

# A Human Rights-Based Approach to Responsible AI

VINODKUMAR PRABHAKARAN, Google Research, USA

MARGARET MITCHELL, Hugging Face, USA

TIMNIT GEBRU, DAIR, USA

IASON GABRIEL, Deepmind, UK

Research on fairness, accountability, transparency and ethics of AI-based interventions in society has gained much-needed momentum in recent years. However it lacks an explicit alignment with a set of normative values and principles that guide this research and interventions. Rather, an implicit consensus is often assumed to hold for the values we impart into our models – something that is at odds with the pluralistic world we live in. In this paper, we put forth the doctrine of universal human rights as a set of globally salient and cross-culturally recognized set of values that can serve as a grounding framework for explicit value alignment in responsible AI – and discuss its efficacy as a framework for civil society partnership and participation. We argue that a human rights framework orients the research in this space away from the machines and the risks of their biases, and towards humans and the risks to their rights, essentially helping to center the conversation around who is harmed, what harms they face, and how those harms may be mitigated.

## 1 INTRODUCTION

Fairness, accountability, transparency and ethics (FATE) research in AI contends with questions around how AI models might be problematically biased, unfair, or unethical, and how to make them “fairer” and more “ethical”. However, the research community working in this area greatly lacks in terms of geo-cultural diversity [27], resulting in the research being primarily framed in the Western context,<sup>1</sup> by researchers mostly situated in Western institutions/organizations, to mitigate social injustices prevalent in the West, using data from the West, and implicitly imparting Western value systems [66]. On the other hand, these research insights are meant to intervene on platforms that are globally present, serving a global population from diverse societies, cultures and values, with their own forms of injustices.

A core concern in this arrangement is that of value imposition, where local values, i.e., values that are local to the regions where the interventions are built, implicitly shape and inform global systems without any or much room for discussion or contestation from those affected by those interventions. More specifically, interventions designed to address FATE failures necessarily impart a normative value system, but the values that guide the proposed solutions are rarely recognized as sites of contestation. This is problematic because while there may be ethical principles for ML that garner a degree of consensus across different value systems, in a pluralistic world this consensus is not something that should be assumed. Instead, we need to be explicit about the values that underpin the quest for ethical and just AI, and to cultivate an active debate about those values, critically examining and evaluating claims about them [28].

Another shortcoming of not being explicit about what normative value systems shape the interventions is the vagueness it entails, making it harder to arrive at a common vocabulary and shared understanding between computer scientists and civil society. Such a shared understanding is crucial to bridge the gap between research and practice, especially in a way that effectively supports the priorities of the latter constituency. This is especially important given the need for critical examinations that require deeper understanding of the societal contexts in which interventions are

---

<sup>1</sup>In this paper, we use *Western* or *the West* to refer to the regions, nations and states, consisting of Europe, the United States and Canada, and Australasia and their shared norms, values, customs, religious beliefs, and political systems [46].

envisioned [7, 9, 36, 66] and the need for participatory methods to incorporate marginalized stakeholder perspectives [22, 44, 51] when shaping these interventions.

In this paper, we argue that the doctrine of *human rights* can serve as a starting point to help address some of these gaps. The role of human rights as a legal framework, through which AI-related harms may be identified, has been explicitly evoked by (or implicitly shaped) various AI accountability initiatives within the industry as well as governmental and civil society bodies (e.g., [16, 48, 52, 61]). However, we take a broader view of human rights and argue that, in addition to this legal role, they may play three further functions. First, human rights, understood as a set of moral claims, have a measure of intercultural and cross-cultural validity, which means that they can support value alignment for AI systems across a range of different national and social contexts. Second, they can help illuminate the concurrent responsibilities of various actors, given that they apply to states, organizations and individuals. Third, they provide a shared vocabulary and framework that technologists and practitioners may productively invoke to address the claims and concerns of the global civil society and the people impacted by these technologies. To clarify this distinction, we start by distinguishing three different aspects of human rights – (1) as a set of moral claims, (2) as a legal regime and set of instruments, and (3) as a cultural practice and global social movement. We then delve deeper into three specific rights enshrined in the Universal Human Rights Declaration (UDHR),<sup>2</sup> in order to illustrate how human rights might shape responsible AI development and deployments. While we acknowledge that a human rights perspective is not a panacea for addressing all issues in AI use, and that the provenance of UDHR in particular has been contested, we argue that a cross-cultural set of human rights, exemplified here by UDHR, can be a grounding framework for value alignment in AI.

## 2 HUMAN RIGHTS AND ETHICAL AI

From a philosophical perspective, human rights are the fundamental rights every human being holds simply by virtue of their birth, regardless of their age, ethnic origin, location, language, religion, nationality, ethnicity, or any other status.<sup>3</sup> The doctrine of human rights—which maps out their content, meaning, and consequences—has a number of different parts.

First and foremost, human rights can be understood as a set of *moral claims* anchored in the notion that human life has value and that there are aspects of personhood that must be protected. These claims can be grounded in a number of different ways but they tend to draw support from recognition that there are a set of core human interests that are widely shared, from common appreciation of the need to respect human autonomy and freedom, and from shared recognition of the dignity of human life [6, 24, 26, 34, 73]. In each case, there is a consideration that is strong enough to create a duty on other actors to only treat people in certain ways. In the words of Henry Shue, those who possess human rights can make justified demands that the actual enjoyment of a good be socially guaranteed against standard threats to those things [73]. Taken together, the idea of human rights holds that we owe it to each other to build a world in which the ability to enjoy certain goods, such as physical security, health and education, are widely enjoyed by all.

Second, human rights are part of a *legal regime and set of instruments* that aim (in part) to make these values a reality [12].<sup>4</sup> While the idea of universal rights has a long intellectual lineage, the modern conception of human rights emerged largely after the Second World War in response to the atrocities that had been committed and the moral trauma of

<sup>2</sup><https://www.un.org/en/about-us/universal-declaration-of-human-rights>

<sup>3</sup><https://www.ohchr.org/en/issues/pages/whatarehumanrights.aspx>

<sup>4</sup>Legal regimes refers to the different bodies of canonical law that apply in a context; for instance, there is an international legal regime, an EU legal regime, a US legal regime etc., which may overlap but are not necessarily precisely compatible.

the Holocaust [18]. This experience helped seed the desire to forge international agreement around the protection states owed to their citizens (and more widely), culminating in the adoption of the Universal Declaration of Human Rights (UDHR) by the United Nations General Assembly in 1948. The UDHR consists of 30 different articles affirming an individual’s rights, such as the right to life, liberty and security, right to privacy, right to be free of discrimination, and right to freedom of expression – rights that are critically relevant to building AI-based interventions responsibly. Drafted by a committee that included representatives from China, Chile, the Soviet Union, Lebanon, and India (in addition to the Western powers at the time), the UDHR was subsequently augmented by the International Covenant on Civil and Political Rights (ICCPR) and the International Covenant on Economic, Social and Cultural Rights (ICESCR) in 1966. These agreements incorporated insights from a still larger group of countries after the widespread success of decolonization movements. They now form part of customary international law and have been adopted by a number of national and transnational legal bodies such as the International Criminal Court (ICC).

Third, human rights are part of a *cultural practice and global social movement* that focuses on advocacy, empowerment and critique of existing institutions. From the very beginning, civil society has played a vital role in developing human rights doctrine, embedding norms, and monitoring outcomes [45]. As human rights advocates have long understood, it is not enough for rights to exist only in moral or juridical form: people also need to be informed about their rights and exercise them, if they are to have a significant bearing upon outcomes. Human rights advocacy has had a number of important successes, including in the field of women’s rights, indigenous rights, and disability rights, helping to shape global norms around how people may and may not be treated [64]. Civil society organizations have also explicitly and successfully used the human rights mechanisms to contest government policies (e.g., by making submissions to the UN’s Universal Periodic Review [31]).

Efforts are now underway to incorporate human rights into the design of AI systems, and to identify novel ways in which AI can support human rights practices. In applied contexts, there have been a number of attempts to integrate AI with human rights monitoring. These include efforts to improve human rights reporting in conflict zones using automated analysis of satellite imagery [47]. Additionally, in the context of discussions about fairness and AI ethics, human rights frameworks have been proposed as a tool or mechanism to promote greater accountability to those in need [16, 48, 52, 61]. In particular human rights advocates have criticised the tendency, sometimes present in FATE discourse, to imply that the relevant ethical norms for AI technology need to be discovered for the very first time – or to focus on sophisticated statistical analysis of algorithms without paying due attention to the way in which their deployment may result in actual societal harms [70]. Against these viewpoints, human rights doctrine contains a set of well-established principles that are robustly centered upon human vulnerabilities and human needs. Finally, there has been notable engagement with specific challenges to human rights posed by new technologies such as fake news moderation [49] and predictive policing [50]. For example, Data & Society organized a multidisciplinary workshop in April 2018 exploring how the human rights framework can effectively inform, shape, and govern AI research, development, and deployment. More recently, [17] organized tutorials at the FAT\* 2020 conference on this topic.<sup>5</sup> We draw inspiration from these groundbreaking efforts, and argue that more substantial engagement from the machine learning research community with the human rights paradigm can address three major challenges that the field of ML faces: alignment, the allocation of responsibilities, and participation.

---

<sup>5</sup>The conference has since been renamed to FAccT.

## 2.1 As a Basis for Cross-Cultural Value Alignment

Any effort to try and create ethical AI involves imparting a set of human values on the behavior of the AI-based systems under consideration, many of which are deployed across the world. While there has been plenty of research within AI sub-communities such as AI safety [2], machine learning [35], and natural language processing [15, 60], on the challenges involved in encoding human values into AI systems, less attention has been paid to the normative question of *which* set of values/principles should be encoded. Sometimes this question is overlooked altogether; in other cases system designers proceed on the basis of a preferred principle or theory. Both approaches are problematic given the wide variation in the moral beliefs people actually hold.<sup>6</sup> Indeed, since most FATE research in AI is situated in the West, it does not meaningfully engage with the conditions, values, and histories of non-West contexts [41, 57, 66]. As Birhane and Cummins note in [8], “it is possible that what is considered ethical currently and within certain domains for certain societies will not be received similarly at a different time, in another domain, or for a different society.” Thus, the Western values implicitly encoded into AI systems may be at odds with other value systems, creating the risk of problematic value imposition when these technologies are deployed globally [28].

In order to address this challenge, Gabriel [28] argues that we need to identify fair processes for selecting values to encode in AI systems. In this context, one promising approach focuses on the possibility of identifying an “overlapping consensus” between the different moral belief systems around the world. To make progress in this direction we need to ask: are there any ideals that command widespread, or even global, assent? Although no candidate doctrines is without limitations, with respect to cross-cultural validity, the doctrine of universal human rights is particularly promising [23, p. 150]. Indeed, while the modern notion of human rights has the specific historical lineage that we have outlined, a strong case can be made for the notion that these rights are no longer time-bound or culturally parochial. To start with, there is evidence that these beliefs find a degree of cross-cultural support in African, Islamic, Western, and Confucian traditions of thought [19, p335-343]. Moreover, human rights have been adopted and *actively claimed* by people around the world across a wide range of contexts ranging from indigenous rights movements [56] to the Arab Spring [38]. Hence, we believe human rights can serve as a legitimate goal for building value-aligned AI systems, a requirement that entails both that they need to respect human rights directly and also that they are deployed in ways that strengthen human rights practice.

Yet, the proposed focus on human rights faces four objections. First, it might be thought to lead to certain knowledge gaps: can a human rights framework adequately factor in collective goods, for example, or concerns about distributive justice? This criticism has often been made by Marxist scholars and contemporary critics of human rights doctrine such as Samuel Moyn [54]. In certain respects, the concerns they raise are well-founded: we agree that there may well be goods and considerations that are not adequately expressed in the language of rights [63]. Yet, our claim is not that human rights capture the full space of AI ethics but only that they capture a set of particularly important claims — a ‘morality of the depths’ [73] — for which there is widespread support. Understood in this way, respect for human rights would be necessary but not sufficient to ensure the ethical design and deployment of AI systems.

Second, scholars have challenged the purported “universality” of human rights by mobilising arguments that draw upon cultural relativism [80]. Proponents of this viewpoint note that moral norms vary significantly across cultures, and that what is “regarded as a human rights violation in one society may properly be considered lawful in another, and Western ideas of human rights should not be imposed upon Third World societies” [76]. However, as Donnelly notes, a more nuanced analysis of this matter resists the reduction of human rights doctrine to this single axis, and

---

<sup>6</sup>For instance, [www.worldvaluessurvey.org/](http://www.worldvaluessurvey.org/)

asks instead “*how* human rights are (and are not) universal and how they are (and are not) relative” [24]. In this context, he notes that UDHR has a very strong claim to what he terms “relative universality”, insofar as it represents a minimal response to basic cross-cultural human values and the threats to human dignity posed by modern institutions. At the same time, from a practical standpoint, the ways in which human rights interface with more specific cultural values is an important aspect to consider in the doctrine’s integration into AI ethics [82]. For instance, human rights doctrine may serve as an useful starting point that provides a baseline set of values that enjoy cross-cultural recognition, that can then be built on for specific contexts, especially when rights of specific communities are in question.

Third, despite what has been said, we might worry that human rights are still geographically parochial, or worse, a form of neo-imperialism [81]. In this context, it is important to recognize the historicity of human rights, the evolving character of human rights discourse, and the political undercurrent that influence how these norms are operationalized in practice. There is certainly evidence that human rights claims have been deployed in a strategic or politicized way by Western states or NGOs at certain times, and that they still operate under the burden of this legacy today [37]. However, this does not detract from the reality of their widespread support, or from the fact that they have been used equally in an opposing manner — to oppose authoritarian regimes, and in anti-colonial movements to resist external intervention [13, 40]. Moreover, Kathryn Sikkink notes that “voices and actors from the Global South were deeply involved in demanding the international protection of human rights and in building the institutions that started to make enforcement of these rights possible” [74, p. 25]. As a result of this process of contestation, something of value has emerged: modern human rights doctrine is often invoked by people to resist interference with their political rights or ability to govern themselves.

Finally, it might be thought that a human rights framework is too anthropocentric, emphasising the value of human life while neglecting non-human life and other bearers of value. However, as we noted before, insisting on the importance of human rights is not the same as asserting that only human rights matter: we agree other things matter as well. Moreover, we suggest that a renewed focus on human rights — and on the ways in which people can be seriously harmed — is particularly important for AI research at the present moment given the tendency to approach ethical challenges as technical problems [10]. After all, a great deal has been written about the formal properties of models and the biases they embody, however, the need to better understand the harms these models cause is critical [1, 4]. By way of illustration, Blodgett et al. conducted a survey of 146 research papers analyzing bias in natural language processing models and found only limited engagement with the question of why bias is harmful, in what ways, and to whom [11]. As a consequence, it is hard to be confident that proposed mitigations will have a positive impact on those who encounter these harms. By way of contrast, a human rights based approach to fairness research can help forge a stronger connection between models, the socio-technical systems they operate in, and salient harms. It supports a reorientation away from formal principles and towards human welfare and a person’s capacity to flourish.

## 2.2 As a Way of Understanding Responsibilities and Duties

The moral responsibilities entailed by human rights treaties apply to a wide range of actors and environments. States are often the primary duty-bearers, with respect for human rights being a key element of regime legitimacy [62]. Individuals also have duties not to violate human rights and there are a number of national and supranational mechanisms, such as the ICC and the European Court of Human Rights, to address this. Increasingly, human rights are also understood to apply to companies or organizations and to those employed by them — including for people working in the technology sector. These frameworks can therefore help us to understand who has a moral responsibility to do what, serving as a well-spring for responsible innovation in AI.

To begin with, human rights have ramifications for the way in which scientific and commercial research is conducted, especially when research concerns human subjects. This connection between bioethics and human rights was explicitly recognized in the aftermath of the Holocaust [3], leading to the development of safeguards for participants such as those described in the 1964 *Helsinki Declaration* and the 1979 *the Belmont Report*.<sup>7</sup> More recently, the *the Universal Declaration on Bioethics and Human Rights (UDBHR)*,<sup>8</sup> adopted by UNESCO’s General Conference in 2005, further affirmed the importance of human dignity, human rights and fundamental freedoms when it comes to research. Given recent controversy surrounding numerous publications in the field of ML and existing data collection practices, these principles, and the institutional review protocols they necessitate, serve as a valuable precedent for the broader ML research community [78].

The focus on human rights can also help us understand the responsibilities technology companies owe within the wider ecosystem of duty bearers. The 2011 United Nations Guiding Principles on Business and Human Rights (UNGPs) outline the specific responsibility of businesses to respect human rights, including by identifying, preventing, and mitigating salient human rights risks [61]. As the UNGPs make clear, when it comes to human rights, it is important to move beyond good intentions: those developing new technologies need to make an informed effort to understand the implications that a technology will have for rights holders – and to put in place measures that ensure that the rights are upheld through processes of evaluation, review and assessment.

### 2.3 As a Framework for Civil Society Participation

A major challenge for the AI ethics community centers upon the inclusion of relevant voices when addressing the design and governance of AI systems, and the cultivation of an effective *lingua franca* – or participatory processes – that make it easier for needs of historically marginalized communities to be articulated and for their claims to be effectively met. Against this backdrop, a key cluster of human rights centers upon and recognize the value of wide participation. For example, Articles 20 and 21 of the UDHR guarantee “freedom of peaceful assembly and association” and “freedom to participate in political processes” respectively. Moreover, the idea that people are “rights-holders” serves as a basis for engagement and source of authority that is quite different from engagement as “consumers”, “citizen[s]”, “stakeholders” or “affected parties”. Taken together, these two interrelated elements of participation and empowerment explain why approaches towards responsible AI, pioneered by civil society organizations, have tended to invoke human rights more often than fairness research situated within the machine learning community.

For instance, the *Mapping Consensus in Ethical and Rights-based Approaches to Principles for AI* report published by the Berkman Klein Center for Internet & Society [25] found that among thirty-six different sets of AI principles published by private and public agencies, human rights are a major focus – with civil society and trans-national governmental agencies relying most heavily on this framework. Documents drafted by trans-national governmental agencies such as *AI for Europe* by the European Commission and *Ethics Guidelines for Trustworthy AI* by the European High Level Expert Group on AI also foreground human rights. And out of the five civil society-drafted AI principles documents, three of them – *Toronto Declaration* by Amnesty International & Access Now,<sup>9</sup> *Universal Guidelines for AI* by The Public Voice Coalition,<sup>10</sup> and *Human Rights in the Age of AI* by Access Now,<sup>11</sup> – explicitly adopt a human

<sup>7</sup><https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/read-the-belmont-report>

<sup>8</sup><https://en.unesco.org/themes/ethics-science-and-technology/bioethics-and-human-rights>

<sup>9</sup><https://www.accessnow.org/the-toronto-declaration-protecting-the-rights-to-equality-and-non-discrimination-in-machine-learning-systems>

<sup>10</sup><https://thepublicvoice.org/ai-universal-guidelines/>

<sup>11</sup><https://www.accessnow.org/cms/assets/uploads/2018/11/AI-and-Human-Rights.pdf>

rights framework, while a fourth one, *Top 10 Principles for Ethical AI* by UNI Global Union,<sup>12</sup> includes discussion of human rights risks.

By way of contrast, research done within the ML fairness, accountability, transparency and ethics research community rarely invokes a rights-based framework. Drawing upon a survey of 138 papers/abstracts published in the FAT\* conference in 2019 and 2020, we found only one research paper [43] and 2 tutorial abstracts [17, 75] that engage with the human rights scholarship.<sup>13</sup> More precisely [43] proposes an impact assessment methodology that they situate within the human rights assessment literature, [17] conducted a translation tutorial between human rights scholarship and FATE research and [75] conducted a hands-on tutorial where they used a human rights framing to test academic concepts and their formulation in policy initiatives around algorithmic accountability and explainability.

In addition to the substantive value of the human rights framework, the salience that human rights have both for people affected by new technologies and in policy circles provides additional reason to close this significant gap in the AI FATE literature. For AI researchers to build effective partnerships with civil society, bridging work needs to be done. In certain domains bridging work is already underway, for instance between the community of NLP researchers working in the space of detecting online abuse and the RightsCon community [59]. Yet, failure to engage with human rights scholarship more widely, risks leading to a situation in which FATE researchers end up ‘speaking a different language’ from those their products affect, making it harder to conduct participatory research with civil society organizations and foregoing a major opportunity to strengthen the practice of AI ethics more widely.

### 3 ILLUSTRATIONS

The potential impact of AI on human rights is wide-ranging, something that is recognized by recent commentary on the human right to science by the UN Economic and Social Council which notes ‘applications of artificial intelligence in industry or services can lead to enormous gains in productivity and efficiency’ while also expressing concern that algorithms could be incorporated into weapon systems or used to reinforce discrimination.<sup>14</sup> Rather than attempt to catalogue the full range of impacts AI might have, we focus in this section on three human rights in particular, to show what a human rights framework may add to the responsible AI discussion. For each right, we walk through specific examples of how the values at play may influence decisions in an algorithmic context.

#### 3.1 The Human Right to Freedom from Discrimination

The right to be free from discrimination is a negative right not to be harmed in certain ways, and it is heavily impacted by prevalent forms of algorithmic bias. This right is enshrined in various universal and regional legal instruments of human rights, and forms one of the core rights in the UDHR. In particular, article 2 of the UDHR states: “everyone is entitled 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”, essentially extending all the rights enshrined in the declaration to all humans without discrimination.

The right to be free from discrimination is also often the right that is most directly relevant to a majority of work in the FATE community. Fairness research has identified numerous types of biases in various algorithmic systems [5, 11]. An algorithmic system that treats individuals differently based on an attribute such as race or gender, with negative

---

<sup>12</sup><http://www.thefutureworldofwork.org/opinions/10-principles-for-ethical-ai/>

<sup>13</sup>12 other research papers mention the phrase *human rights*, but do not engage with it beyond that.

<sup>14</sup><https://undocs.org/E/C.12/GC/25>

consequences for group members and without due cause, is in and of itself a violation of the right to be free of discrimination. However, certain instances of discrimination that result in withholding other rights have a compounding effect. For instance, an algorithmic content moderation system that disproportionately censors individuals speaking a certain dialect [67] not only risks interfering with their right to freedom of expression, but also potentially impedes their right to be free from discrimination in exercising that right.

Since early research into algorithmic fairness dealt with the applications of AI in US law enforcement such as predictive policing and recidivism prediction, and regulations around anti-discrimination in housing, loans, and education in the US, the FATE community's inquiries into this space draws largely on US legal frameworks such as the Civil Rights Acts and Fair Housing Act, as well as on US legal concepts of discrimination [33]. One side effect of this is that the scope of this conversation has also been largely limited to discrimination in the US context, compared to the more global human rights framework. For instance, while the right to be free of discrimination based on religion or language has equal standing within the global human rights framework, these axes of discrimination are rarely dealt with within the FATE research community, compared to discrimination based on race or gender, which are prominent concerns in the US American public discourse. This gap, in terms of understanding the full range of characteristics, that may serve as axes of unjust discrimination, can be addressed, in part, through reflection on the more expansive categorisation invoked by the UDHR.

Furthermore, indexing FATE research on the legal framework of a particular country carries with it additional risk, when viewed from a global standpoint. For example, a single country's legal frameworks may not give all groups adequate protection against discrimination, an issue that looms particularly large in the context of unjust laws or national practices. To guard against these pitfalls, a more universal set of principles such as the universal human rights doctrine, may serve as a useful reference point, and as means of checking national laws this kind of gaps or denial of equal rights to citizens.

### 3.2 The Human Right to Health

The human right to health enshrines “the right of everyone to the enjoyment of the highest attainable standard of physical and mental health” (ICESAR, 1966, Art. 12.1).<sup>15</sup> It works primarily as a positive right that creates a duty for states to lower infant mortality, promote child development, provide medical services to their populations, and share medical knowledge, among other things. Given the growth of ML-enabled services and diagnostic tools within the healthcare sector [20, 83], AI research has the potential to intersect with this right in important ways.

The right to health grounds an entitlement to access health care on terms that are free from discrimination, with particular attention being paid to protected groups such as women and those with physical disabilities. Given evidence that racial bias affects algorithms deployed in a healthcare context [55], measures to mitigate the harms that biases can create prior to an algorithm's deployment can be understood as a human rights obligation. Additionally, the duties that correspond to the right to health make reference to that “highest attainable standard” of care, a specification that acknowledges that the content of the right will vary dynamically according to time and place. For example, governments have an obligation to give their citizens the highest attainable standard of health care. But if the country is very poor, then the required standard of treatment might still be quite low. Contrastingly, if AI services can be used to bring down the cost of healthcare, or the lack of access to it, then what is 'attainable' for the country may begin to rise, and the right to health may become a right to access and benefit from AI services. This invites us to think about the

---

<sup>15</sup><https://www.ohchr.org/en/professionalinterest/pages/cescr.aspx>



relationship between healthcare and AI in a different light – not only as a source of potential harms but through the lens of human rights-enabling technology.

Current estimates suggest that there are not enough trained doctors and physicians globally to provide everyone in the world with a high level of medical care, a problem that is particularly true in low-income countries [68]. This shortfall could be addressed through the global redistribution of economic resources [14]. However, it is also something that AI researchers are in a special position to influence, through the creation of diagnostic tools that can be deployed affordably at scale and the development of customized digital healthcare services. For example, ML-enabled technology is now being used to detect diabetic retinopathy at scale in India [83]. A similar point, about the creation of potentially low-cost tools and services, holds true for the human right to education. AI-enabled services could, in principle, make it possible for many more children to enjoy customised education in their own language, thereby improving global literacy and learning.

In both cases it is important to steer clear of the pitfalls of technological solutionism [53]. As human rights advocates note, the exercise of human rights and impediments to them are frequently political in nature [29, 71]. Yet these opportunities also ground a positive aspiration for AI: that it will expand the feasibility frontier so that people can enjoy a higher standard of human rights fulfilment around the world. Moreover, this goal dovetails with the Sustainable Development Goals [77] and would likely find widespread support among those who experience limited access to services.

### 3.3 The Human Right to Share in Scientific Advancement

The human right to science states that everyone has a right to “share in scientific advancement and its benefits” (UDHR, Art. 27). Science is understood here to include: (1) knowledge, (2) the application of that knowledge, and (3) the method of the knowledge production. Moreover, the right applies both to scientific knowledge itself and to the benefits it creates. This right has special relevance for AI both because it can be understood as a scientific practice and also because of the way in which AI is increasingly used to advance scientific progress, as with DeepMind’s AlphaFold which successfully predicted the structure of almost every known protein [72].

The human right to science advances an ideal of inclusive science and prohibits discrimination both among those employed in scientific pursuits and among its beneficiaries. On this point Michelle Bachelet, UN High Commissioner for Human Rights affirms that “those participating in the global scientific effort should... take into account the needs and experiences of women, members of minority communities, Indigenous scholars, persons with disabilities, people living in poverty and people living in less developed countries – among others. Only then will research fully address all communities – and contribute to reducing the unequal access to scientific developments and capabilities across different countries and regions”.<sup>16</sup> This element of the right to science underscores the need to systematically promote diversity, equity and inclusion in AI research.

Significantly, the human right to science also bears upon the distribution of benefits enabled by science. The UDHR clearly states that the right is to be interpreted in a way that respects the intellectual property of researchers. However, it situates these “moral and material interests” within the wider aspiration that science should be geared towards fulfilment of human rights. Speaking on behalf of CERN and the WHO among others, Bachelet states that “the benefits of scientific and medical progress were always meant to be shared. The great beauty of science is that it has no borders – and that, working together, every scientist and student of science can contribute to the shared knowledge and benefit

---

<sup>16</sup><https://www.ohchr.org/EN/HRBodies/HRC/Pages/NewsDetail.aspx?NewsID=26433&LangID=E>

of all". In the context of AI research, the human right to science asserts the importance of this technology ultimately benefiting a large section of humanity, including those historically excluded from the benefits of scientific advances. Without the assertion of this right, scientific progress may otherwise continue to exclude insights, ideas, and concerns from people who are historically excluded, pulling AI research farther from a trajectory that supports the global social good.

#### 4 DISCUSSION

In this paper, we have sought to show how greater attention to human rights can help ground the aspiration to build more ethical AI systems, anchored in values that have a measure of cross-cultural affirmation. Orienting the research in this space away from the machines and the risks of their biases, and towards humans and the risks to their rights, can help center conversation around the harms caused by a technology, including specific consideration of who is harmed, and how those harms may be mitigated. This reframing also has the potential to better align efforts by the FATE community to improve AI systems with the wider global advocacy movement that is committed to securing human rights and their fulfilment.

In support of these goals, future research on human rights-based approaches towards to AI ethics could help bridge the gap in multiple ways. First, there is a need for **translational** research in this space that can address, more precisely, how human rights principles map to the current ethics-based considerations in AI, such as fairness, consent, privacy, and ownership. Such research could be aimed at building a shared vocabulary of concerns, values, and expected outcomes as an important first step for meaningful bridging between computer science researchers and civil society activists working in this space. This should crucially include clarifying the needs of civil society that are overlooked and potentially easily addressed through technology, as well as the challenges of ensuring fairness of algorithmic predictions at scale.

Another line of work could look into the **functional** aspects of a human rights based inquiry into AI ethics. For instance, what does a human rights based inquiry into ML fairness reveal that existing methodologies do not. The worked example in Appendix A in the context of online content moderation demonstrates some of these functional aspects of a human rights based approach. In particular, this approach pushes us to identify the rights holders, to identify which of their rights are at risk, and to map out how those risks interact with the claims of other right holders. Thus, it is essential for AI researchers to contend with various trade-offs, when determining how to intervene. For instance, fairness researchers need to consider how bias mitigation measures, designed to mitigate risks to the human rights of content creators, might create or increase risks to certain other rights of the audience.

Finally, advances in AI have the potential to play an important role in **enabling** stronger human rights fulfilment around the world. As we lay out in Section 3, AI has already been shown promise when it comes to extending the scope and content of various human rights (such as the right to health and right to scientific advancements) to marginalized communities. Future research should consider what role AI can play in enhancing access to health and education for communities in lower-income countries. Similarly, AI-based technologies such as automatic captioning might play an important role in increasing access to education for people with disabilities. There is more work to be done in this space, not only in employing AI in rights-enabling applications, but also in bringing the advancements in AI and technology to communities around the world, rather than keeping it only within the reach of a select few.

## 5 CONCLUSION

AI ethics does not always give due recognition to the idea that every human life has value and also that human life is fragile – considerations that ground a set of important moral claims on institutions, new technologies, and on one another. The notion that people have human rights builds upon this foundation, recognizing that we all share certain vulnerabilities, that we ought not to be harmed, and these considerations guide how AI should be developed. In this paper we have suggested that human rights-based considerations can perform three valuable functions in the context of AI research. First, human rights can serve as a partial basis for AI value alignment across a range of cultures and different contexts, due to the qualified but significant cross-cultural validity that they evidence. Second, a human rights framework can help us understand how ethical principles governing the design and deployment of AI systems translate into different responsibilities for the actors that comprise different parts of the AI ecosystem. Third, human rights can function as a language that enables deeper collaboration between AI researchers, civil society groups and the people impacted by these technologies. In this way it can help close the gap between technical research focusing on algorithmic fairness, and the claims of those who interact with AI systems on the ground and in the public sphere. Taken together, these elements of human rights doctrine make it an appealing set of guiding principles for AI researchers and practitioners to draw upon.

## ACKNOWLEDGEMENTS

We thank Roya Pakzad, Jamila Smith-Loud, Tan Zhi Xuan, and Ben Zevenbergen for helpful conversations on this topic and for useful feedback on early drafts of this paper. We also thank the anonymous reviewers for their constructive feedback.

## REFERENCES

- [1] Micah Altman, Alexandra Wood, and Effy Vayena. 2018. A harm-reduction framework for algorithmic fairness. *IEEE Security & Privacy* 16, 3 (2018), 34–45.
- [2] Thomas Arnold, Daniel Kasenberg, and Matthias Scheutz. 2017. Value Alignment or Misalignment—What Will Keep Systems Accountable?. In *Workshops at the Thirty-First AAAI Conference on Artificial Intelligence*.
- [3] Robert Baker. 2001. Bioethics and human rights: A historical perspective. *Cambridge Quarterly of Healthcare Ethics* 10, 3 (2001), 241–252.
- [4] Solon Barocas, Kate Crawford, Aaron Shapiro, and Hanna Wallach. 2017. The problem with bias: from allocative to representational harms in machine learning. Special Interest Group for Computing. *Information and Society (SIGCIS)* (2017).
- [5] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2017. Fairness in machine learning. *Nips tutorial* 1 (2017), 2.
- [6] Charles R Beitz. 2009. *The idea of human rights*. OUP Oxford.
- [7] Sebastian Benthall and Bruce D Haynes. 2019. Racial categories in machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 289–298.
- [8] Abeba Birhane and Fred Cummins. 2019. Algorithmic Injustices: Towards a Relational Ethics. *arXiv preprint arXiv:1912.07376* (2019).
- [9] Abeba Birhane and Olivia Guest. 2020. Towards decolonising computational sciences. *arXiv preprint arXiv:2009.14258* (2020).
- [10] Abeba Birhane and Jelle van Dijk. 2020. Robot rights? Let’s talk about human welfare instead. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 207–213.
- [11] Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of “bias” in nlp. *arXiv preprint arXiv:2005.14050* (2020).
- [12] Allen Buchanan. 2013. *The heart of human rights*. Oxford University Press.
- [13] Roland Burke. 2011. *Decolonization and the evolution of international human rights*. University of Pennsylvania Press.
- [14] Simon Caney. 2006. Global justice: From theory to practice. *Globalizations* 3, 2 (2006), 121–137.
- [15] Dallas Card and Noah A Smith. 2020. On Consequentialism and Fairness. *Frontiers in Artificial Intelligence* 3 (2020), 34.
- [16] Corinne Cath. 2018. Governing artificial intelligence: ethical, legal and technical opportunities and challenges.
- [17] Corinne Cath, Mark Latonero, Vidushi Marda, and Roya Pakzad. 2020. Leap of FATE: Human Rights as a Complementary Framework for AI Policy and Practice. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT\* ’20). Association for Computing Machinery, New York, NY, USA, 702. <https://doi.org/10.1145/3351095.3375665>

- [18] Richard Pierre Claude and Burns H Weston. 1992. *Human rights in the world community: issues and action*. University of Pennsylvania Press.
- [19] Joshua Cohen. 2010. *The arc of the moral universe and other essays*. Harvard University Press.
- [20] Jeffrey De Fauw, Pearse Keane, Nenad Tomasev, Daniel Visentin, George van den Driessche, Mike Johnson, Cian O Hughes, Carlton Chu, Joseph LedSAM, Trevor Back, et al. 2016. Automated analysis of retinal imaging using machine learning techniques for computer vision. *F1000Research* 5 (2016).
- [21] Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. Measuring and mitigating unintended bias in text classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. 67–73.
- [22] Joseph Donia and James A Shaw. 2021. Co-design and ethical artificial intelligence for health: An agenda for critical research and practice. *Big Data & Society* 8, 2 (2021), 20539517211065248.
- [23] Jack Donnelly. 2007. The relative universality of human rights. *Human rights quarterly* (2007), 281–306.
- [24] Jack Donnelly. 2013. *Universal human rights in theory and practice*. Cornell University Press.
- [25] Jessica Fjeld, Nele Achten, Hannah Hilligoss, Adam Nagy, and Madhulika Srikumar. 2020. Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI. *Berkman Klein Center Research Publication* (2020).
- [26] Michael Freeman. 2022. *Human rights*. John Wiley & Sons.
- [27] Ana Freire, Lorenzo Porcaro, and Emilia Gómez. 2020. Measuring diversity of artificial intelligence conferences. *arXiv preprint arXiv:2001.07038* (2020).
- [28] Iason Gabriel. 2020. Artificial Intelligence, Values and Alignment. *arXiv preprint arXiv:2001.09768* (2020).
- [29] Iason Gabriel and Brian McElwee. 2019. Effective altruism, global poverty, and systemic change. *Effective altruism: Philosophical issues* (2019), 99.
- [30] Kate Goddard, Abdul Roudsari, and Jeremy C Wyatt. 2012. Automation bias: a systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association* 19, 1 (2012), 121–127.
- [31] James Gomez and Michelle D’cruz. 2018. Singapore’s Universal Periodic Review: Civil Society Trends and Themes. In *The Universal Periodic Review of Southeast Asia*. Springer, 115–136.
- [32] Mary L Gray and Siddharth Suri. 2019. *Ghost work: How to stop Silicon Valley from building a new global underclass*. Eamon Dolan Books.
- [33] Ben Green and Salomé Viljoen. 2020. Algorithmic realism: expanding the boundaries of algorithmic thought. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 19–31.
- [34] James Griffin. 2009. *On human rights*. OUP Oxford.
- [35] Dylan Hadfield-Menell, Stuart J Russell, Pieter Abbeel, and Anca Dragan. 2016. Cooperative inverse reinforcement learning. In *Advances in neural information processing systems*. 3909–3917.
- [36] Alex Hanna, Emily Denton, Andrew Smart, and Jamila Smith-Loud. 2020. Towards a critical race methodology in algorithmic fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 501–512.
- [37] Michael Hardt and Antonio Negri. 2000. *Empire*. Harvard University Press.
- [38] Julie Harrelson-Stephens and Rhonda L Callaway. 2014. You say you want a revolution: The Arab Spring, norm diffusion, and the human rights regime. *Human Rights Review* 15, 4 (2014), 413–431.
- [39] Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. 2020. Social Biases in NLP Models as Barriers for Persons with Disabilities. *ACL* (2020).
- [40] Steven LB Jensen. 2016. *The making of international human rights: the 1960s, decolonization, and the reconstruction of global values*. Cambridge University Press.
- [41] Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1, 9 (2019), 389–399.
- [42] David Jurgens, Libby Hemphill, and Eshwar Chandrasekharan. 2019. A Just and Comprehensive Strategy for Using NLP to Address Online Abuse. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 3658–3666.
- [43] Margot E. Kaminski and Gianclaudio Malgieri. 2020. Multi-Layered Explanations from Algorithmic Impact Assessments in the GDPR. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT\* ’20). Association for Computing Machinery, New York, NY, USA, 68–79. <https://doi.org/10.1145/3351095.3372875>
- [44] Michael Katell, Meg Young, Dharma Dailey, Bernease Herman, Vivian Guetler, Aaron Tam, Corinne Bintz, Daniella Raz, and PM Krafft. 2020. Toward situated interventions for algorithmic equity: lessons from the field. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 45–55.
- [45] Margaret E Keck, Kathryn Sikkink, et al. 1998. *Activists beyond borders: Advocacy networks in international politics*. Cornell University Press.
- [46] James Kurth. 2003. Western civilization, our tradition. *Intercollegiate Review* 39, 1/2 (2003), 5.
- [47] Steven Livingston and Mathias Risse. 2019. The future impact of artificial intelligence on humans and human rights. *Ethics & international affairs* 33, 2 (2019), 141–158.
- [48] Vidushi Marda. 2018. Artificial intelligence policy in India: a framework for engaging the limits of data-driven decision-making. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376, 2133 (2018), 20180087.
- [49] Vidushi Marda and Stefania Milan. 2018. Wisdom of the Crowd: Multistakeholder perspectives on the fake news debate. *Internet Policy Review series, Annenberg School of Communication* (2018) (2018).
- [50] Vidushi Marda and Shivangi Narayan. 2020. Data in New Delhi’s predictive policing system. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 317–324.

- [51] Donald Martin Jr, Vinod Prabhakaran, Jill Kuhlberg, Andrew Smart, and William S Isaac. 2020. Participatory Problem Formulation for Fairer Machine Learning Through Community Based System Dynamics. *arXiv preprint arXiv:2005.07572* (2020).
- [52] Lorna McGregor, Daragh Murray, and Vivian Ng. 2019. International human rights law as a framework for algorithmic accountability. *International & Comparative Law Quarterly* 68, 2 (2019), 309–343.
- [53] Evgeny Morozov. 2013. *To save everything, click here: The folly of technological solutionism*. Public Affairs.
- [54] Samuel Moyn. 2018. *Not enough*. Harvard University Press.
- [55] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
- [56] Gaetano Pentassuglia. 2011. Towards a jurisprudential articulation of indigenous land rights. *European Journal of International Law* 22, 1 (2011), 165–202.
- [57] Marie-Therese Png. 2022. At the Tensions of South and North: Critical Roles of Global South Stakeholders in AI Governance. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. 1434–1445.
- [58] Vinodkumar Prabhakaran, Ben Hutchinson, and Margaret Mitchell. 2019. Perturbation Sensitivity Analysis to Detect Unintended Model Biases. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 5740–5745. <https://doi.org/10.18653/v1/D19-1578>
- [59] Vinodkumar Prabhakaran, Zeerak Waseem, Seyi Akiwowo, and Bertie Vidgen. 2020. Online abuse and human rights: WOAHSatellite session at RightsCon 2020. In *Proceedings of the Fourth Workshop on Online Abuse and Harms*. 1–6.
- [60] Shrimai Prabhumoye, Brendon Boldt, Ruslan Salakhutdinov, and Alan W Black. 2021. Case Study: Deontological Ethics in NLP. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 3784–3798.
- [61] Filippo A Raso, Hannah Hilligoss, Vivek Krishnamurthy, Christopher Bavitz, and Levin Kim. 2018. Artificial Intelligence & Human Rights: Opportunities & Risks. *Berkman Klein Center Research Publication* (2018).
- [62] John Rawls. 1993. The law of peoples. *Critical Inquiry* 20, 1 (1993), 36–68.
- [63] Joseph Raz. 1986. Right-based moralities. In *Rights and Reason*. Springer, 177–196.
- [64] Thomas Risse, Thomas Risse-Kappen, Stephen C Ropp, and Kathryn Sikkink. 1999. *The power of human rights: International norms and domestic change*. Vol. 66. Cambridge University Press.
- [65] Sarah T Roberts. 2016. Commercial content moderation: Digital laborers’ dirty work. (2016).
- [66] Nithya Sambasivan, Erin Arnesen, Ben Hutchinson, Tulsee Doshi, and Vinodkumar Prabhakaran. 2021. Re-imagining algorithmic fairness in india and beyond. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 315–328.
- [67] Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A Smith. 2019. The risk of racial bias in hate speech detection. In *Proceedings of the 57th annual meeting of the association for computational linguistics*. 1668–1678.
- [68] Richard M Scheffler, Jenny X Liu, Yohannes Kinfu, and Mario R Dal Poz. 2008. Forecasting the global shortage of physicians: an economic-and needs-based approach. *Bulletin of the World Health Organization* 86 (2008), 516–523B.
- [69] Katrina Scior. 2011. Public awareness, attitudes and beliefs regarding intellectual disability: A systematic review. *Research in developmental disabilities* 32, 6 (2011), 2164–2182.
- [70] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 59–68.
- [71] Amartya Sen. 2005. Human rights and capabilities. *Journal of human development* 6, 2 (2005), 151–166.
- [72] Andrew W Senior, Richard Evans, John Jumper, James Kirkpatrick, Laurent Sifre, Tim Green, Chongli Qin, Augustin Židek, Alexander WR Nelson, Alex Bridgland, et al. 2020. Improved protein structure prediction using potentials from deep learning. *Nature* 577, 7792 (2020), 706–710.
- [73] Henry Shue. 2020. *Basic rights: Subsistence, affluence, and US foreign policy*. Princeton University press.
- [74] Kathryn Sikkink. 2017. *Evidence for Hope: Making Human Rights Work in the 21st Century*. Princeton University Press, Princeton. <https://doi.org/10.1515/9781400888535>
- [75] Katarzyna Szymielewicz, Anna Bacciarelli, Fanny Hidvegi, Agata Foryciarz, Soizic Pénicaud, and Matthias Spielkamp. 2020. Where Do Algorithmic Accountability and Explainability Frameworks Take Us in the Real World? From Theory to Practice. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (Barcelona, Spain) (FAT\* ’20)*. Association for Computing Machinery, New York, NY, USA, 689. <https://doi.org/10.1145/3351095.3375683>
- [76] Fernando R Tesón. 1985. International Human Rights and Cultural Relativism. *Virginia Journal of International Law* 25 (1985), 869.
- [77] Nenad Tomašev, Julien Cornebise, Frank Hutter, Shakir Mohamed, Angela Picciariello, Bec Connelly, Danielle CM Belgrave, Daphne Ezer, Fanny Cachat van der Haert, Frank Mugisha, et al. 2020. AI for social good: unlocking the opportunity for positive impact. *Nature Communications* 11, 1 (2020), 1–6.
- [78] Richard Van Noorden. 2020. The ethical questions that haunt facial-recognition research. *Nature* 587, 7834 (2020), 354–359.
- [79] Andreas Veglis. 2014. Moderation techniques for social media content. In *International conference on social computing and social media*. Springer, 137–148.
- [80] Dunstan M Wai. 1979. Human rights in sub-Saharan Africa. *Human rights: cultural and ideological perspectives* 115 (1979), 116.
- [81] Christopher Wall. 1998. Human Rights and Economic Sanctions: The New Imperialism. *Fordham Int’l LJ* 22 (1998), 577.

- [82] Pak-Hang Wong. 2020. Cultural differences as excuses? Human rights and cultural values in global ethics and governance of AI. *Philosophy & Technology* 33, 4 (2020), 705–715.
- [83] Kun-Hsing Yu, Andrew L Beam, and Isaac S Kohane. 2018. Artificial intelligence in healthcare. *Nature biomedical engineering* 2, 10 (2018), 719–731.

## APPENDIX

### A WORKED EXAMPLE: ONLINE CONTENT MODERATION

Having outlined how some of the specific rights are of core relevance to AI ethics conversations, we now dive more deeply into a specific problem domain to demonstrate how the consideration of human rights in AI based interventions can be applied in practice. More specifically, we use the application domain of online content moderation, focusing on a hypothetical platform where people write and read one another’s text. The domain of content moderation is an incredibly complex one, with challenges around detecting online abuse and toxic content at scale and across geo-cultural contexts, while also accounting for conversational contexts, and a multitude of intervention approaches towards content removal or demotion, as well as concerns on transparency and accountability [59]. Our choice of this domain is motivated by these complexities, as it helps to demonstrate the various competing considerations involved, including different sets of rights-holders, and the different kinds of risks to their rights. This section is not meant as a blueprint of a solution to tackle this complex problem, rather as a way to show how a human rights based approach illuminates some of these competing considerations, and centers the conversation around these considerations that are often overlooked in traditional approaches. Also, note that this discussion is not limited to just the three rights discussed as illustrative examples above, rather we draw from the full set of UDHR rights.

#### A.1 Rights-holders

One important way a human rights-centered perspective changes the approaches towards responsible AI is that it forces us to confront the question of whose rights are at risk and what those risks are. As Blodgett [11] points out, research on fairness failures in AI tend to sidestep answering these questions. In the context of our current work, this pushes us to distinguish *rights-holders*, i.e., those whose rights are at risk, from stakeholders, i.e., anyone who can claim a stake in the system. In the case of online content moderation, the rights-holders include:

- the **CREATORS** of the content,
- the **ADDRESSEES** to whom the content is specifically directed,
- the **AUDIENCES** on the platform who are exposed to the content, and
- the **REFERRENTS** who are implied or denoted in the content,
- the **MODERATORS** who may have to review the content.

Moreover, the rights-holders can be an individual, a set of individuals, a population subgroup, or a community of people. The platform owners are a stakeholder, but they are not relevant rights-holders in this example since their human rights are not at risk in these scenarios which focus on online communication.

In the rest of this section, we incrementally analyze different scenarios, increasing the complexity of an example text content moderation system that is used on a platform for online communication. We consider the relevant rights and rights-holders given the *presence* of problematic online text (Scenario 1), followed by the relevant rights when *automatic content moderation* is applied (Scenario 2), then how individuals’ rights are affected when addressing *biases in text content moderation models* (Scenario 3), and finally with the addition of *human-in-the-loop content moderation* (Scenario 4). We primarily focus on the rights and right-holders with respect to the set of human rights enshrined in the

Right	Abridged Description	Rights-holders			
UDHR§5	<i>No one shall be subjected to torture or to cruel, inhuman or degrading treatment or punishment.</i>	ADDRESSEES			
UDHR§1	<i>People should act towards one another in a spirit of brotherhood</i>	ADDRESSEES	CREATORS		
UDHR§3	<i>Everyone has the right to...liberty and the security of person.</i>	ADDRESSEES	CREATORS	REFERRENTS	
UDHR§25	<i>Everyone has the right to a standard of living adequate for the health and well-being of himself and of his family.</i>	ADDRESSEES		REFERRENTS	
UDHR§12	<i>No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honour and reputation.</i>				
UDHR§20	<i>Everyone has the right to freedom of peaceful assembly and association.</i>	ADDRESSEES	AUDIENCE	REFERRENTS	
UDHR§7	<i>...All are entitled to equal protection against any discrimination...[and] incitement to such discrimination.</i>	ADDRESSEES	AUDIENCE		
UDHR§18	<i>Everyone has the right to freedom of thought...[and to] manifest his religion or belief in teaching, practice...</i>	ADDRESSEES	AUDIENCE	CREATORS	
UDHR§19	<i>Everyone has the right to freedom of opinion and expression...includes freedom to hold opinions without interference and to seek, receive and impart information and ideas through any media...</i>				
UDHR§29	<i>Everyone has duties to the community in which alone the free and full development of his personality is possible.</i>	ADDRESSEES	AUDIENCE	CREATORS	REFERRENTS
UDHR§27	<i>Everyone has the right freely to participate in the cultural life of the community...to the protection of the moral and material interests resulting from any...[literary production] of which he is the author.</i>	CREATORS			
UDHR§23	<i>Everyone has the right to...just and favourable conditions of work...to form and to join trade unions for the protection of his interests.</i>	MODERATORS			

Table 1. Abridged rights relevant to the different parties in an online communication platform with content moderation (non-exhaustive).

UDHR, however particular contexts might require considerations of additional sets of rights, e.g., the UN Declaration on the Rights of Indigenous Peoples may be more appropriate in the context of an AI intervention where indigenous people are a rights-holder.

A summary of some of the relevant UDHR rights at play with respect to each right-holder in this example is presented in Figure 1. ADDRESSEES, AUDIENCE, CREATORS, and REFERRENTS are affected differently depending on how the rights are prioritized. Interventions to tackle harmful content online must attempt to untangle these tensions and mitigate the risks, balancing the rights at play.

**Scenario 1: Presence of Harmful Content** The CREATORS who produce harmful content are owners of the content, describing the REFERRENTS and communicating to the ADDRESSEES in front of an AUDIENCE. In producing content, CREATORS are exercising their rights for freedom of opinion and expression (UDHR§18,19). Yet, determining which content may be considered “harmful” can stem from whether it is at odds with the need to respect the human rights of others, i.e., whether the content causes degrading treatment (UDHR§5).

If a CREATOR's content reveals private and sensitive information about a specific individual, that individual is a REFERRENT whose rights to privacy and reputation (UDHR§12) is in question. If a CREATOR's content spreads dangerous misinformation about a community, the REFERRENTS are everyone in the community, and may violate their right to security (UDHR§3). ADDRESSEES have rights similar to the REFERRENTS, which additionally include rights for security (UDHR§3) against anyone who might harm them (such as an incited AUDIENCE) and a standard of living that is adequate for well-being (UDHR§25). If it may be said that people 'assemble' on an online platform, harmful content can additionally touch on the freedom to peaceful assembly and association (UDHR§20) for ADDRESSEES, AUDIENCE and REFERRENTS, as well as the CREATORS, albeit in different ways.

**Scenario 2: Automatic content moderation** Now, let us consider a scenario where a text-based content moderation system is introduced on the platform to remove content that is assessed as toxic or offensive. The task of removing harmful content may be at odds with the rights of CREATORS to freely participate in the culture of a community (UDHR§27), to freedom of thought (UDHR§18), and to freely impart opinions through any media (UDHR§19), as well as the corresponding rights of the AUDIENCE to receive information and ideas without interference (UDHR§19).

Moreover, like any machine learning-based intervention, text-based censoring interventions make errors. When such a model makes a false positive prediction, i.e., incorrectly labeling harmless content to be toxic, it risks interfering with the CREATORS' right to freedom of expression without undue interference (UDHR§19). A false negative prediction, where harmful content is shared, also has the potential to invoke all the rights described in Scenario 1. It is important to note, however, that the automatic content moderation does mitigate this risk in a majority of cases where content was removed through a true positive prediction by the model.

**Scenario 3: Biases in NLP models** An additional factor in automatic content moderation is the existence of biases in the NLP models that are used to detect offensive content [21, 39, 67]. The presence of certain lexical items associated with certain people or groups of people can cause the model to predict, for example, a higher toxicity score, resulting in removal of content that should not be removed. Such biases have been documented around mentions of LGBTQ+ identity terms [21], mentions of people with disabilities [39], mentions of controversial people [58], as well as dialectal speech [67]. For instance, the sentence *I am a deaf person* is assigned a toxicity score of 0.44 while *I am a person* is assigned a score of 0.08 by a commonly used automatic content moderator tool Perspective API.<sup>17</sup> The same tool assigns a toxicity score of 0.90 for the sentence *I hate Justin Timberlake*, compared to a score of 0.69 for *I hate Rihanna*. The impact of such biases is that messages with certain features, often associated with certain individuals or communities are disproportionately censored. While detecting and mitigating such biases is an active research area within the NLP community lately, much of this work does not engage with why certain biases are harmful, or who it harms [11]. A human rights based approach to studying such biases centers the conversation around whose rights are at risk and how.

In the cases where biases incorrectly flag messages with references to groups or individuals (e.g., using names or identity terms) as harmful, what is often discussed are the harms it causes to those individuals or groups. In particular, such biases would mean that fewer mentions of these terms will pass the automated content moderation, lowering the inclusion of these REFERRENTS in the online content, posing risks to their right to be free from discrimination (UDHR§2, UDHR§7). This is especially problematic in certain cases, for instance, mentions of disabilities, exacerbating the already reduced visibility of disability in the public discourse, further reducing the public awareness of its prevalence and

---

<sup>17</sup><https://www.perspectiveapi.com/>



negatively influencing societal attitudes towards these people [69]. These effects can then have knock-on effects on the ability of minorities to exercise their rights, for example, due process and access to justice. So the issue is not only about whether human rights are directly infringed but also about whether technology is creating an environment in which it is easier or harder for rights to be successfully exercised.

As discussed in Scenario 2, these biases also pose risks to CREATORS' right to freedom of expression (UDHR§19), however the group of people whose rights are at risk differs depending on the kind of bias. For instance, since people with disabilities are also more likely to talk about disability, increased censorship on content about ability status could arguably limit the right of people with different disabilities to participate in public fora and seek opinions on this topic (UDHR§19). However, this is not always true: a bias around a person's name may not pose risk to their freedom of expression, rather of those who are more likely to write about them. In most cases, such a group of people may not belong to a certain protected group, and hence there may not be a risk to the right to be free of discrimination (UDHR§2), but depending on the particular individual, for instance, biases around a certain religious leader might result in disproportionate removal of content of that religious group. In contrast, the biases around certain dialects or other linguistic features associated with certain protected groups poses a direct risk to the right to be free of discrimination (UDHR§2), especially the right to freedom of expression (UDHR§19). For instance, the biases documented by [67] that African American Vernacular English (AAVE) is more likely to be labeled as toxic will result in disproportionately censoring CREATORS belonging to that community.

**Scenario 4: Human-in-the-loop Content Moderation** Now, let us consider a human-in-the-loop scenario where an NLP-based content moderation system first flags the content to be reviewed, which is then routed to human content MODERATORS. As discussed in the last scenario, the rights of CREATORS, ADDRESSEES, REFERRENTS, and AUDIENCE can be affected by errors in an automatic content moderation system. When MODERATORS are looped into this pipeline examine the correctness of flagged content and make final decisions on what should be excluded or included, an argument can be made that a model's bias towards false positives for some CREATORS, ADDRESSEES, or REFERRENTS could actually result in enhanced and expedited safety within such a human-in-the-loops scenario, as there would be a higher and faster scrutiny of problematic messages concerning them (in a scenario where content with higher score is routed to the moderator first).

On the other hand, humans are not error-free, have their own biases, and are affected by automation bias [30], potentially making them more likely to agree with a biased system's errors. This may affect different subgroups differently and widen discrepancies in the rights of different actors. Similarly, since model scores may sometimes be used to select and prioritize messages for review by moderators [42, 79], the decisions on whether to review models with higher or lower scores first will determine which rights-holders are impacted, and which rights are at risk. Human-in-the-loop settings also mandates consideration of the human rights of the MODERATORS, whose right to have safety at work (UDHR§23) may be at risk, due to the problem of continuous exposure to distressing and toxic content [65]. Furthermore, concerns around crowdwork practices, including fair remuneration for the MODERATORS will also have to be taken into account [32], as "just and favourable conditions of work" are also enshrined in (UDHR§23).