

When is algorithmic discrimination wrong?

Karoline Reinhardt 

Abstract: This paper investigates the moral implications of algorithmic discrimination from a philosophical perspective, focusing on the concept of discrimination as delineated by Deborah Hellman and extending the analysis through insights from Immanuel Kant's moral philosophy. First, the paper establishes that algorithms have to discriminate in a descriptive sense. Highlighting their ubiquitous presence and profound impact on various aspects of life, it is argued for the relevance of morally wrong algorithmic discrimination. Drawing from Hellman's framework, discrimination is understood as morally problematic when it is demeaning, irrespective of whether the affected individuals perceive it as such. This approach is particularly relevant to algorithmic discrimination, where individuals may not be aware of their unequal treatment due to the opacity of algorithmic processes. Moreover, the paper addresses the role of intentionality in discrimination, arguing that algorithms, as problem-solving structures, operate without intentionality. Emergent discrimination, a phenomenon observed in algorithmic systems, further underscores the importance of understanding discrimination beyond intentional acts. By invoking Kant's notion of respect owed to others, the paper argues that algorithmic discrimination is morally wrong when it makes demeaning distinctions, as respect is owed to every individual. In conclusion, the paper advocates for a nuanced understanding of algorithmic discrimination, drawing from philosophical frameworks to differentiate morally unproblematic from problematic distinctions.

Keywords: algorithms, discrimination, Hellman, Kant

1. Introduction

¹ According to Article 2 of the Universal Declaration of Human Rights, everyone has the right to

all the rights and freedoms set forth in this Declaration, without distinction of any kind, such as race, colour, sex, language, religion, political or other opinion, national or social origin, property, birth or other status. This article is often referred to as the "prohibition of discrimination."

2 However, as is often the case in philosophy, a notion that seems uncontroversial at first glance can lead to greater systematic difficulties upon closer examination: Is discrimination always group-based? Are these the only categories under which someone can be discriminated against? Are these categories not historically contingent – those categories under which discrimination has occurred and continues to occur – but do they tell us anything conceptually or systematically about what discrimination actually is? Another problem that arises is who can actually discriminate against someone: the state and state institutions, individuals. Does the notion of discrimination require an intention to discriminate against somebody?¹

3 All of these questions take on a new face when we consider what has been termed algorithmic discrimination since the advent of the digital age.

4 In academic debate, the topic of algorithmic discrimination is discussed by many disciplines. In addition to approaches from computer science, there are particularly legal² and sociological³ analyses. I will consider the topic here from a decidedly philosophical, primarily moral philosophical perspective: The "wrong" in the title is thus to be understood as a moral "wrong." In my argumentation in algorithm ethics as a subfield of applied ethics, I will start, but also draw on arguments from political ethics and the history of moral philosophy and make them fruitful for the current question of algorithmic discrimination.

5 In the literature, when it comes to algorithmic discrimination, you will find descriptions of specific cases of algorithmic discrimination,⁴ problematic definitions of target variables, the role of individual bias of developers,⁵ training data quality and data distortion,⁶ or labeling processes.⁷ What I want to present here is an approach that determines when exactly distinctions implemented in algorithmic systems or made by them are morally wrong.

6 One might ask: Isn't discrimination always wrong? Upon closer examination, however, it quickly becomes apparent that many forms of discrimination (in a certain sense of the word) are (indeed) morally unproblematic, some are even morally required. At the same time, determining what distinguishes these required and unproblematic discriminations from morally problematic discriminations is not an easy task.

7 Following this introduction (section 1), I will demonstrate that algorithms always have to dis-

1. For a comprehensive overview of the current debate, see the articles in ([Lippert-Rasmussen 2017](#)).

2. For instance ([Citron and Pasquale 2014](#)), ([Crawford and Schultz 2014](#)) or ([Hellman 2020](#)).

3. For instance ([Hagendorff 2019](#)).

4. For instance ([Eubanks 2018](#)), ([O'Neil 2016](#)).

5. For instance ([Heesen, Reinhardt, and Schelenz 2021, 135](#)).

6. For instance ([Heesen, Reinhardt, and Schelenz 2021, 134](#)), ([Friedman and Nissenbaum 1996](#)).

7. Cf. ([Reinhardt 2020, 272-274](#)).

criminate in a certain respect. I will call this the descriptive meaning of the notion of discrimination (section 2). After that I will turn to the question of how we can distinguish between morally problematic discrimination and morally unproblematic discrimination (section 3). For this, I will first build on an approach by Deborah Hellman ([Hellman 2008](#)). Hellman argues that discrimination is morally problematic when it is demeaning. I will apply her understanding of morally problematic discrimination to phenomena of algorithmic discrimination and highlight the particular strengths of this approach. However, this approach also, lacks an explanation of why demeaning is immoral, as I will show. I will sketch how we might fill this gap by referring to Immanuel Kant's considerations on the respect owed to others as formulated in the *Metaphysics of Morals* (section 4). Finally, I will summarize the results (section 5).

2. Why Algorithms Must Discriminate

8 What do I mean when I say that algorithms must discriminate? This question can be answered by looking at what an algorithm actually is.

2.1 What is an Algorithm?

9 There are many definitions of what an algorithm is (cf. inter alia ([Hill 2016](#)), ([Mittelstadt et al. 2016](#))). Fundamentally, an algorithm is a set of instructions that, in a finite number of steps, possibly repeating certain steps or sequences of steps, is intended to lead to a solution to a problem. An algorithm can be computerized – or not.

10 One example for a non-computerized algorithm is the Basic Life Support algorithm as you might learn it, for example, in a first aid course. The Basic Life Support (BLS) algorithm is a series of steps designed to provide immediate care to individuals experiencing cardiac arrest or other life-threatening emergencies. It consists of the following steps:

11 1) Assess the Scene: Ensure that the scene is safe for both you and the victim. Look for any potential hazards that could harm you or the person in need of assistance. 2) Check Responsiveness: Gently shake the victim's shoulders and ask loudly, "Are you okay?" Look for any response, such as movement or sound. If there is no response, the person may be unresponsive and in need of immediate help. 3) Activate Emergency Medical Services (EMS): If the victim is unresponsive, shout for help and instruct someone to call emergency services. If you are alone, make the call yourself. 4) Check Breathing: Place your ear close to the victim's mouth and nose while looking at their chest. Look, listen, and feel for breathing for about 5-10 seconds. If the victim is not breathing or only gasping, it indicates a need for Cardiopulmonary Resuscitation (CPR). 5) Start CPR: If the victim is not breathing or only gasping, begin CPR. Position the victim on their back on a firm surface. Place the heel of one hand on the center of the victim's chest, then place the other hand on top of the first. Interlock your fingers and posi-

tion yourself directly over the victim's chest. Perform chest compressions by pushing down hard and fast, aiming for a rate of 100-120 compressions per minute. After 30 compressions, give two rescue breaths by tilting the victim's head back slightly, lifting the chin, pinching the nose shut, and giving two full breaths into the victim's mouth. 6) Continue CPR: Perform cycles of 30 chest compressions followed by two rescue breaths. Continue CPR until help arrives or the victim shows signs of life, such as breathing or movement. 7) Use an Automated External Defibrillator (AED): If an AED is available, follow the device's instructions to deliver a shock to the victim's heart if advised. Resume CPR immediately after delivering the shock. 8) Continue Care: Continue to monitor the victim's condition and provide care until EMS personnel arrive and take over.

12 The Basic Life Support Algorithm is an algorithm – just one that is executed by humans. When we talk about algorithms today, we usually mean things that look a bit different: The term algorithm is now often associated with computerized variants. However, fundamentally, computerized digital algorithms are still a sequence of specified steps in a finite number to solve a defined problem. Even algorithms based on machine learning are in many ways just that, except that they are based on statistical probabilistic models derived from training data: Classical algorithms adhere to deterministic principles, wherein a prescribed sequence of steps is executed to achieve a definitive solution to a given problem. These algorithms rely on logical rules to navigate through predefined paths towards an exact outcome. In contrast, probabilistic algorithms introduce probability into the problem-solving process by incorporating probabilistic models. Rather than aiming for a singular solution, probabilistic algorithms offer a range of potential outcomes, each associated with a probability of occurrence. Thus, while classical problem-solving algorithms, known from written arithmetic, inevitably lead to one solution when correctly applied, predictive algorithms, which operate statistically, generate probable outputs.

13 Computerized algorithms are used today in many places and ways; every Google search is based on them, but they are also used in human resources, insurance and banking, or in law enforcement.⁸ They are used to determine eligibility for social benefits, to determine the urgency with which a homeless person needs housing, or the risk of neglect or child abuse within families.⁹ So, if there is something like algorithmic discrimination, then it is not a marginal phenomenon. Moreover, it potentially concerns sensitive areas of life and profound effects on people's resource allocation and general life standard. However, what does it mean for an algorithm to discriminate and when is this discrimination problematic. Lets turn to the notion of discrimination first.

2.2 What is Discrimination?

14 In everyday language, the term "discrimination" has a negative connotation. When we talk

8. For an impressive presentation of examples cf. (O'Neil 2016).

9. Cf. (Eubanks 2018).

about discrimination today, we mostly mean morally and legally problematic unequal treatment of people. We use the term (often rightly) to criticize, to point out injustices and wrongs. However, there is also a descriptive usage closer to the word's origin. The word "discrimination," stems from the Latin verb *discriminare*, which initially simply means "to separate, distinguish." In English, this meaning has been preserved more than in German, for example, in formulations such as "she has a discriminating taste", or "we need to discriminate between methods and solutions", or "we have to discriminate reliably between legitimate and illegitimate cases". In this understanding, "to discriminate" simply means to make a distinction.¹⁰

15 This can be based on certain properties of a person without being *per se* problematic: In Germany, you must be at least 40 years old to become president; you must have passed the second state examination in law to become a judge; you must be registered as seeking work to receive unemployment benefits. I am, for instance, still too young to become president, am not allowed to work as a judge, because my Master of Arts does not qualify me for that, but fortunately enough, I am nevertheless employed and therefore do not receive unemployment benefits. The distinctions made can have profound effects on my resource allocation and my living situation. However, they are not morally problematic for that reason alone.

16 I point to this descriptive meaning of discrimination, as "making distinctions," here – not to play down the harm and suffering caused by unjust discrimination in any way. What I am concerned with is to emphasize that it is not always easy to answer what is problematic about the distinctions made and why certain distinctions are morally wrong: Distinctions are something we make constantly in our daily lives. Some result in unequal treatment of people. Some of these distinctions resulting in unequal treatment are morally irrelevant: For example, everything being equal to assign students to a group task according to the first letter of their surname in a class setting. Other distinctions that lead to unequal treatment might be regarded as morally good, but not morally obligatory: A teacher pays more attention to a student who does not understand the subject matter as well as other students.

17 In certain instances, however, differential treatment based on distinction made between people is not only permissible but morally obligatory: For instance, it would be entirely appropriate to offer an aspirin to an otherwise healthy adult aunt experiencing headaches. However, under no circumstances should an aspirin be administered to a one-year-old child, even if they complain of severe headaches, due to the potential risk of triggering Reye's syndrome, which can be potentially fatal.

18 Making distinctions, even distinctions that result in unequal treatment are not in themselves morally questionable. Therefore, we need further criteria to understand what makes a distinction made wrong. In the digital age, we further need criteria that can also be applied to algorithms, since, as we have seen, for on the one hand algorithms are based on distinctions, and for the other hand are used in many contexts of life in which they have a profound impact on the

10. On the debate about the appropriateness of a generic or a normative concept of discrimination, cf. ([Altman 2020](#)).

resources available to persons and on their life prospects. As we have seen, an algorithm is a sequence of specified steps in a finite number to solve a defined problem. As such, it must make distinctions – it must discriminate. Otherwise, it could not fulfill its function: Every algorithm processes data in some way. This processing always involves some form of discrimination. For example, an image processing algorithm trained to identify cars on pictures, data about the geometry cars must be processed in such a way that the algorithm can identify a in an image. The algorithm must discriminate between what is a car and what is not a car, what is the configuration of features that constitutes a care, and so on. The key question is: When is algorithmic discrimination morally wrong?

3. When Is Discrimination Wrong?

¹⁹ In her book "When Is Discrimination Wrong?", Deborah Hellman delves into the question of what exactly renders certain differentiations morally wrong. Her proposition suggests that the crucial point lies in whether the differentiation is demeaning: "Discrimination is wrong when it demeans" (Hellman 2008, 33).¹¹ However, Hellman argues that whether a particular differentiation or unequal treatment is demeaning depends on the social context in which it occurs: "Whether distinguishing among people demeans any of those affected is determined by the social context in which the action occurs" (Hellman 2008, 27). The exact nature of what is demeaning is context-sensitive. For Hellman, it is immaterial whether harm arises from the differentiation or not. This argument is compelling as there exist differentiations that yield no disadvantages for the affected individuals, perhaps even advantages, yet are still demeaning.¹² The example that Hellman refers to illustrate her point, is the following: Nelson Mandela vividly recounts in his autobiography "Long Walk to Freedom" that under the South African apartheid regime, prisoners identified as "black" were required to wear shorts, while those classified as "white" or belonging to the "Asian" group wore long pants (Mandela 1997, 515-516). Despite the practicality of shorts in South Africa's climatic conditions, as Hellman aptly emphasizes, this practice was demeaning. Firstly, the classification itself within the apartheid regime served to degrade one group over another. Secondly, although shorts were not disadvantageous under the circumstances, they nevertheless symbolized infantilization: In the English Upper Class tradition, male offspring wear shorts until approximately 8 years old. This legacy of British colonialism manifested in this disparate treatment: boys wear shorts, not adult men. Hence, the connotation of infantilization. Such a treatment although advantageous is thus according to Hellman still demeaning.

11. This approach is indeed close to Margalit's "Humiliation" (1996), which Hellman also discusses (Hellman 2008, 56-57). For critique against this approach: cf. (Arneson 2013, 91-94), (Lippert-Rasmussen 2014, 134-139).

12. Against this (Lippert-Rasmussen 2014, 165-183).

20 Although Hellman is not concerned with algorithmic discrimination, her approach is, as I will argue, particularly attractive, especially concerning algorithmic discrimination. I will focus on two advantages: The first advantage is that, contrary to other approaches, for Hellman, it is not crucial whether the affected individuals feel demeaned – they may or may not. This point is persuasive for several reasons. Firstly, we acknowledge that individuals can internalize demeaning differentiations and no longer perceive them as wrong, even when they are the ones affected. Secondly, concerning algorithmic discrimination – and here I go beyond Hellman – many affected individuals may not even be aware of their unequal treatment and the underlying parameters due to the widely discussed opacity of complex computerized algorithmic systems.¹³ Therefore, under these circumstances, they might not develop a sense of it being demeaning.

21 The second advantage is that while some argue that morally problematic discriminations require a discriminatory intent, Hellman asserts that this is not decisive. This part of her account is compelling because we do not need to have the intention to discriminate against someone in a demeaning manner to still do so. This is due in part to the fact that many of our stereotypes and prejudices are not readily apparent to us as are the structural preconditions of discrimination. Nonetheless, the distinctions we make can be morally impermissible. Furthermore, and here again, I go beyond Hellman's own account: This point is also advantageous concerning algorithmic discrimination because intentions are irrelevant for algorithms. Algorithms operate as problem-solving structures – potentially computerized. Hence, if we were to assume that discrimination is morally problematic only when someone intends to demean someone, then algorithms would be exempt.¹⁴ Furthermore, there is a phenomenon called "emergent discrimination" in algorithm research: This term refers to the emergence of morally problematic classification and unequal treatment patterns in the machine learning process, that result from the interlinking of data and seemingly "harmless" proxies under changed conditions of application.¹⁵ Here, nobody – not even the developers – intends to demean a particular group of people yet demeaning unequal treatment may occur through the interplay of certain data sets when algorithms are employed – especially machine learning algorithms.¹⁶ Hellman's approach could accommodate this phenomenon as it does not hinge entirely on the intention behind a differentiation.

22 After highlighting the advantages¹⁷ of Hellman's approach, I would like to point out a gap: ac-

13. Cf. (Koch 2020).

14. This does not affect the point that all technology is created intentionally and that intentions and values are therefore inscribed in it: Cf. Benjamin

15. Cf. (Mann and Matzner 2019). Cf. also the definition by (Friedman and Nissenbaum 1996) of emergent bias: Friedman and Nissenbaum focus on specific application scenarios, referring to "emergent bias" when a bias emerges under particular application conditions.

16. For a more extensive discussion on whether intentions make a difference regarding whether an action is rendered morally problematic by them when it would otherwise be considered unproblematic, see also (Scanlon 2000).

17. Nevertheless, there are further assessments by Hellman that I do not share. For instance, she contends that it is necessary for a person to be in a position of power over another in order to demean them at all (Hellman 2008, 38). Conversely, I would argue that in some cases, this asser

ording to Hellman, what constitutes the specific demeaning aspect varies depending on the context. However, what remains context-independent is that demeaning is morally problematic. Thus, we require a context-independent justification for why demeaning is morally problematic. Hellman briefly hints at this within her argument.¹⁸ I would like to sketch an alternative in the remaining paragraphs.

4. Kantian Perspective on Demeaning Discrimination

²³ To further strengthen the argument that demeaning discrimination is morally wrong, we could for instance turn to Immanuel Kant's reflections on the respect owed to others in the *Metaphysics of Morals*. Here Kant derives various duties, from the Categorical Imperative. For him, there are both perfect and imperfect duties, directed towards oneself and others, respectively. While his argumentation raises a number of questions, it suffices for our purposes here, to note that one of these duties is the duty to respect others. Kant formulates this duty in the second part of the *Metaphysics of Morals*, the Doctrine of Virtue, under the heading: "On the Duties of Virtue towards Others from the Respect owed to Them."

²⁴ According to Kant, this duty constitutes a perfect duty towards others. What renders a duty perfect and sets it apart from an imperfect duty is subject to contentious debate within Kantian scholarship. A common interpretation is that such duties allow no exceptions or no leeway. In any case, as Kant elucidates, the duty to respect others is, interestingly, a negative duty (§ 41), a term that leaves less room for interpretation: It is a duty of omission. This means, as Kant writes: "I am not bound to worship others, i.e., to show them positive esteem" [*Ich bin nicht verbunden, andere zu verehren, d.i. ihnen positive Hochachtung zu beweisen*] (Kant 1900, 467).

²⁵ What is more, the duty to respect does not oblige one to feel a sense of respect either. Rather, one is obliged "to acknowledge practically the dignity of humanity in every one" [*die Würde der Menschheit in jedem anderen praktisch anzuerkennen*] (Kant 1900, 462). "Practically" here, as in other contexts in Kant's philosophy, refers to actions. The duty of respect for others is not about the feeling of respect but rather about refraining from any actions that would undermine the respect that we owe them. Demeaning others would violate this duty and is therefore to be refrained from. Every individual, as Kant states, has a "lawful claim" [*gesetzmäßigen Anspruch*] to the fulfillment of this negative duty (Kant 1900, 464), meaning it is not merely meritorious to fulfill this duty but rather we owe it to others. Hence, Kant also refers to it as "due respect" [*schuldige Achtung*].

tion is redundant, as the social context of the categories and their history already "speaks" to this power dynamic. In other cases, however, it is even incorrect: I can demean others through my actions and their significance without being in a power relationship with the individual in question. Simply because I fail to express the respect due to them.

18. Hellman argues that demeaning behavior is significant because it violates the moral equality of all individuals. For her, demeaning behavior

26 Building upon this concept of "due respect," algorithmic discrimination is therefore wrong when it makes demeaning distinctions, precisely because respect is owed to every individual, which, as Kant puts it, must be practically acknowledged. There is a moral obligation to do so.

27 I am aware that these remarks still leave many questions unanswered, calling for a more thorough investigation, both with regard to the respective interpretations in Kant as well as with regard to how can we apply this idea to algorithmic discrimination. For instance, who is under an obligation to uphold respect for others – and can we transfer this duty to technology. Or, are we rather dealing here with a derivative duty that we have as bystanders of immoral actions, i.e. even though we ourselves might not violate the duty of respect owed to others, are we under an obligation to respond when we witness the violation of this perfect duty. What would the appropriate response be? How could one argue for such a duty with respect to violations of perfect duties in general in Kantian terms? How could we apply such a reasoning to algorithmic discrimination? Who are, for instance, the addressees of such a response? Here is not the place to answer these questions. However, I aim to have illustrated the potential fruitfulness of pursuing further inquiry in this direction.

5. Conclusion

28 In the public and political debate on Artificial Intelligence and Machine Learning, we often encounter calls for non-discrimination or "discrimination-free" algorithms. Less frequently it is elaborated what this would entail. This is particularly problematic because algorithms inherently "distinguish" – or indeed "discriminate" – by necessity: Making and applying distinctions is essential to algorithmic processes and applications. Hence, what is needed is an approach capable of meaningfully differentiating morally unproblematic from morally problematic distinctions – one that is also applicable to algorithmic discrimination.

29 I have argued, here, that Deborah Hellman's approach formulated in *When is discrimination wrong?* (Hellman 2008) is uniquely suited for application to algorithmic discrimination, unlike other approaches. Her idea is that discrimination is wrong, when it is demeaning. The empirical question of when a distinction and its accompanying treatment, however, are demeaning is context-sensitive. On the contrary, the justification for why precisely distinctions that are demeaning are morally problematic is not context-sensitive. Here, I have extended beyond

carries a comparative element. In contrast, I would want to argue that no such comparative element is necessary to establish its moral dubiousness; rather, it suffices that someone is not accorded the respect they are owed. Here, I concur with Frankfurt: "being treated with appropriate respect and consideration and concern have nothing essentially to do with the respect and consideration and concern that other people that other people are shown" (Frankfurt 1997, 7). Often, we can only discern the demeaning nature of an action in comparison; it is only then that it becomes apparent to us, but it exists independently of this comparison. The shorts worn by black prisoners in South Africa during the apartheid regime would have been demeaning even if white prisoners had not received long pants.

Deborah Hellman's approach by referencing Immanuel Kant's considerations in the *Metaphysics of Morals* to provide a context-independent answer: algorithmic discrimination is wrong when it makes demeaning distinctions because of the respect is owed to every individual.

Bibliography

- Altman, Andrew. 2020. "Discrimination." In *The Stanford Encyclopedia of Philosophy*. Winter 2020 Edition. <https://plato.stanford.edu/archives/win2020/entries/discrimination/>.
- Arneson, Richard. 2013. "Discrimination, Disparate Impact, and Theories of Justice." In *Philosophical Foundations of Discrimination Law*, edited by Deborah Hellman and Sophia Moreau, 87–111. Oxford: Oxford University Press.
- Citron, Danielle K., and Frank Pasquale. 2014. "The Scored Society. Due Process for Automated Predictions." *Washington Law Review* 89:1–33.
- Crawford, Kate, and Jason Schultz. 2014. "Big Data and Due Process." *Boston Legal College Law Review* 55:93–128.
- Eubanks, Virginia. 2018. *Automating Inequality*. New York: St. Martin's Press.
- Frankfurt, Harry. 1997. "Equality and Respect." *Social Research* 64 (1): 3–15.
- Friedman, Batya, and Helen Nissenbaum. 1996. "Bias in Computer Systems." *ACM Transactions on Information Systems* 14 (3): 330–47.
- Hagendorff, Thilo. 2019. "Maschinelles Lernen Und Diskriminierung. Probleme Und Lösungsansätze." *Österreichische Zeitschrift Für Soziologie* 44:53–66.
- Heesen, Jessica, Karoline Reinhardt, and Laura Schelenz. 2021. "Diskriminierung Durch Algorithmen Vermeiden: Analysen Und Instrumente Für Eine Demokratische Digitale Gesellschaft." In *Diskriminierung Und Antidiskriminierung. Beiträge Aus Wissenschaft Und Praxis*, edited by Gero Bauer, Maria Kechaja, Sebastian Engelmann, and Lean Haug, 129–47. Bielefeld: transcript.
- Hellman, Deborah. 2008. *When Is Discrimination Wrong?* Cambridge: Harvard University Press.
- . 2020. "Measuring Algorithmic Fairness." *Virginia Law Review* 106 (4): 811–66.
- Hill, Robin K. 2016. "What an Algorithm Is." *Philosophy & Technology* 29 (1): 35–59.
- Kant, Immanuel. 1900. "Die Metaphysik Der Sitten." In *Gesammelte Schriften*, edited by Königlich

- Preußische Akademie der Wissenschaften, 203–493. Akademie-Ausgabe, VI. Berlin: Akademie-Verlag.
- Koch, Heiner. 2020. “Intransparente Diskriminierung Durch Maschinelles Lernen.” *Zeitschrift Für Praktische Philosophie* 7 (1): 265–300. <https://doi.org/https://doi.org/10.22613/zfpp/7.1.9>.
- Lippert-Rasmussen, Kasper. 2014. *Born Free and Equal? A Philosophical Inquiry into the Nature of Discrimination*. Oxford: Oxford University Press.
- , ed. 2017. *The Routledge Handbook of the Ethics of Discrimination*. London: Routledge.
- Mandela, Nelson. 1997. *Der Lange Weg Zur Freiheit*. Frankfurt/M.: Suhrkamp.
- Mann, Monique, and Tobias Matzner. 2019. “Challenging Algorithmic Profiling: The Limits of Data Protection and Anti-Discrimination in Responding to Emergent Discrimination.” *Big Data & Society* 6 (2): 1–19.
- Mittelstadt, Brent Daniel, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter, and Luciano Floridi. 2016. “The Ethics of Algorithms: Mapping the Debate.” *Big Data & Society* 3 (2).
- O’Neil, Cathy. 2016. *Weapons of Math Destruction*. London: Penguin.
- Reinhardt, Karoline. 2020. “Between Identity and Ambiguity. Some Conceptual Considerations on Diversity.” *Symposion* 7 (2): 261–83.
- Scanlon, Thomas Michael. 2000. “Intention and Permissibility.” *Proceedings of the Aristotelian Society* 74:301–17.