that have learned to leverage the added skillsets and perspectives of their increasingly diverse workforces through building cultures of inclusion have experienced significant gains in productivity and profitability (Katz & Miller, 2017; Miller & Katz, 2002; Page, 2007). The addition of a new class of worker, driven by AI, promises to challenge the path to greater inclusion by having the potential to exponentially increase disruption not just in organizations but in society, government, and our everyday lives.

Robots and other machines powered by computerized algorithms are already working alongside humans in factories around the world. Some, with self-programming machine-learning capabilities, are performing customer service functions, implementing marketing strategies, and making consequential decisions that can determine the opportunities we see, the jobs we get, the products we buy, the prices we pay, and the treatment we receive from officers of the law and the courts. Robots are the visible manifestations of artificial intelligence—the hands and feet of AI. Many of the manifestations and influences of AI are less visible, however, and some of these are proving to be problematic.

What is Artificial Intelligence Learning from Humans?

Although feared by some, the great hope of many people was that AI would give us faster, wiser, fairer decisions and actions without the downsides of human error, fatigue, or bias. Through the magic of machine learning, it would speed customer service transactions, unstick the gridlock of governmental and organizational bureaucracies, eliminate traffic jams, improve medical diagnoses and treatments, and relieve us of the burden of countless boringly repetitive tasks.

But who is teaching the machine? And once activated, what will the machine teach itself and other machines, especially if what it learns is based on human history, the content of the Internet, and the biases, fears, and unexamined assumptions of its coders, programmers, and model builders? Many OD practitioners are trained to identify manifestations of bias, oppression, and discrimination in organizational systems and culturally influenced data. But the program developers who write the algorithms that drive the machines rarely receive such training (Mundy, 2017). Without such knowledge, they can overlook the danger that the data used to inform the AI machine-learning process may have culturally determined biases already baked in. For example, AI-driven risk-assessment tools currently in use in some places sift through racially biased arrest records and historical crime data to help courts make decisions and police departments determine which neighborhoods should receive greater scrutiny and coverage. In doing so, they are actively reflecting, perpetuating, and magnifying racial inequities caused by societal prejudice (Crawford, 2016).

Bias Is Baked into the Data

It is too late to merely worry that human biases might cross over into the computerized programs affecting many individual lives and organizational functions. Our biases are baked right into our language and the language-usage data AI systems learn from (Caliskan, Bryson, & Narayanan, 2017). To cite a readily observable phenomenon, AI-driven language translation tools routinely add gendered stereotypes in translating from gender-neutral languages:

> Google Translate converts these Turkish sentences with gender-neutral pronouns: "O bir doctor. O bir hemşire." to these English sentences: "He is a doctor. She is a nurse." We see the same behavior for Finnish, Estonian, Hungarian, and Persian in place of Turkish. Similarly, translating the above two Turkish sentences into several of the most commonly spoken languages (Spanish, English, Portuguese, Russian, German, and French) results in gender-stereotyped pronouns in every case (Caliskan et al., 2017).

In 2015, Google's photo app—powered by AI and machine learning processes—identified black people in some photos as gorillas (Barr, 2015). That same year, a Carnegie Mellon University study determined that AI-driven, search-based advertising promising employment assistance for obtaining high-paying jobs—for $200,000 and higher—targeted significantly fewer women than men (Spice, 2015).

In bail and sentencing hearings in courtrooms across the U.S., AI-driven software systematically—and mistakenly—rates black people as higher recidivism risks than white people (Angwin, Larson, Mattu, & Kirchner, 2016). Based on AI-driven calculations, insurance companies routinely charge residents of zip codes with large minority populations up to 30% more than residents from whiter neighborhoods with similar accident costs (Angwin, Larson, Kirchner, & Mattu, 2017).

Outcomes like these violate our expectations. We assume machines must be inherently fair and objective, that they cannot help but analyze data without bias or malice. But it is easy to forget that the programming that drives the way AI analyze data is originally created by humans. The people who create the algorithms belong to an industry culture that has bias against women and African Americans, even if based solely on their conspicuous underrepresentation (Clark, 2016; Mundy, 2017). Undoubtedly, few programmers would intentionally embed bias in their work, but it is hard to address problems you do not see, and impossible to avoid doing things you do not even know you are doing. Racist and sexist assumptions are ingrained in the wider societal culture, and perhaps even more so in the tech industry subculture (Mundy, 2017; Tiku, 2017).

Computers Learn Bias the Same Way People Do

Machine learning is a process by which computers sift through and process enormous amounts of data with a goal of identifying underlying patterns in the data, which is basically the same way humans learn (Emspak, 2016). In both cases, the results are most often used to predict future actions and behaviors. For early human learning, the prediction can involve what kinds of vocalizations and facial expressions are most likely to elicit a hug, food, or a diaper-change. For a machine-learning computer, the prediction is likely to involve which humans to target for product advertising and which advertising messages are most likely to produce sales, but it can also involve who to loan money to, who to hire, who to promote, and who are the greatest risks for committing crimes or appearing for trials.

Humans start processing data as infants, and we learn the expectations of our society from the actions and words of all the people with whom we come into contact. If there are biases in our upbringing, we can sometimes learn to overcome them if we consciously decide to do so. We can learn to identify patterns of unfairness and discrimination in other people's attitudes and behavior, and we can seek out additional sources of information to fact-check biased claims and act to correct them. But with up to 98% of our own attitudes and decisions arrived at through unconscious processes, it is harder to identify the biases we hold implicitly (Staats, Capatosto, Wright, & Jackson, 2016). Without training and vigilance, AI programmers and model-builders cannot help but perpetuate these implicit, unconscious biases in their work.

In machine learning, computers can only process the data they receive, and they may be restricted to considering only specific facets of that data as part of their initial human-sourced programming. Add in the fact that virtually all data available for analysis, including language itself, has roots in human perception and interpretation, and it becomes clear that bias in machine learning is inevitable. Like a child, a machine-learning computer builds its vocabulary and "intelligence" through pattern recognition (Bornstein, 2016)—for instance, in how often terms and value judgments appear together on the Internet and other sources (Caliskan, et al., 2017). The word "nurse" is vastly more often accompanied by female gendered pronouns than by male gendered pronouns. African-American names are often surrounded by words that connote unpleasantness because people on the Internet say awful things, not because African Americans are unpleasant.

Companies that offer AI services to other companies may tout the speed and capability of their processes, but unless they offer transparency in the development of their algorithms and the training of their people, there is no way for their client organizations to know if the AI package includes baked-in biases. OD practitioners working to eliminate institutionalized "isms" in organizational interactions and systems need to be aware of the potential of AI to institutionalize those "isms" in ways that are much harder to detect, challenge, or change.

Prejudices produce actions that, in turn, produce data. For instance, it is widely acknowledged that arrest and incarceration data reflect societal biases against people of color, a pattern that is readily seen in the way drug laws have been enforced. While whites and African Americans are equally likely to use illegal drugs (Lopez, 2015), African Americans are roughly three times as likely to be arrested and prosecuted for possession of illegal drugs (Common Sense for Drug Policy, 2014). A similar skewing of "objective" data can be seen in percentages of women serving on corporate boards and in senior management positions (Warner, 2014). Without specific instructions to consider these kinds of patterns as evidence of bias, machine-learning computers are likely to use these data to predict that African Americans are three times as likely as whites to be carrying illicit drugs (which can be used as a justification for racial profiling and stop-and-frisk practices), and that women lack certain leadership qualities.

Don't Ask, Because We Can't Tell

Because machines are assumed to be fair and unbiased, machine-produced predictions, and the resulting recommendations and decisions, are less likely to be questioned as biased than if they had come from human agents (The AI Now Report, 2016). Not only is it less likely a machine's decision will be questioned, its decision is also significantly harder to question than a human's. AI-developers such as Google and Amazon consider their algorithms to be proprietary information, and they protect them vigorously. Moreover, particularly in advanced machine-learning systems, the details of any individual prediction may be based on literally billions of individual digital processes and, as such, are opaque even to the original coders (Bornstein, 2016; Knight, 2017). In other words, while humans may be asked to account for and justify what seem like biased decisions, machines may not be able to provide such explanations—and neither will their creators.[1]

Companies that offer AI services to other companies may tout the speed and capability of their processes, but unless they offer transparency in the development of their algorithms and the training of their people, there is no way for their client organizations to know if the AI package includes baked-in biases. OD practitioners working to eliminate institutionalized "isms" in organizational interactions and systems need to be aware of the potential of AI to institutionalize those "isms" in ways

1. The European Union's General Data Protection Regulation (GDPR), which goes into effect in May 2018, is meant to protect the right of individuals to know how their personal data is used. There is a view that the GDPR includes a "right of explanation" as to how outputs are generated from machine learning models. If true, companies that are building these models may need to demonstrate that they have removed bias from those outputs. More information is available at: http://www.eugdpr.org/

that are much harder to detect, challenge, or change.

Bias In, Bias Out: Coder Culture Resists Change

As detailed above, AI-driven decision-making processes can produce biased outcomes that reflect the same sets of "isms" OD practitioners and others have been working to ameliorate for decades. The evidence suggests that if the biases exist in the wider society, they will be "learned" by AI systems that use the collective behavior and data of the wider society to learn from.

This would be less of a problem if the programmers writing the algorithms on which machine-learning systems run were more aware of the biases that exist in the wider society, and by extension, in the data sets produced by that society. Greater awareness would make them better able to ensure their coding efforts include strategies for identifying patterns of bias in societally-influenced data and safeguards against existing, documented biases. Making such awareness more normative within the tech industry will be a challenging undertaking. Of course, as might be expected in the tech industry, "there's an app for that," with a proliferation of anti-bias apps and training workshops that try to reduce bias itself to an algorithm. But there continues to be unwillingness among some tech companies to change core parts of their culture (Mundy, 2017).

Celebration of the tech industry's coding community as an elite, exclusive, meritocratic club seems to be a deeply entrenched ethos, sometimes defended by claims that the sparse numbers of women and African Americans are a consequence of a reluctance to "lower our standards" (Mundy, 2017). Racial stereotyping is a well-acknowledged problem within the software industry (Tiku, 2017). Gender stereotyping, in contrast, seems to attract more attention as well as greater pushback when attempts are made to address it (Wakabayashi, 2017). In recent years, the tech industry has produced an increasing number of reports on their companies' diversity numbers, but little in the way of positive change in those numbers or the cultures that have produced and sustained them. Studies have shown that women leave the tech industry at twice the rate that men do, and that the percentage of computer science degrees earned by women has decreased from 37% in 1984 to 18% in 2014 (Alba, 2017). Some diversity education programs at tech companies have seemed to produce boomerang effects, with declines in diversity at some of the companies in which such programs were enacted (Alba, 2017).

Not Just a Tech Issue: AI's Expanding Presence

It may seem tempting to focus warnings about bias and discriminatory implications of AI solely on the tech industry, but AI-driven processes and services are already part of the routine experience of everyday life inside organizations of all sizes in all industries. (How many times have you Googled something today?) In fact, people in organizations outside the tech industry are even less likely to question the algorithms and machine-logic on which AI-influenced decisions are made than within the tech industry. Without a keen awareness of the potential for baked-in bias in their AI-driven systems, some organizations are at risk of inadvertently becoming party to actions that have a discriminatory effect on their customers or their team members, with potentially dire bottom-line consequences in either case. This may already be influencing hiring practices, in which AI is increasingly used in talent sourcing and acquisition. AI is being used to make the candidate-selection process faster and more efficient (Alsever, 2017), and to root out human biases (Captain, 2016), but because it relies on human-programmed choice trees and human-generated data in deciding which candidates are the best "fits," the process also can rule out some of the diversity organizations are—or ought to be—seeking (Ghosh, 2017).

There is an upside in all this for those seeking to address issues of inclusion and diversity in organizations, however. The potential for bias in AI systems can actually be a useful tool for OD practitioners. By raising concerns about machine-based biases in organizational practices, we may also be able to raise awareness of how unconscious bias is carried like an "equal opportunity virus" (Dasgupta, 2013) by all the humans of the organization. Considering its effects, of course, bias might be more accurately considered an "un-equal opportunity virus."

AI and OD: What's Around the Corner

The rippling effects of AI promise to impact virtually all facets of organizational life, from decisions about who to hire and promote, to design and marketing of products and services, to each organization's competitive position and reputation in the global marketplace. Instead of disregarding it as too technical for our purview, OD practitioners need to see AI as a critical element of the organization that needs to be analyzed and addressed in regard to its effects on institutionalized "isms" and people's ability to do their best work.

There are more implications for the role of OD in addressing issues of AI than can be covered in any single article. Some of the AI-related issues OD practitioners should anticipate facing include:

AI-fueled entrepreneurship. As access to the tools of AI becomes more widespread, it is likely to spur the growth of entrepreneurial start-ups that focus on applying the potential of AI to solve an ever-widening array of personal and commercial needs (Lee, 2017). The role of OD will be critical in assisting these start-ups to avoid the toxic-culture missteps of tech start-ups like Uber (Noguchi, 2017) and SoFi (O'Connor, 2017).

Worker disruption and displacement. Robots powered by AI-systems are already replacing people in manufacturing plants, warehouses, banks, and supermarkets throughout the world. Other types of jobs will inevitably be replaced or displaced as AI systems become more sophisticated. Challenges for the practice of OD are likely to include working to create a culture that enables people to work effectively with robots and advanced AI: *How will workers* *react and relate to non-human co-workers? Will work teams accept an AI as a teammate or an agent of management?* OD practitioners will almost certainly need to prepare the organization and its people for widespread role-changes and potentially stressful rounds of outplacement and downsizing. The shapes of the changes to come are difficult to predict, but preparing organizations and the people in them for inevitable and increasingly rapid AI-related change is a necessity.

For practitioners of OD, the challenge will be not just to assist organizations to recognize and address the inherent dangers presented by AI, but also to recognize the potential of AI to integrate inclusive algorithms into the fabric of their existence. Today, our task is to identify and root out the biases and inequities of human society that are being absorbed through machine-learning processes and presented as objective and unquestionable reality. This is no small task! However, we would be remiss if we did not also address the positive potential of AI.

How to Add Inclusion to the AI Algorithm: AI x I = AI²

To address the issue of bias in AI, it will be essential to address the culture of the coders as well as the code. Following are a few suggestions for changing the culture of the tech industry to be more inclusive and more aware of the potential for bias in its members and their code.

A strategy for creating culture change within tech organizations and among coders. Before AI model builders—and those working in partnership with them—can be expected to root out biases and inequities from their algorithms and AI-based products, they will need the competence and capability to recognize and address those biases and inequities. They will also need to accept that those biases and inequities are real, harmful, and consequential. Any efforts to address the prevailing practices and mindsets of the tech industry in this regard must start with awareness that some aspects of coder-culture have deep-seated resistance to change, as noted above. The following elements might better position such a culture-change strategy for success.

Education. This may be an occasion to brandish Churchill's "those who fail to learn from history are doomed to repeat it." Claims regarding "lowering our standards" were exposed decades ago as pretexts for excusing the exclusion of women, people of color, and other undesirables (Cross, Katz, Miller, & Seashore, 1994). It will be vital to help those involved with AI to gain greater competence in recognizing bias in themselves and societally produced data. Although many organizations are doing education/training on unconscious bias, that alone will not solve this issue. It has to go beyond personal awareness to scrutiny of how the data itself may be reflecting biases and therefore to reimagine how to use AI's data-crunching abilities to avoid perpetuating longstanding patterns of discrimination.

Education in this direction could include exposing tech industry members to evidence of their own biases, as well as documentation of biases in the data used in machine-learning applications. Motivation for change could be addressed with additional education regarding the value-added and return-on-investment of inclusive practices (e.g., Katz & Miller, 2017; Page, 2007) as well as the costs of bias-centered lawsuits and public relations disasters.

Socialization. People cannot adopt a cultural norm of inclusive behaviors until they experience that norm. To accomplish this, it will be necessary to establish pilot groups that practice and model inclusive mindsets and actions, and to nurture these groups with education and organizational support. Ideally, these pilot groups will grow and eventually form the core of each organization's new culture.

Certification. A program that requires and provides certification of competence for recognizing bias and practicing inclusive behaviors seems a particularly apt accountability tool for the software industry. AI programmers could be required to pass multicultural competence tests or attend education programs that address bias, diversity, inclusion, and the practice of self-as-instrument. They might also undergo periodic recertification processes that could include 360-degree reviews from a diverse group including their team leaders, colleagues, and direct reports.

A strategy for overseeing code quality and addressing grievances. Because of the specialized nature of this field, few people possess the competence to recognize defects or flaws in computer programs, and fewer can trace potential problems with the deep processes involved in machine learning. This has created problems with regard to accountability and redress of issues that affect people's lives and livelihoods, and suggests a need for creation of at least two sets of human-staffed resources:

Organizational and industry-wide peer-review boards. To protect the integrity of the organizations producing the code, there needs to be a process for some of the AI-based products to have their code (and the results of pilot runs for deep-process machine-learning applications) reviewed by an independent diverse panel of experts before being released into the public sphere.

Organizational and industry-wide AI grievance panels. It should be assumed that AI applications will produce unexpected and unintended inequities. Each organization that produces AI-based products could establish a standing panel to address grievances from consumers and others affected by their products, either directly or indirectly. For consumers who are not satisfied with the redress given them by the manufacturing organization, there could be an industry-wide appeals panel that would hold organizations accountable.

A strategy that requires immediate action. Regardless of the industry, OD practitioners cannot wait for a world-changing robot apocalypse to sound the alarm or to start addressing the issues of AI. We need to be mindful that this is happening now, and at a pace that is accelerating. We cannot settle for a "let the buyer beware" market for AI products. We must enable the buyer to beware when the organizations we support are purchasing such products until we are sure anti-bias safeguards are in place and the awareness of the programmers and sellers is at a level that they have made their products "safe" for our diverse world.

We need to be willing to get into the messy work of understanding how bias is being built into these systems. We need to be willing to venture outside our comfort zones in questioning the fitness and objectivity of algorithms we may not have the technological savvy to understand, but whose biased effects we can and need to identify.

Conclusion: This is Just the Beginning

Whether you believe AI has the potential to create an Eden-like utopia (Lee, 2017) or bring about the extinction of humankind (Dowd, 2017) or something in between, it is clear that AI will exert greater and greater influence over virtually all aspects of individual and organizational life (The AI Now Report, 2016). For practitioners of OD, the challenge will be not just to assist organizations to recognize and address the inherent dangers presented by AI, but also to recognize the potential of AI to integrate inclusive algorithms into the fabric of their existence.

Today, our task is to identify and root out the biases and inequities of human society that are being absorbed through machine-learning processes and presented as objective and unquestionable reality. This is no small task! However, we would be remiss if we did not also address the positive potential of AI. Consider applying the power of AI to any of these "what ifs":

» *What if*, instead of equating data with purely objective facts, AI routinely identified patterns that could be the result of societal or organizational biases and discrimination, and sounded alarm bells?
» *What if*, instead of selecting only the job candidates who fit our existing organization profile, AI selected an array of candidates who provide the perspectives we currently lack?
» *What if*, instead of showing us only the news we are likely to be most interested in, AI showed us the news we most need to see to be well-rounded, responsible citizens?

These are the kinds of questions an inclusive, culturally competent AI coding and consuming community would ask about how AI could enhance human interaction. What the results might be, we can only imagine.

References

Alba, D. (2017, March 31). Hey tech giants: How about action on diversity, not just reports? *Wired*. Retrieved from *https://www.wired.com/2017/03/hey-tech-giants-action-diversity-not-just-reports/*

Alsever, J. (2017, May 19). How AI is changing your job hunt. *Fortune*. Retrieved from *http://fortune.com/2017/05/19/ai-changing-jobs-hiring-recruiting/*

Angwin, J., Larson, J., Kirchner, L., & Mattu, S. (2017, April 5). Minority neighborhoods pay higher car insurance premiums than white areas with the same risk. *ProPublica*. Retrieved from *https://www.propublica.org/article/minority-neighborhoods-higher-car-insurance-premiums-white-areas-same-risk*

Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016, May 23). Machine bias: There's software used across the country to predict future criminals. And it's biased against blacks. *ProPublica*. Retrieved from *https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing*

Barr, A. (2015, July 1). Google mistakenly tags black people as 'gorillas,' showing limits of algorithms. *The Wall Street Journal*, retrieved from *https://blogs.wsj.com/digits/2015/07/01/google-mistakenly-tags-black-people-as-gorillas-showing-limits-of-algorithms/*

Caliskan, A., Bryson, J.J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science, 356*, 183–186. DOI: 10.1126/science.aal4230 (Supplemental Materials: *www.sciencemag.org/*content/356/6334/183/suppl/DC1)

Captain, S. (2016, May 18). Can artificial intelligence make hiring less biased? *Fast Company*. Retrieved from *https://www.fastcompany.com/3059773/we-tested-artificial-intelligence-platforms-to-see-if-theyre-really-less-*

Clark, J. (2016, June 23). Artificial intelligence has a 'sea of dudes' problem. *Bloomberg Technology*, retrieved from *https://www.bloomberg.com/news/articles/2016-06-23/artificial-intelligence-has-a-sea-of-dudes-problem*

Common Sense for Drug Policy. (2014). "Race and Prison." *Drug War Facts*. Retrieved from *http://drugwarfacts.org/chapter/race_prison#sthash.WRkTtM10.dpbs*

Crawford, K. (2016, June 25). Artificial intelligence's White Guy problem. *The New York Times*, retrieved from *https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html*

Cross, E.Y., Katz, J.H., Miller, F.A., & Seashore, E.W. (Eds.) (1994). *The promise of diversity: Over 40 voices discuss strategies for eliminating discrimination in organizations*. Burr Ridge, IL: Irwin Professional Publishing.

Dasgupta, N. (2013). Implicit attitudes and beliefs adapt to situations: A decade of research on the malleability of implicit prejudice, stereotypes, and the self-concept. *Advances in Experimental Social Psychology, 47*, 233–279. *dx.doi.org/10.1016/B978-0-12-407236-7.00005-X*

Dowd, M. (2017, April). Elon Musk's billion-dollar crusade to stop the A.I. apocalypse. *Vanity Fair*, April 2017. Retrieved from *https://www.vanityfair.com/news/2017/03/elon-musk-billion-dollar-crusade-to-stop-ai-space-x*

Emspak, J. (2016, December 29). How a machine learns prejudice. *Scientific*

American, retrieved from *https://www.scientificamerican.com/article/how-a-machine-learns-prejudice/*
Ghosh, D. (2017, October 17). AI is the future of hiring, but it's far from immune to bias. *Quartz at Work*. Retrieved from *https://work.qz.com/1098954/ai-is-the-future-of-hiring-but-it-could-introduce-bias-if-were-not-careful/*
Katz, J.H., & Miller, F.A. (2017). Leveraging differences and inclusion pays off: Measuring the impact on profits and productivity. *OD Practitioner, 49*(1), 56–61.
Knight, W. (2017, April 11). The dark secret at the heart of AI: No one really knows how the most advanced algorithms do what they do. That could be a problem. *MIT Technology Review.* Retrieved from *https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/*
Lee, T.E. (2017, May 18). Artificial intelligence is getting more powerful, and it's about to be everywhere. *Vox*. Retrieved from *https://www.vox.com/new-money/2017/5/18/15655274/google-io-ai-everywhere*
Lopez, G. (2015, October 1). Black and white Americans use drugs at similar rates. One group is punished more for it. *Vox*. Retrieved from *https://www.vox.com/2015/3/17/8227569/war-on-drugs-racism*
Miller, F.A., & Katz, J.H. (2002). *The inclusion breakthrough: Unleashing the real power of diversity.* San Francisco, CA: Berrett-Koehler Publishers, Inc.
Mundy, L. (2017, April). Why is Silicon Valley so awful to women? *The Atlantic*. Retrieved from *https://www.theatlantic.com/magazine/archive/2017/04/why-is-silicon-valley-so-awful-to-women/517788/*
Noguchi, Y. (2017, June 6). Uber fires 20 employees after sexuial harassment claim investigation. *NPR*. Retrieved from *http://www.npr.org/sections/thetwo-way/2017/06/06/531806891/uber-fires-20-employees-after-sexual-harassment-claim-investigation*
O'Connor, C. (2017, September 12). SoFi CEO Mike Cagney resigns following sexual harassment lawsuit. *Forbes*. Retrieved from *https://www.forbes.com/sites/clareoconnor/2017/09/12/sofi-ceo-mike-cagney-resigns-following-sexual-harassment-lawsuit/#6847d9b565be*
Page, S.E. (2007). The empirical evidence. In S.E. Page, *The difference: How diversity creates better groups, firms, schools, and societies* (pp. 313–337). Princeton, NJ: Princeton University Press.
Resnick, B. (2017, April 17). How artificial intelligence learns to be racist. *Vox,* retrieved from *https://www.vox.com/science-and-health/2017/4/17/15322378/how-artificial-intelligence-learns-how-to-be-racist*
Spice, B. (2015, July 7). Questioning the fairness of targeting ads online: CMU probes online ad ecosystem. *Carnegie Mellon University News*, retrieved from *http://www.cmu.edu/news/stories/archives/2015/july/online-ads-research.html*
Staats, C., Capatosto, K., Wright, R.A., & Jackson, V. W. (2016). *Implicit bias review, 2016 edition*. Ohio State University: Kirwan Institute for the Study of Race and Ethnicity. Retrieved from *http://kirwaninstitute.osu.edu/my-product/2016-state-of-the-science-implicit-bias-review/*
The AI Now Report. (2016, September 22). The social and economic implications of artificial intelligence technologies in the near-term. *AI Now* (Summary of public symposium). Retrieved from *https://artificialintelligencenow.com/media/documents/AINowSummaryReport_3_RpmwKHu.pdf*
Tiku, N. (2017, October 3). Why tech leadership has a bigger race than gender problem. *Wired*. Retrieved from *https://www.wired.com/story/tech-leadership-race-problem/*
Wakabayashi, D. (2017, August 7). Google fires engineer who wrote memo questioning women in tech. *The New York Times*. Retrieved from *https://www.nytimes.com/2017/08/07/business/google-women-engineer-fired-memo.html*
Warner, J. (2014, March 7). Fact sheet: The women's leadership gap. *Center for American Progress*. Retrieved from *https://www.americanprogress.org/issues/women/reports/2014/03/07/85457/fact-sheet-the-womens-leadership-gap/*

Frederick A. Miller and Judith H. Katz are CEO and Executive Vice President (respectively) of The Kaleel Jamison Consulting Group, Inc., one of *Consulting Magazine*'s Seven Small Jewels in 2010. They have partnered with Fortune 50 companies globally to elevate the quality of interactions, leverage people's differences, and transform workplaces. Katz sits on the Dean's Council, College of Education at the University of Massachusetts, Amherst, and the Board of Trustees of Fielding Graduate University. Miller serves on the boards of Day & Zimmermann, Rensselaer Polytechnic Institute's Center for Automated Technology Systems, and Hudson Partners. Both are recipients of the OD Network's Lifetime Achievement Award and have co-authored several books, including *Opening Doors to Teamwork* and *Collaboration: 4 Keys that Change EVERYTHING* (Berrett-Koehler, 2013) as well as a book on workplace psychological and emotional safety, to be published in Fall 2018. Miller can be reached at *fred411@kjcg.com*. Katz can be reached at *judithkatz@kjcg.com*.

Roger Gans, MA, ABD, is a writer, consultant, and educator who specializes in strategic communication. He has been a long-time thinking and writing partner of Miller, Katz, and KJCG. An adjunct professor in the management and communication departments of the Sage Colleges, his doctoral dissertation examines how pro-social advocacy campaigns can exacerbate engagement disparities in civic affairs, health care, and the workplace. His current consulting projects include promoting health care services on Eastern Long Island (NY) and development of a youth addiction services program in Iowa. Gans can be reached at *rgans@albany.edu*.