

# A Relational Justification of AI Democratization

Bauke Wielinga<sup>1</sup>, Stefan Buijsman<sup>2</sup>

<sup>1</sup>Delft University of Technology

<sup>2</sup>Delft University of Technology

b.wielinga@tudelft.nl, s.n.r.buijsman@tudelft.nl

## Abstract

While much has been written about what democratized AI should look like, there has been surprisingly little attention for the normative grounds of AI democratization. Existing calls for AI democratization that do make explicit arguments broadly fall into two categories: outcome-based and legitimacy-based, corresponding to outcome-based and process-based views of procedural justice respectively. This paper argues that we should favor relational justifications of AI democratization to outcome-based ones, because the former additionally provide outcome-independent reasons for AI democratization. Moreover, existing legitimacy-based arguments often leave the why of AI democratization implicit and instead focus on the how. We present two relational arguments for AI democratization: one based on empirical findings regarding the perceived importance of relational features of decision-making procedures, and one based on Iris Marion Young’s conception of justice, according to which the main forms of injustice are domination and oppression. We show how these arguments lead to requirements for procedural fairness and thus also offer guidance on the how of AI democratization. Finally, we consider several objections to AI democratization, including worries concerning epistemic exploitation.

## Introduction

With the increasing impact of AI systems such as Large Language Models, social media and gig-economy algorithms, and algorithmic decision-making systems such as government fraud detection systems have had in recent years, it is no surprise that calls have been made to democratize AI (e.g. Wong 2020; Holm 2023). Democratizing AI can refer to a range of phenomena such as widening its use, sharing its profits, and crowdsourcing its design (Seger et al. 2023). We will be concerned primarily with the democratizing of AI governance, design and implementation. Democratizing AI governance means that “a wider community of stakeholders and impacted populations” (Seger et al. 2023) should have influence in decisions about “balancing AI-related risks and benefits to determine if, how, and by whom AI should be used, developed and shared” (Seger et al.

2023). Following Himmelreich (2023), we will take it that calling for AI democratization means that “AI should be subject to novel or different forms of democratic governance” beyond existing democratic state infrastructures, or common practices of companies developing AI. Throughout, we use the term AI in a broad sense, encompassing both complex LLM’s and simple rule-based systems. We conceive of these systems, moreover, as parts of larger socio-technical systems, such as executive agencies, social media networks, etc.

Much existing literature on AI democratization has focused on what democratized AI should look like, rather than on why we should be democratizing it in the first place. For example, Sætra, Borgehund and Coeckelbergh (2022) and Seger et al. (2023) urge careful consideration of what we mean by democracy. Wong (2020) and Buhman and Fieseler (2023) propose deliberative democracy-inspired approaches to democratizing AI, and Noorman and Swierstra (2023) propose Warren’s (2017) problem based approach as a framework for how to democratize AI in particular contexts. While we do not mean to criticize such approaches to AI democratization, we do think that this emphasis on how to democratize AI has meant a lack of emphasis on why AI should be democratized. Arguments in favor are typically brief, and fall into two categories.

First, there are calls which focus on the benefits that AI democratization might bring: using the input of affected persons to improve AI systems might result in AI that is more effective at meeting the needs of users and in the reduction of associated risks and harms (see, e.g. Murphy and Taylor 2023; Rahwan 2018). We argue in section 1 that such calls rely on outcome-based views of procedural justice and as such miss important reasons to consider AI democratization. Specifically, they miss principled reasons to listen to those whose lives will be affected by AI systems, because it is the only way for AI developers to relate to them as equals, and to avoid domination and oppression.

Second, other calls have given more principled reasons to consider AI democratization (e.g. Wong 2020; Zimmerman;

Rosa and Kim 2020), but, as mentioned, these calls have rather focused their attention on how to democratize AI, and lack detailed theoretical accounts of why AI should be democratized. This is unfortunate because close attention to why we should democratize AI is important. Not only do some argue against democratizing AI, and is it thus important to provide detailed argument, it can also help us ground the discussion on how to democratize AI, as we will show.

In this paper, we will argue that calls for AI democratization should be based primarily on relational considerations. We will make two such relational arguments by drawing on the work of Meyerson and Mackenzie (2018) and Young (1990). Section 1 argues that existing scholarship calling for AI democratization tends to rely on either outcome-based or legitimacy-based considerations. Section 2 contrasts these kinds of justifications with relational ones, and makes a relational argument in favor of AI democratization, based on empirical evidence. Section 3 details a second relational argument for AI democratization based on Young's *Justice and the Politics of Difference*. Section 4 discusses the implications of this account for what democratization should look like in practice. Finally, section 5 discusses additional objections and puzzles.

## Section 1: Outcome and Legitimacy-Based Justifications of AI Democratization

Outcome-based justifications of procedural justice assess a procedure by considering how effectively the procedure achieves its intended purpose.<sup>1</sup> In the case of a trial, for example, an outcome-based account would suggest that procedural justice is necessary insofar as it helps reach an accurate and fair decision: the guilty and only the guilty are convicted, and allotted a proportionate punishment (Meyerson and Mackenzie 2018). In this case there is a fairly clear standard for what criteria the outcome of a trial should meet, but no fool-proof way to reach such outcomes (Rawls 1999). In turn, procedures are justified by showing that they manage to reach the correct outcomes often enough. Thus, outcome-based justifications of AI democratization have to assume some set of goals for the governance and development of AI systems. Such goals might include, for example, AI systems that improve lives, do not discriminate, are widely used, maximally profitable and have minimal environmental impact. We should then democratize AI, on an outcome-based account, insofar as democratization promotes the intended outcome of the process of AI governance. This leaves democratization up for debate, insofar as some other model of governance better realizes these goals.

---

<sup>1</sup> Following Rawls (1999) we can additionally distinguish between three kinds of procedural justice: perfect, imperfect and pure. Imperfect procedural justice is what is at stake here, as AI systems do not guarantee correct/just outcomes (as would be the case in perfect procedural justice,

This mode of reasoning is found throughout current AI ethics literature, both in and outside of the AI democratization debate. The idea that seems to be implicit here is that what ultimately matters for the (democratic) legitimacy of an AI system is how often it leads to the right outcomes. Democratization is then instrumental to ensuring that decision-making procedures more often yield those right outcomes.

Outside the democratization debate, Holm (2023) is a good example of this approach, as he argues that “in the context of referral decisions the legitimacy of an algorithmic decision procedure can be fully accounted for in terms of the instrumental values of accuracy and fairness” (Holm 2023, p.1). Indeed, he defends a view called Procedural Instrumentalism on which only instrumental properties, with accuracy being the most important, matter for legitimacy. Corbett-Davies et al. (2023) likewise argue that maximizing the utility/accuracy of an algorithm is, ultimately, what matters. They introduce a consequentialist framework for designing equitable algorithms. On their view, disparities between groups are acceptable as long as the consequences of the decision-making procedure are maximally beneficial to society.

Within the AI democratization debate, similar ideas can be found in Murphy and Taylor (2023) paper (p.1), which “highlight[s] the importance of participation in AI development to ensure its legitimacy, considering the capacity of this technology to shape reality” and gives a number of examples of how democratic participation can improve outcomes, such as in the case of worker participation. “The general result of this involvement is that work conditions are more commodious, production improves, and employees have better attitudes toward their jobs” (Murphy and Taylor 2023, p.4). Here, it is the outcome of the participation that ultimately grounds the call for democratization.

Similarly, Rahwan (2018) argues for a society in the loop approach to AI development, which “is about embedding the values of society, as a whole, in the algorithmic governance of societal outcomes” (Rahwan 2018, p.7). The justification for this democratization is that it will lead to an “aggregation of societal preferences and fair allocation of resources [...] that rational actors would be willing to vote for” (Rahwan 2018, p.11). While this is not a purely outcome-based justification, Rahwan sees both algorithmic aggregation of societal preferences, and decision-making systems programmed to act as though behind a veil of ignorance, as possible ways to implement this preference aggregation. Such a procedure would at least fall short on a relational account of AI democratization, since it lacks consideration for the procedural value of participation and relationality.

where cutting a cake and giving the last piece to the person cutting it is the typical example) nor do AI systems lack a standard of correct outcomes (as would be the case for pure procedural justice, for which betting procedures are the typical example).

Attention to why AI should be democratized also matters for a second group of authors in the literature. These often get closer to a relational or at least process-based justification of AI democratization, but generally focus on the how-question rather than the why-question of AI democratization. Noorman and Swierstra (2023) for example stress that AI is part of political systems and that we therefore need to consider how democratic theory interacts with AI. This leads to a discussion on “how political practices can be improved to strengthen empowered inclusion, collective will formation, and collective decision capacity” (Noorman and Swierstra 2023, p.577) with AI. These considerations are not just about the AI outcomes, but rather focus on how best to realize a democratic system with the advent of AI. Erman and Furendal (2022) similarly proceed from democratic theory in their argument that AI governance needs democratic (normative) legitimacy. On their view, legitimacy requires a minimum of democracy, and they then proceed to fill how this democratic governance should be shaped. As little further detail is offered on the substantive theory of legitimacy (other than compatibility with a range of such theories) there is again a focus on how to democratize rather than why.

This focus on how is also present in other contributions. Buhman and Fieseler (2023) discuss deliberative governance of AI, motivated by a view on which developers have responsibilities for governance that enables systems to both do no harm and to do good. They then argue that distributed deliberation is a promising way to meet this responsibility. Zimmerman, Rosa and Kim (2020) likewise argue that to fix AI we should turn to (technologically facilitated) collective deliberation. In these works, as in e.g. Wong’s (2020) call for democratization of algorithmic fairness, there is typically a reliance on the political nature of AI.

We agree with this point, yet also think that more substantial arguments for why AI should be democratized are available and helpful. Most importantly, such an insight into the reasons for democratization can help evaluate the different proposed remedies. Deliberation, for example, comes in many forms from technologically mediated mass deliberation to McQuillan’s (2018) call for people’s AI councils. However, such appeals to legitimacy and political nature are not sufficiently unpacked to make out exactly what kind of justification the authors have in mind.

Of course, there are nuances between the different accounts discussed here and it is important to realize that even many discussions focused on outcomes the aim is not per se to improve the accuracy of a given AI system, but rather to improve societal outcomes as we develop and use AI. We hope to show with this brief overview of the current debate that outcome-based thinking both in and outside justifications of AI democratization is common, and that more explicit attention for how we justify AI democratization is needed, because, as we discuss below, it ultimately determines how we evaluate whether measures such as stakeholder participation are successful.

We argue that relational theory can help answer the why-question of AI democratization. This highlights also that empowered inclusion itself matters, and not just the outcomes that may (or may not) be improved through democratic deliberation. We therefore work out the case for a more relational perspective on AI democratization in the rest of the paper, starting with Meyerson and Mackenzie (2018) and moving on to Young (1990).

## Section 2: Relational Theory and Empirical Studies on Legitimacy

Relational theory is a group of normative theories that center the relationships, interconnectedness and interdependence between people when thinking about ethics and political philosophy (Birhane 2021). They emphasize that humans do not live, form identities and senses of self-worth in isolation from each other, but rather that we are born into and shaped by various relationships that we have to others. Usually, the premise of relational theories is that relations among humans should be characterized by certain features such as equal standing or moral worth, and certain attitudes such as care. Additionally, “relational theory draws specific attention to the moral significance of social oppression, marginalisation, and inequalities of social power and standing, and their effects on individuals’ self-identities, autonomy, and sense of self-respect” (Meyerson and Mackenzie 2018). Examples of relational theories are relational egalitarianism and care ethics.

In relational accounts of procedural justice the standard by which to evaluate both procedures and outcomes is found in some conception of appropriate ways to relate. Such relational accounts are process-based (Meyerson and Mackenzie 2018) in that the procedure matters independently from the outcome: if the right outcome was arrived at through coercion, say, then this would violate any plausible conception of good relationships. On the other hand, a fair procedure alone does not guarantee relational procedural justice, since the outcome should also be in accord with some plausible conception of good ways of relating. Thus relational accounts of procedural justice take into account both the procedures themselves, and their outcomes.

Meyerson and Mackenzie (2018) argue that in the legal setting (i.e. a court case), relational factors are important in determining whether a procedure is perceived as just. They consider factors such as how respectfully people are spoken to, whether they are listened to and whether they feel like judges care about them. These factors are important, they argue, “because satisfactory interpersonal treatment by group authorities symbolically communicates the information that we possess value or status in the eyes of our community, which in turn supports a sense of self-respect or self-worth” (Meyerson and Mackenzie 2018). This means that when a judge treats you as someone worthy of respect, you are more likely to think of yourself as worthy of respect,

and when authorities do not seem to care about your opinion about how you are treated, this undercuts your self-respect. Hence it is not only important that one's procedural rights are respected, but also that this happens in such a way that this respect is also communicated to the affected person.

Meyerson, Mackenzie and MacDermott (2021) support the importance of relational considerations with a range of empirical studies going back to the work of Thibaut and Walker (1975). These studies indeed highlight that people are more satisfied with outcomes they receive when a process is perceived as fair, regardless of whether the outcome is favorable to them. It is how they were treated and thus what the procedure looks like, that helps determine people's perception of the legitimacy of the procedure. In particular, these empirical studies have found that people care about what is called "voice": the opportunity to be heard during the procedure or after outcomes have been finalized (Burke and Leben 2007; Folger 1977; Lind, Kanfer and Earley 1990) as well as "benevolence", i.e. the willingness of decision makers to consider one's needs and the clear communication of this willingness (Tyler and Blader 2003; Lind and Tyler 1988).

These general findings have been studied more specifically in the context of algorithms by Lee et al. (2019), who investigated people's perception of an algorithm based on different kinds of information and control over the decision making. 'Standards clarity', where the basic functioning of the algorithm (rules-based in this particular study) was explained to participants, failed to improve their perception of the fairness of the algorithm. "Many understood how the algorithm worked and how fairness in division was operationalized based on this step-by-step description, and they later used the knowledge to interpret the input and output matrix table. At the same time, participants told us that the standards clarity alone did not make them trust the algorithm or see the results as being automatically fair" (Lee et al. 2019, p.13). Outcome control, where participants got to experiment with changing outcomes, were able to discuss alternatives with the rest of the group and possibly deviate from the algorithm's recommendation, did significantly improve people's perceptions of fairness even though 80% of the groups made no changes to the allocation of goods. The interpretation here is that precisely aspects such as voice (being able to discuss) and benevolence (seeing that there were no better allocations available) contributed to this improvement. However, one may also interpret it as showing that participants realized that the outcomes were optimal and legitimizing the procedure on that basis.

In another empirical study regarding the legitimacy of algorithms, Martin and Waldman (2022) found that when arbitrary or morally problematic grounds are behind an algorithmic decision, it is not viewed as legitimate regardless of whether the outcome is positive or negative. If the reasons for the decision were deemed in order, then algorithms leading to positive outcomes were judged as more legitimate

than those leading to negative outcomes (decisions that negatively affected the subject or that they disagreed with). Now, this study focused only on single outcomes of the procedure, and thus it cannot rule out that we actually only care about the reasons behind a decision because the wrong reasons would lead to more incorrect decisions overall. Thus the study is not wholly incompatible with outcome-based justifications of procedural justice.

These empirical findings, in and outside the AI context, create a plausible basis for thinking that relational aspects of procedures do in fact matter for their fairness. Thus, we should incorporate them into the justification for AI democratization. A first relational argument for AI democratization can then start from this point: we have (empirical and philosophical) reasons to consider certain relational aspects important for the legitimacy of decision-making procedures that include AI. Based on the work of Meyerson Mackenzie and MacDermott (2021) we highlight here the importance of voice and benevolence as relational aspects that need to be incorporated in the development and implementation of AI for these procedures to be just.

From that follows a call for a type of democratization: if affected parties are to have an opportunity to be heard during the procedure or after outcomes have been finalized then we need a systematic way for them to be included in the development and implementation of AI systems. Moreover, this inclusion of affected parties' voice cannot be one where they are merely given a say without being taken seriously. The need for benevolence in the procedure means that there has to be willingness of decision-makers to consider one's needs, as well as clear communication of this willingness, reflected in the way in which affected parties' perspectives are heard and included. While this may affect outcomes for the better (and, as noted, this would be important from a relational perspective), it additionally matters from the relational perspective whether people are listened to and feel heard independently of whether outcomes are improved. While there will be a trade-off to be made in certain circumstances, on a relational account procedural justice can sometimes trump an optimal outcome.

To get a better grasp on how the inclusion of perspectives should happen concretely we find it helpful to expand the scope of the relational argument somewhat first. The next section will therefore look at Young's conception of justice and the requirement that relationships are characterized by non-domination. From that requirement of non-domination follows a similar need that AI systems' development and implementation is governed democratically by a group of people, insofar as that system shapes the actions of that group.

### Section 3: AI Democratization and the Politics of Difference

Social justice for Young concerns “the degree to which a society contains and supports the institutional conditions necessary for the realization” of two values, which, she claims, comprise the good life: “(1) developing and exercising one’s capacities and expressing one’s experience [...] and (2) participating in determining one’s action and the conditions of one’s action” (Young 1990, p.37). Thus, a society is just to the degree that its laws, institutions and other relevant conditions are such that they allow and realize self-expression and self-determination for all. Justice for Young is not limited to distributive justice, because it concerns many things that are not in a straightforward sense distributable goods. Rather, it consists of institutional conditions (i.e. conditions that are non-incident, spread out over time and regard the structure of society as opposed to singular events) that foster a certain kind of positive freedom, and relationships among people that are free of dominating power.

The two values comprising the good life correspond to the two main kinds of injustice that occur when those values are not realized: Domination, “the institutional constraint on self-determination” (Young 1990, p.37) and oppression, “the institutional constraint on self-development” (Young 1990, p.37).

Domination consists in institutional conditions which inhibit or prevent some people from determining their actions or the conditions of their actions. Persons live within structures of domination if other persons or groups can determine without reciprocation the conditions of their action, either directly or by virtue of the structural consequences of their actions. Thorough social and political democracy is the opposite of domination (Young 1990, p.38).

Thus, domination occurs when people in a society cannot determine for themselves what they will do or under what conditions they will do it. This occurs not only when an autocratic government, say, determines what job every person should have, without the person’s opinion making a difference, but also in more subtle ways, for instance in the bureaucratization of everyday life. “Formalizing collective action through explicitly articulated rules and procedures [...] separates it from normative inquiry and commitment. Decisions and actions will be evaluated less according to whether they are right or just than according to their legal validity” (Young 1990, p.77). The dominated “experience themselves as subject to the unreciprocated authority of others” (Young 1990, p.78).

It’s a somewhat open question how severe the degree of influence over people’s action has to be in order to instantiate Young’s definition of oppression. However, if we assume that less than total control suffices, there is good reason to think that some currently existing AI systems fit the

first part of the definition, regarding the partial determination of courses of action. If so, then we should ensure through democratic control that we do not end up with domination of affected persons by guaranteeing that there is reciprocation of control over the developer or employer’s actions.

The first part of this argument then goes as follows: there are many examples of AI systems shaping people’s actions or the conditions of their actions in significant ways. One potent example is found in the algorithms that govern social media feeds and that power online search engines. Use of such systems is for many people almost inevitable because of how ingrained they are in everyday life, and the lives of millions of people are shaped by them, for example through the information that such algorithms present and do not present (which, Noble’s 2018 book argues, can shape stereotypes), by how addictive the systems are designed to be, and by the way they shape social relationships. As a result, these systems not only shape people’s actions and the conditions of their actions, but the fact that such conditions apply also inhibits people in determining their actions. Thus, in the absence of democratic governance, they dominate.

Moreover, such systems meet at least the second of three conditions suggested by Himmelreich (2023), each of which implies that they are in need of democratic legitimation: they have pervasive impacts on people’s lives. Himmelreich argues that AI never has pervasive impact by itself, because it is always part of some institution which already has the relevant powers and impact. However, social media networks and search engines are only possible and can only be so pervasively impactful because of their algorithms. Hence, it seems fair to say that they should be democratized.

One can also think of various algorithms employed by government agencies, for example to check the compliance of citizens with social security rules, or to do predictive policing. There are many examples of such systems, from Northpointe’s reoffending risk assessment tool (Angwin et al. 2023), various “fraud-detecting” systems used by the Dutch government (Klaasen en van Dijk 2023; Volkskrant 2021; Rengers et al. 2023), or Austria’s profiling system for job-seekers (Allhutter et al. 2020). Such systems have enormous impacts on the lives and actions of citizens, but they are often not developed because of policy choices that governments make, but, as illustrated in the case of Dutch municipalities by Choi, van Eck and Hukshorn (2021), within the prerogative of civil servants. Insofar as such systems are developed within the prerogative of civil servants, citizens do not participate in determining how those systems will shape their actions, and thus those systems enact domination on them. Since these systems are not only impactful, but also coercive, they also need to be democratized according to Himmelreich’s (2023) criteria.

More examples of algorithms which prevent people from determining their actions are readily available, such as the use of AI in recruitment, law enforcement, border control and credit scoring. As such, they meet the first part of

Young's definition in how they relate to decision subjects. Of course, such systems would not be dominating insofar as a democratically elected government ordered their implementation. In that case, at least some groups of affected citizens would indirectly have been involved in determining how the new AI system would shape their actions, and so it would not dominate them. However, such a system could still be oppressive towards those same groups, for instance if it reinforces harmful stereotypes through discriminatory policing. Furthermore, it would also still be oppressive towards those affected groups of citizens whose perspectives were not heard and given due weight in the decision-making process leading to the implementation of the system. In terms from the empirical literature discussed by Meyerson and Mackenzie (2018): such systems could still be oppressive insofar as the processes implementing them did not give voice to all affected parties, or was not equally benevolent. Thus the mere presence of democratic governance need not be enough to ensure procedural justice.

In short, since AI systems frequently are part of institutional conditions that constrain actions, we need to ensure that there is reciprocity of authority and control, and thus a democratization of the governance of AI. Another important question is how exactly AI should be democratized. Here Young's view on oppression is a good starting point:

Oppression consists in systemic institutional conditions which prevent some people from learning and using satisfying and expansive skills in socially recognized settings, or institutionalized social processes which inhibit people's ability to play and communicate with others or to express their feelings and perspective on social life in contexts where others can listen (Young 1990, p.38).

Oppression thus occurs when people in a society are not able to develop and express themselves, for example because some people or groups are not allowed to work or go to school. Oppression is not only something that can occur when a tyrant exercises illegitimate power, but also "the vast and deep injustices some groups suffer as a consequence of often unconscious assumptions and reactions of well-meaning people in ordinary interactions, media and cultural stereotypes, and structural features of bureaucratic hierarchies and market mechanisms" (Young 1990, p.184). Domination and oppression overlap partially, but at least not insofar as the development of skills and the expression of opinions is compatible with a lack of self-determination.

Young (1990) identifies five faces or aspects of oppression, each of which, she argues, is a sufficient condition to speak of oppression: exploitation, marginalization, powerlessness, cultural imperialism, and violence. We will discuss three of these faces of oppression, as they pertain directly to the question of what AI democratization should look like.

The first face of oppression is exploitation. "[T]he injustice of exploitation," writes Young, "consists in social processes that bring about a transfer of energies from one group

to another to produce unequal distributions, and in the way in which social institutions enable a few to accumulate while they constrain many more" (Young 1990, p.53). The typical example of exploitation is how, according to Marx, workers are exploited by capitalists because in their interaction wealth is systematically generated through the labor of the workers to the benefit of their employers in the form of profit. Young's formulation is more abstract, so as to also capture processes where money is not or not directly involved, such as the exploitation of women in traditional families, where they expend their energy towards household work which is not rewarded with respect, but which frees men to pursue work that is better rewarded and acknowledged.

This more abstract concept of exploitation additionally can be used to conceptualize *epistemic exploitation*, which occurs when the (epistemic) energies of one group are used to systematically and unreciprocally (epistemically) benefit another group. An example of epistemic exploitation discussed by Cooper (2021), is a situation where Black university staff are structurally expected to do more work to help universities understand better how to handle diversity issues, while this benefits the university and not them. This concept is useful for the design of a participatory democratic process because, as we will see, such processes risk being epistemically exploitative.

For the purpose of democratizing AI, the most important face of oppression in powerlessness. Of the faces of oppression, it comes most close to being the same as domination, and can be seen as a condition applying to those in liberal societies, who face the most severe domination:

[D]omination in modern society is enacted through the widely dispersed powers of many agents mediating the decisions of others. To that extent many people have some power in relation others, even though they lack the power to decide policies or results. The powerless are those who lack authority or power even in this mediated sense, those over whom power is exercised without their exercising it (Young 1990, p.56).

Young sees powerlessness as a kind of class division, "a position in the division of labor and the concomitant social position that allows persons little opportunity to develop and exercise skills" (Young 1990, p.56). Powerlessness is not quite the same as domination, since it is a total lack of power, whereas many who are dominated still have some power over yet other people themselves: "All who do not participate in deciding the rules and choosing the authorities that order their actions and the conditions of their actions experience domination. This includes all who are powerless. Many who experience domination nevertheless exercise considerable initiative, authority, and creativity in their work, and this is what distinguishes them from those subject to the oppression of powerlessness" (Young 1990, pp.219-220).

The final face of oppression, cultural imperialism, happens when the culture of some social groups is dominant in a society, and determines the kinds culturally informed perceptions and associations all groups in that society have. Young writes that “[t]o experience cultural imperialism is to experience how the dominant meanings of a society render the particular perspective of one’s own group invisible at the same time as they stereotype one’s group and mark it out as the Other” (Young 1990, p.59). The perspectives and meanings of the oppressed group are supplanted in a sense by those of the dominant group, and seemingly neutral cultural expression tend to assume the perspective of the dominant group, whereas if oppressed groups are represented, this is either from the perspective of the dominant group, as a stereotype, or it is perceived as very different, standing out from the seemingly neutral norm. This is what Young calls “the universalisation of a dominant group’s experience and culture, and its establishment as the norm” (Young 1990, p.59).

Cultural imperialism is relevant for our discussion, as its presence in the wider culture within which a participatory democratic process takes place may distort that process in favor of privileged groups and against oppressed groups. Hence, Young argues, in societies “where social group difference exists and some groups are privileged while others are oppressed, social justice requires explicitly acknowledging and attending to those group differences in order to undermine oppression” (Young 1990, p.3). Democratic processes that wish to truly include the voices of underprivileged social groups should make arrangements for their systematic representation (Young 1990, p.184). By explicitly representing, for example, racial and ethnic minority groups in decision-making procedures, we can strive to counter-act mechanisms of cultural imperialism that might otherwise silence their voices.

Taking these different elements together, Young’s account of justice has at least two implications about what the AI democratization demanded by the above argument should look like: it should be *representative* of the affected by the systems, especially those groups suffering any form of oppression. Here, the five faces of oppression serve as a good criterion of what social groups especially need to be represented in the procedures, as these are the groups who would otherwise lack reciprocal power over the decision-making procedures around impactful AI systems. Secondly, if democratization is to act as a remedy to domination, it should involve empowerment, in the sense that representation should be effective. In other words, the voice and perspective of those affected groups have to be given due weight in the decision-making process and be considered on their merits. Interestingly, these are very much the same conclusions that we saw from the relational account of Meyerson, Mackenzie and MacDermott (2021), derived partly from empirical studies on people’s perceptions of legitimacy. The requirement of representation mirrors that of giv-

ing people a voice, whereas the effectiveness of that representation mirrors the need for benevolence, where people’s values and needs are taken seriously by decision-makers. We thus have two different approaches, one more empirical and one more purely normative, converging on the same precepts for AI governance. We need to give people a voice in the development and implementation of AI systems, and this voice must be taken seriously.

Importantly, this can diverge from the outcome-based justifications of AI democratization discussed in section 1. Giving people a voice and taking it seriously is something that we can expect to improve outcomes, but even if outcomes aren’t improved it is still important to listen to affected parties. It is not the effect they have on the outcome that matters, on our view, but their actual inclusion. Furthermore, improved outcomes can also be obtained in ways that are still exploitative and thus oppressive, as we discuss in the next section. It is therefore important to look just as much at how we can understand the requirement for democratization more fully, in a way that offers people a voice and treats them benevolently without leading to domination and oppression.

#### **Section 4: Shaping Effective Representation for AI Governance**

This section will consider some criteria that AI democratization should meet if it is to be implemented justly. The section starts by discussing some general points, before treating in detail Pohlhaus’ (2020) worries regarding pernicious inclusions in epistemic systems as requirements for just AI democratization.

If we need further representation of affected parties in the development and implementation of AI, then a crucial question is *how exactly* affected groups should be *represented* in decision making processes. There is no one agreed upon conception of political representation (Dovi 2018), and representation can take many different shapes. Hanna Pitkin (1967) has distinguished various different criteria by which we could judge representatives: e.g. should we ask whether representatives are accepted as representatives by their constituents, whether they *resemble* the represented, or whether they realize their interests (Dovi 2018; Pitkin 1967).

Each of these considerations seems to be important from a relational perspective, but different ways of organizing a participatory process may promote these different concerns to different degrees. Representatives could be chosen, for example, through elections or lotteries, or on a volunteer basis, and this has an effect on how well they represent. Representatives chosen through election might count on more acceptance by the represented for example, but may not actually resemble their constituents as much as representatives chosen through lottery. We do not offer any way to navigate this dilemma, but do mention these puzzles here, because they will have to be faced when any kind of democratization,

of AI or otherwise, is implemented. Hence, we do not see representation as a problem for AI democratization specifically, but as a problem for any kind of non-direct system of governance.

Still, we can make some more headway in the characteristics of this representation by considering some objections to Young's proposal for democratic pluralism: Oppression, as mentioned, includes epistemic exploitation and from this we can get further requirements on how representation should be shaped. In this we follow Pohlhaus (2020), who has argued that inclusion of agents who are "epistemically marginalized on account of relationships of domination and oppression" in epistemic systems risks inflicting further epistemic injustice by exploiting those agents (Pohlhaus 2020, p.1).

This can happen in several ways. First, epistemic systems may "coercively direct the epistemic agency of marginalized knowers, in ways that asymmetrically serve the epistemic interests of dominantly situated knowers" (Pohlhaus 2020). Thus, by being included in an epistemic system (say, a committee assessing how, and, ideally, whether, a particular generative AI system ought to be implemented), those knowers may end up serving their own interests, as representatives of those affected by the AI system, much less than they do the owners of the AI system, whose generative AI tool can now be implemented more smoothly and profitably.

Second, "semi-inclusions" may occur, situations where a knower is included "in the norms of judgement that govern testimonial exchange", but only to the degree that they can then be excluded according to that system's own norms, for example of what counts as a trustworthy or competent knower.

Thirdly, inclusion in epistemic systems can be unjust, when, through previous exclusions, that system has epistemic gaps which only members of previously excluded groups can fill. Then it can occur that members of those groups are continually expected to help address those gaps. Pohlhaus likens expecting such help to expecting those who use wheelchairs to devote themselves to becoming architects so that they can make buildings more wheelchair-accessible. This is "an infringement on epistemic autonomy, [and] also can be deployed in ways that disregard the epistemic labor involved" (Pohlhaus 2020, p.11). This is akin to Himmelreich's (2023) third objection to AI democratization: it is resource intensive, whereas the most marginalized people in a society lack precisely those resources: "They lack the time, the money, the relevant social capital, or simply the freedom to associate and organize effectively.[...] This taints the democratic process" (Himmelreich 2023, p.1340).

Lastly, Pohlhaus discusses situations where "the recognized systems available are inapt for attending to the experienced world one is expected to represent" (Pohlhaus 2020, p.12). These are situations where one is included in an epistemic system to provide some particular perspective, but the

norms and customs of that system are not capable of recognizing the experiences one is supposed to bring to light. Pohlhaus mentions as an example the case of microaggressions discussed by Fatima (2017), in which, when people of color are asked to speak about microaggressions, they are then asked to justify for every single instance of a microaggression, that it was not a misunderstanding. This is nigh impossible to do, and so despite being included, and being able to give testimony, there is little point in doing so for the person in question, as the epistemic system in place is not capable of recognizing their experience.

The difficulty for AI democratization is then to avoid these pitfalls. We think that this is feasible, but challenging. Participants can be taken seriously. Epistemic systems, though resilient, can be changed, and epistemic labor can be rewarded, and not be taken for granted. For example, in the case of executive agencies dealing with unemployment benefits, there will be many affected persons who are unable to work, and so also to provide epistemic labor. In such cases, participatory representation will be of limited use, and alternative or supplementary ways will have to be found to represent the needs and interests of those affected persons. One way to ensure that epistemic labor is not taken for granted, is to financially compensate participants, which could also help to make participation accessible to a wider group of people.

If participants are taken seriously, then it is also less likely that participation in the democratic governance of AI will not serve the interests of affected groups, and if participation is rewarded, then it will be no more exploitative than other kinds of paid labor. In the case of micro-aggressions, for example, actors in the epistemic system could work to learn more about micro-aggressions, so that they will not apply unrealistic evidential standards. If no effective participation is possible, then asking for participation is unfair, and so some alternative approach should be taken, to ensure that the needs and interests of affected groups are recognized, and in such a way that this does not come across to the affected as an exclusion.

Taken more broadly, decision-makers could, for example, do extra work themselves to increase their capacity for empathy with various affected social groups. Getting acquainted with these groups and their lived experience can be valuable for a better understanding of their needs and values, and perhaps through informal conversations it is possible to gauge their opinions regarding the (kind of) AI system in question. To the degree that it is possible to build empathy in this way, and to the degree that such representation can be effective, this could be a good way to democratize AI without requiring significant epistemic labor from oppressed groups. In practice, however, we think that ensuring (perceived) effectiveness in the sense defined above requires some degree actual democratic participation, and that substitutes are (1) necessary when adequate rewards for epistemic labor are not possible, and (2) desirable in combination with democratic participation.

Further ways to limit the risk of epistemic exploitation may be to focus AI democratization on the governance aspect, rather than the design aspect, as democratic AI governance is more a transfer of control to affected persons, whereas democratic AI design requires more epistemic labor. Additionally, the labor involved in doing AI governance is more likely to be towards one's own end and needs, rather than disproportionately to that of others. Thus, by focusing on democratic AI governance, and on democratic AI design only insofar as it facilitates democratic AI governance, the risk of epistemic exploitation can be reduced.

With regard to Pohlhaus' point concerning respect for autonomy, the issue is complicated: on the one hand, participatory democracy is a very direct way to enhance autonomy. On the other hand, Pohlhaus distinguishes two ways in which inclusion in an epistemic system can harm autonomy: Within the epistemic system, participants' autonomy can be harmed because of that system's entrenched structures, practices, norms, etc. Just participatory democracy would have to find ways to limit this, such as ensuring that the views expressed by participants are recorded, so that their effectiveness can be checked, and by giving participants explicit freedom to focus their attention on issues that matter to them.

Outside of the epistemic system, the worry is that participation may not be entirely voluntary, because there is moral risk involved in not participating. E.g.: if I do not participate and lend my perspective, my perspective may not be adequately considered, and the AI system may end up doing harm to others. Additionally, declining invitations from government institutions that decide over important aspects of your life, may not always be (or be perceived as) a free choice. This is a worry that can perhaps not be fully removed, but it could be lessened by some of the measures discussed above, such as properly rewarding participatory epistemic labor and enhancing empathy between decision-makers and affected parties.

In the end, there will be cases where democratization of AI governance is the only way to meaningfully increase the autonomy of persons significantly affected by new AI systems, and recognition of their interests in the decision-making procedure regarding the implementation of those systems is the only way to avoid unjust domination by those systems. Rather than abandon democratic AI governance because of the risk to the autonomy of participants, the kinds of scenarios described by Pohlhaus should be kept in mind, and avoided to the best degree possible, so that the autonomy of affected persons, including participants, can best be safeguarded.

## Section 5: Discussion

A question might be asked concerning the scope of these arguments for AI democratization: when can we say that AI is (sufficiently) democratically governed, and when is it not?

Looking at the European Union, for example, all of the affected parties have the right to vote in a democratic country that then enacts legislation on AI, such as the AI Act. Echoing Himmelreich's (2023) redundancy objection to AI democratization, is such legislation enough to say that one cannot, as an EU-citizen, be dominated by AI systems?

We consider the answer to be no: laws such as the AI act are insufficient to ward off arbitrary meaningful interference in people's daily lives, as any such law is likely to leave considerable space for choices that are significant for the lives of the people affected by AI systems, without offering appropriate reciprocal power. If a democratically legitimate law forbids or explicitly allows certain matters, then, at least in principle, no further democratization is necessary, but insofar as no such law specifically allows something which has dominating influence over people's, then further democratic decision-making is necessary. For example, if a law were made that allowed only generative AI systems to be built that include watermarks that make their productions recognizable, this would not mean that therefore it has been democratically decided that, say, translators can be fired and replaced by generative AI systems, so long as their outputs include watermarks. Thus, an AI act or similar legislation is not enough, because there is still potential for domination within the systems and design options of systems on which it is neutral.

Furthermore, if the political process that produced the law in question did not sufficiently take into account the needs and interests of affected groups, then those groups can still be dominated by systems and design options that the law explicitly allows or forbids. In both these ways, AI democratization is not redundant as Himmelreich's (2023) second objection to AI democratization has it, since there are limitations to what laws cover, and how well they represent the needs and interests of all affected groups and persons.

Finally, it might be argued, if the reason we should adopt relational rather than outcome-based considerations is that outcome-based considerations put democratization up for debate, aren't relational justifications just as bad, since they also consider the outcomes of a procedure? While it is true that relational justifications take into account the outcomes of a procedure, their benefit compared to outcome-based views is that they additionally take into account relevant aspect of the procedure. Indeed, for procedures where there is enough at stake that democratic decision-making becomes necessary, it would be odd not to look at outcomes at all.

Lastly, we disagree with Himmelreich's (2023) argument that AI democratization is theoretically ill-founded. Himmelreich argues that calls for AI democratization are motivated by cases of algorithmic injustice, such as when AI leads to discriminatory hiring practices, or entrenches stereotypes. There is a mismatch here, Himmelreich argues, because democratization won't necessarily solve those issues. This is false however, at least insofar as a democratic process that is both representative and effective in the senses defined above, effectively representing oppressed groups, is

likely to contribute to solving such issues. This seems plausible. On the other hand, Himmelreich's objection is an outcome-based consideration, and perhaps it is true that there are better ways to avoid these problems. More fundamentally though, calls for democratizing AI should be based primarily on considerations that also (relationally) regard the process, not just the outcome.

This helps us, as mentioned in the beginning of the paper, to evaluate proposals on how to democratize AI. A good proposal to democratize AI should ensure effective representation, that is, representation that gives affected parties a voice that is treated benevolently and in a way that prevents domination. Collective decision-making on its own is thus unlikely to be sufficient, especially if the decision-making happens through majority vote. Many of the issues around (epistemic) oppression and the need for benevolence towards the voice of minorities are not resolved purely through majority voting. As such, the preference aggregation discussed by Rahwan (2018) has a risk of not listening sufficiently to the broad set of values of affected parties, at least in the examples given that focus on distributions of goods or outcomes and Rawlsian maximin solutions. There is a risk there that interests are taken into account, but by the majority group that decides which procedure is fair, as opposed to through giving people a voice in the procedure itself. Instead, attention to features such as empowered inclusion, discussed by Noorman and Swierstra (2023), is needed. Empowered inclusion here does need to mean more than ensuring that people can benefit from or participate in AI-powered socio-technical systems. Empowered inclusion also has to refer to an empowerment of affected parties such that their values and interests are taken into account effectively.

We thus have a set of criteria on which to evaluate different proposals for how to democratize AI. By paying attention to the underlying reason for this move, in our case coming from relational theory, we not only get a more principled motivation for including affected parties but also much-needed standards for that inclusion. This still leaves a lot of work to determine how we can in fact best organize effective representation. Both empirically and conceptually we have a way to go before that question can be answered. We believe, however, that the current arguments can help to steer those efforts.

## Conclusion

Calls to democratize AI are widespread. However, they frequently rely on either outcome-based justifications of democratization or only focus on how AI should be democratized rather than why. We have argued that relational theory, based on both empirical and normative considerations, provides a good account of why AI democratization is important. Improved outcomes are a part of that, but do not

give the entire story. Rather, in order to legitimize AI systems and prevent domination, effective representation is needed. Effective representation has to give voice to affected parties and demonstrate benevolence towards their values. Moreover, we must do this in a way that avoids oppression, thus suggesting further criteria for how effective representation is to be shaped. As a result, getting a firmer grasp on the normative foundation of AI democratization also helps to evaluate different concrete proposals for how to make it happen.

## References

- Allhutter, D.; Florian, C.; Fabian F.; Gabriel G.; and Astrid M. 2020. Algorithmic Profiling of Job Seekers in Austria: How Austerity Politics Are Made Effective. *Frontiers in Big Data* 3: 5 <https://doi.org/10.3389/fdata.2020.00005>.
- Angwin, J.; Larson, J.; Mattu, S. and Kirchner, L. 2023. Machine Bias. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>. Accessed: 26-09-2023.
- Birhane, A. 2021. Algorithmic Injustice: A Relational Ethics Approach. *Patterns* 2(2): 100205. <https://doi.org/10.1016/j.patter.2021.100205>.
- Burke, K. and Leben, S. 2007. Procedural Fairness: A Key Ingredient In Public Satisfaction. *Court Review: The Journal of the American Judges Association* 44(1) <https://digitalcommons.unl.edu/ajacourtreview/226>.
- Buhmann, A. and Fieseler, C. 2023. Deep Learning Meets Deep Democracy: Deliberative Governance and Responsible Innovation in Artificial Intelligence. *Business Ethics Quarterly*, 33(1), 146–179. <https://doi.org/10.1017/beq.2021.42>.
- Choi, W.; van Eck, M. and Hukshorn, H. 2021. Hoe Gemeenten Besluiten Over Algoritmen en Mensenrechten. <https://hooghiemstra-en-partners.nl/onderzoek-hoe-gemeenten-besluiten-over-algoritmen-en-mensenrechten/>. Accessed: 13-05-2024.
- Cooper, B. 2015. Black America's Hidden Tax: Why This Feminist of Color Is Going on Strike. *Salon*. [https://www.salon.com/2015/02/25/black\\_americas\\_hidden\\_tax\\_why\\_this\\_feminist\\_of\\_color\\_is\\_going\\_on\\_strike/](https://www.salon.com/2015/02/25/black_americas_hidden_tax_why_this_feminist_of_color_is_going_on_strike/). Accessed: 13-05-2024.
- Corbett-Davies, S.; Gaebler, J. D.; Nilforoshan, H.; Shroff, R.; and Goel, S. 2023. The measure and mismeasure of fairness. *The Journal of Machine Learning Research* 24(1), 14730-14846. <https://doi.org/10.48550/arXiv.1808.00023>.
- Dovi, S. 2018. Political Representation. In: *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta, Fall 2018. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2018/entries/political-representation/>.
- Erman, E. and Furendal, M. 2022. Artificial Intelligence and the Political Legitimacy of Global Governance. *Political Studies* 72(2). <https://doi.org/10.1177/00323217221126665>.
- Fatima, S. 2017. On the Edge of Knowing: Microaggression and Epistemic Uncertainty as a Woman of Color. In: *Surviv-*

- ing Sexism in Academia: Feminist Strategies for Leadership, edited by K. Cole and H. Hassel, 147–57. New York: Routledge.
- Folger, R. 1977. Distributive and procedural justice: Combined impact of voice and improvement on experienced inequity. *Journal of Personality and Social Psychology* 35(2), 108–119. <https://doi.org/10.1037/0022-3514.35.2.108>.
- Himmelreich, J. 2023. Against ‘Democratizing AI.’ *AI & Society* 38 (4): 1333–46. <https://doi.org/10.1007/s00146-021-01357-z>.
- Holm, S. 2023. Algorithmic Legitimacy in Clinical Decision-Making. *Ethics and Information Technology* 25 (3): 35. <https://doi.org/10.1007/s10676-023-09709-7>.
- Klaassen, S. and van Dijk, R. 2023. Computer Zegt Vrouw: Hoe een Rotterdams Algoritme Jonge, Alleenstaande Moeders Discrimineerde. <https://www.versbeton.nl/2023/03/computer-zegt-vrouw-hoe-een-rotterdams-algoritme-jonge-alleenstaande-moeders-discrimineerde/>. Accessed: 13-05-2024.
- Lee, M. K.; Jain, A.; Cha, H. J.; Ojha, S.; and Kusbit, D. 2019. Procedural Justice in Algorithmic Fairness: Leveraging Transparency and Outcome Control for Fair Algorithmic Mediation. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), 182:1-182:26. <https://doi.org/10.1145/3359284>.
- Lind, E. A. and Tyler, T. R. 1988. The Social Psychology of Procedural Justice. In *Contemporary Sociology* (Vol. 18). <https://doi.org/10.2307/2073346>.
- Lind, E. A.; Kanfer, R. and Earley, P. 1990. Voice, Control, and Procedural Justice: Instrumental and Noninstrumental Concerns in Fairness Judgments. *Journal of Personality and Social Psychology*, 59, 952–959. <https://doi.org/10.1037/0022-3514.59.5.952>
- McQuillan, D. 2018. People’s Councils for Ethical Machine Learning. *Social Media + Society* 4(2). <https://doi.org/10.1177/2056305118768303>.
- Meyerson, D. and Mackenzie, C. 2018. Procedural Justice and the Law. *Philosophy Compass* 13: e12548. <https://doi.org/10.1111/phc3.12548>.
- Meyerson, D.; Mackenzie, C. and MacDermott, T. 2021. *Procedural Justice and Relational Theory: Empirical, Philosophical, and Legal Perspectives*. London; New York: Routledge; Taylor and Francis. <https://doi.org/10.4324/9780429317248>.
- Murphy, J. W. and Taylor, R. R. 2023. To democratize or not to democratize AI? That is the question. *AI and Ethics*. <https://doi.org/10.1007/s43681-023-00313-5>.
- Noble, S. U. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York, NY, US: New York University Press.
- Noorman, M., and Swierstra, T. 2023. Democratizing AI from a sociotechnical perspective. *Minds and Machines* 33: 563-586. <https://doi.org/10.1007/s11023-023-09651-z>.
- Pitkin, H. F. 1967. *The Concept of Representation*. Berkeley, Los Angeles, London: University of California Press.
- Pohlhaus, G. 2020. Epistemic Agency Under Oppression. *Philosophical Papers* 49(2): 233–51. <https://doi.org/10.1080/05568641.2020.1780149>.
- Rahwan, I. 2018. Society-in-the-Loop: Programming the Algorithmic Social Contract. *Ethics and Information Technology*, 20(1), 5–14. <https://doi.org/10.1007/s10676-017-9430-8>.
- Rawls, J. 1999. *A Theory of Justice: Revised Edition*. Cambridge, MA: Harvard University Press. <https://doi.org/10.2307/j.ctvkjb25m>.
- Rengers, M.; Houtekamer, C. and Maleeyakul, N. 2023. ‘Pas op Met Deze Visumaanvraag’, Waarschuwt het Algoritme Dat Discriminatie in de Hand Werkt. Het Ministerie Negeert Kritiek. <https://www.nrc.nl/nieuws/2023/04/23/beslisambtenarenblijven-profileren-met-risicoscores-a4162837>. Accessed: 13-05-2024.
- Sætra, H. S.; Borgebund, H.; and Coeckelbergh, M. 2022. Avoid Diluting Democracy by Algorithms. *Nature Machine Intelligence*, 4(10), Article 10. <https://doi.org/10.1038/s42256-022-00537-w>.
- Seger, E.; Ovadya, A.; Garfinkel, B.; Siddarth, D. and Dafoe, A. 2023. Democratizing AI: Multiple Meanings, Goals, and Methods. arXiv:2303.12642.
- Thibaut, J. W. and Walker, L. 1975. *Procedural Justice : a Psychological Analysis*. Hillsdale, N.J.: L. Erlbaum Associates.
- Tyler, T. R. 2006. *Why People Obey the Law*. Princeton, NJ: Princeton University Press. <https://doi.org/10.2307/j.ctvj66769>.
- Tyler, T. R. and Blader, S. L. 2003. The Group Engagement Model: Procedural Justice, Social Identity, and Cooperative Behavior. *Personality and Social Psychology Review*, 7(4), 349–361. [https://doi.org/10.1207/S15327957PSPR0704\\_07](https://doi.org/10.1207/S15327957PSPR0704_07).
- Volkskrant. 2021. ‘Belastingdienst Gebruikte Algoritme dat Lage Inkomens Selecteerde Voor Extra Fraudecontroles.’ <https://www.volkskrant.nl/gs-bac84336>. Accessed: 13-05-2024.
- Waldman, A. and Martin, K. 2022. Governing Algorithmic Decisions: The Role of Decision Importance and Governance on Perceived Legitimacy of Algorithmic Decisions. *Big Data & Society*, 9(1), <https://doi.org/10.1177/20539517221100449>.
- Warren, M. E. 2017. A Problem-Based Approach to Democratic Theory. *American Political Science Review*, 111(1), 39–53. <https://doi.org/10.1017/S0003055416000605>.
- Wong, P. 2020. Democratizing Algorithmic Fairness. *Philosophy & Technology* 33(2): 225–44. <https://doi.org/10.1007/s13347-019-00355-w>.
- Young, I. M. 1990. *Justice and the Politics of Difference. Revised*. Princeton, NJ: Princeton University Press. <https://doi.org/10.2307/j.ctvc4g4q>.
- Zimmermann, A.; Rosa, E. D.; and Kim, H. 2020, January 9. Technology Can’t Fix Algorithmic Injustice. <https://www.bostonreview.net/articles/annette-zimmermann-algorithmic-political/> Accessed: 13-05-2024.