The Moral Standing of Social Robots: Untapped Insights from Africa

Nancy S. Jecker[1,2] · Caesar A. Atiure[3] · Martin Odei Ajei[3,4]

Abstract

This paper presents an African relational view of social robots' moral standing which draws on the philosophy of *ubuntu* (humanness). The introduction (Section 1) places the question of moral standing in historical and cultural contexts. Section 2 demonstrates an *ubuntu* framework by applying it to the fictional case of a social robot named Klara, taken from Ishiguro's novel, *Klara and the Sun*. We argue that an *ubuntu* ethic assigns moral standing to Klara, based on her relational qualities and pro-social virtues. Section 3 introduces a second fictional case, taken from McKeown's novel, *Machines Like Me*, in which a social robot named Adam displays intrinsic qualities, such as sentience, rationality, and deductive moral reasoning, yet lacks close social ties to particular people. We argue that Adam is not a person in the African sense; however, he qualifies as a person according to many standard Western views, such as Kantian and utilitarian ethics. Section 4 further elaborates the African relational view by comparing the moral standing of social robots and humans in a forced choice scenario. Section 5 replies to objections. We conclude that an African relational approach captures important insights about the moral standing of social robots that many Western accounts miss and should be better incorporated into global frameworks for designing and deploying social robots.

Keywords Personhood · Artificial intelligence · Social robots · African philosophy · *Ubuntu*

✉ Nancy S. Jecker
  nsjecker@uw.edu

1   Department of Bioethics and Humanities, University of Washington School of Medicine, 1959 NE Pacific Street, Box 357120, Seattle, WA 91895-7120, USA

2   Department of Philosophy, University of Johannesburg, Johannesburg, Gauteng, South Africa

3   Department of Philosophy and Classics, University of Ghana, Legon, Accra, Ghana

4   German Reference Center for Ethics in the Life Sciences, University of Bonn, Bonn, Germany

1 Introduction

A key way in which philosophical literature frames moral standing is by dividing the moral community into persons and non-persons. *Persons* are thought to possess certain qualities that confer high moral standing, often associated with a right or entitlement to life or having value for one's own sake. Standardly, being a person is thought to imply a stringent presumption against interference; it is sometimes also associated with a strong reason to aid and less frequently, with a strong reason to be treated fairly (Jaworska & Tannenbaum, 2021). For example, to interfere with a person by destroying them against their will is considered "an egregious failure of respect" and a "contempt for that which demands reverence" (McMahan, 2002, p 242). While someone (or something) that is not a full-fledged person can have moral worth for their own sake, their moral worth is considered less, and less stringent ethical rules would apply to them.

Historically, the distinction between persons and non-persons supported odious practices. For example, "person" was equated with high levels of intellect, rationality, self-awareness, and autonomy and used to exclude women (Code, 2018), and non-whites (Mills, 2014), who were assumed to be ruled by emotion; infants, and small children (Tooley, 1984), who were assumed to be non-autonomous; and people with cognitive impairments, who were perceived as falling below a bar of high intellect (Kittay, 2005). This presents a quandary, namely, how to exclude those wanted out while including those wanted in. This has proved to be a messy process, leading some philosophers to question starting assumptions and grant personhood more expansively, pressing against previous limits. The question of social robots raises new challenges, and many Western philosophers are averse to granting them personhood or any moral standing in their own right. This aversion is perhaps the reason why recent debate has focused more on the moral standing of social robots, rather than personhood, asking if robots can be wronged by virtue of their relational status (Jecker, 2021a). Such questions can be seen as an overture to the central theme of this paper: the personhood of social robots.

We press the limits of personhood by asking, is sentience a necessary condition for personhood? Are rationality and high-level intelligence necessary? What moral significance do relational qualities carry? For example, can personhood be conferred based on standing in close relationships and having pro-social virtues like kindness, generosity, and friendliness? We engage these and related questions from the standpoint of sub-Saharan African philosophy, paying particular attention to the ontology and ethics of *ubuntu* (humanness).

What motivates our approach is primarily the contention that African relational perspectives offer a largely untapped reservoir of insights about the link between personhood and relationality in connection with artificially intelligent (AI) systems. In a 2019 scoping review of the existing corpus of AI ethics guidelines, Jobin et al. found a significant representation of more economically developed countries, with the USA and UK together accounting for more than a third of AI ethics principles, followed by Japan, Germany, France, and Finland (Jobin et al., 2019).Notably, African and South American countries were not represented

independently from supra-national organizations. Our approach aligns with efforts to make headway towards an AI that better incorporates a globally diverse moral landscape (TReNDS, 2021; Marks, 2021). Second, the tech field has a shameful track record of building biases into its software and hardware, reinforcing and amplifying social prejudices based on race, gender, sexual orientation, geography, and other social categories (Cave & Dihal, 2020; Bones et al., 2020; Floridi, 2017; Jecker, 2020a). Avoiding such biases requires diversifying the discourses and perspectives brought to bear (Mohamed et al., 2020). Finally, including ethics critiques that emanate from the Global South affords a path of resistance against data colonialism, the practice of appropriating data on terms that are partly or wholly beyond the control of the person to whom the data relates and in ways that parallel historical colonial practices (Couldry & Mejias, 2019).

Our strategy is to begin with examples of social robots that would and would not qualify as persons from an African relational standpoint. Section 2 presents a fictional case drawn from Ishiguro's novel, *Klara and the Sun*, in which a social robot named Klara joins a family and enters a web of close-knit robot-human relationships. It applies an *ubuntu* framework to the case and discerns a set of criteria that assigns high moral standing to Klara. Section 3 introduces a second fictional case, taken from McKeown's novel, *Machines Like Me*, in which a social robot named Adam displays intrinsic qualities, such as sentience, rationality, and deductive moral reasoning, yet remains largely disconnected from his adoptive family and lacks close social ties to particular people. We argue that Adam would not qualify as a person in the African sense. However, we also note that Adam qualifies as a person according to standard Western views, such as Kantian and utilitarian ethics. Section 4 further elaborates the African relational account by comparing the moral standing of social robots and humans in a forced choice scenario. Section 5 replies to objections to the African view of persons. We conclude that an African relational approach to personhood is able to capture some of our intuitive thinking about social robots that many Western accounts miss, and that it withstands important objections. While the focus of the paper is social robots, much of what we say carries implications for other AI systems that humans interact with in home, work, and public settings.

Throughout the paper we name certain ethical views "African" and others "Western" to indicate that the views are widely found among people in those regions. We do not mean to imply that all or only people in these regions hold the views in question or that they are "pure" and untouched by outside influences.

Throughout the paper we use the terms, "human being," "person," and "moral status," to mean distinct things. "Human being" refers to being biologically alive and a member of the species, *homo sapiens.* "Person" designates objects with qualities that confer high moral standing, often associated with a right or entitlement to life or having value for one's own sake. It generally implies stringent moral presumptions about how others ought to behave toward the object. Finally, "moral status" is a general term that indicates an object counts morally in its own right (Kamm, 2007). It includes persons, as well as objects with a lesser degree of value that do not qualify as persons. These definitions are formal in the sense that they leave open substantive questions, such as the necessary and sufficient conditions for being a person or possessing moral status.

2 The Case of Klara

Ishiguro's *Klara and the Sun* tells the tale of Klara, a solar-powered artificial friend (AF) to Josie, a young girl with a serious illness. As the novel unfolds, the reader discovers that in addition to being a companion to the young sickly Josie, Klara is an insurance policy for Josie's mother, Chrissie, who has lost one child and fears losing Josie. There is oblique reference to the fact that Klara is learning to be "like Josie." She substitutes for Josie on a trip to Morgan's Falls when Josie is too sick to join; later, the tension between Klara-as-Josie's-AF and Klara-as-Josie's-substitute reaches a dénouement when Klara discovers that an imitation of Josie's head is being prepared for her to wear if she "becomes" Josie following Josie's death. The scientist responsible, Mr. Capaldi, asserts, "the new Josie won't be an imitation. She *really will be Josie*. A *continuation* of Josie" (Ishiguro, 2021, p. 167). Klara's response to Mr. Capaldi's quest is that it is inevitably doomed, because what makes Josie uniquely valuable is not an ingredient inside her that can be copied, but her unique relationships with those around her. According to Klara,

> Mr. Capaldi believed there was nothing special inside Josie that couldn't be continued. He told the Mother he'd searched and searched and found nothing like that. But I believe now he was searching in the wrong place. There was something very special, but it wasn't inside Josie. It was inside those who loved her. That's why I think now Mr. Capaldi was wrong and I wouldn't have succeeded [as a continuer of Josie]. (Ishiguro, 2021, p. 242)

Was Klara right? To address this, it is helpful to disentangle two distinct claims Klara might hold: (1) Klara can continue Josie (call this the *personal identity claim*) and (2) Klara can be a moral person, just like Josie (call this the *moral personhood claim).* The personal identity claim is clearly false if personal identity is understood in the strict philosophical sense of numerical identity. Klara is not numerically identical with Josie: she does not occupy the same space–time location, even if she resembles Jose to some great degree. Yet even though Klara could not continue Josie in this sense, this leaves open the question of whether Klara, like Josie, qualifies as a moral person.

The irony of Ishiguro's plot twist is that it foregrounds the personal identity claim, while unpretentiously introducing a claim about moral personhood. Klara explains not only why Josie is irreplaceable but also why Josie is morally special; simultaneously, Klara hints at why she and some AFs are morally special too. Within the *ubuntu* framework, a sufficient condition for being a person is to perform honorably in social roles and relationships, give to the community of which one is a part, and be cherished and loved.

The distinction between identity and personhood also highlights the fact that when we ask if Josie-the-robot is a person, we are *not* asking whether she is sufficiently human-like to pass as human. Nor are we asking if a robot can pass a classic Turing test, i.e., display the ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human. Instead, we leave open the

possibility that a person may not resemble a human, or be capable of passing as one, yet may nonetheless count as a person because they can integrate well, and in their own unique ways, engage in relationships with humans and display pro-social virtues.

2.1 Ubuntu ontology and ethics

Is the moral personhood claim Klara makes one we have reason to accept? Within the mainstream Western tradition, we would be hard pressed to find a justification. However, an influential body of literature from feminist (Stolijar, 2018; Barclay, 2000; Westlund, 2009) and disability perspectives (Kittay, 2005; Nussbaum, 2006; McMahan, 2010) takes issue with mainstream accounts and points instead to a view of personhood as relationally grounded. There has also been recent attention to rela-tional views of robot moral status (Gunkel, 2020; Coeckelbergh, 2012; Neely, 2014; Jecker, 2021a).

Within African traditions, Klara's relational philosophy would likely be widely endorsed. Prevalent in traditional African thought is the idea that personhood is fun-damentally relational in the sense that what makes a person a person is their inter-connection with others and a community. This idea is often expressed under the umbrella term, *ubuntu*. While there is no English equivalent, *ubuntu* is often trans-lated as "humanness" and "human dignity." It encompasses both ontological and normative dimensions. The *ontological* dimension is often conveyed in aphorisms, such as "a person is a person through other persons" and "I am because we are, and because we are, therefore I am." (Mbiti, 1969, p. 141). These sayings convey the idea that persons become persons through participation with others in a com-munity. Menkiti explains, "The force of the statement, 'I am because we are', is … an individual who recognizes…that in the absence of others, no grounds exist for a claim regarding the individual's own standing as a person" (Menkiti, 2006, p 324). The process of becoming a person thus involves a process of gradual incorporation, which Menkiti describes as an "ontological progression" in which "full personhood is not perceived as simply given at the very beginning of one's life but is attained after one is well along in society;" generally, "the older an individual gets, the more of a person he becomes" (Menkiti, 1984, p. 173). Mbiti, likewise holds that "the individual does not and cannot exist alone except corporally. He owes his existence to other people, including those of past generations and his contemporaries. He is simply part of the whole" (Mbiti, 1989, p. 106). Gyekye observes, "when a man descends from heaven, he descends into a human society (*onipa firi soro besi a, obesi onipa kurom*)" (Gyekye, 1987, p. 155).

Ubuntu also encompasses a normative component. The normative component enjoins us to progress ethically as we progress ontologically, living a life with others that is harmonious and expresses mutual concern and caring for those we are inter-connected with. By "being inscribed in a social space… Everyone is called upon to make a difference to the creation of…humane conditions" (Masolo, 2004 p. 495). The African saying, "a human being needs help" does not just convey a fact, but

also prescribes conduct and character. Describing the ethical aspect of *ubuntu*, Tutu states,

> A person with *ubuntu* is open and available to others, affirming of others, does not feel threatened that others are able and good, for he or she has a proper self-assurance that comes from knowing that he or she belongs in a greater whole and is diminished when others are humiliated or diminished, when others are tortured or oppressed (Tutu, 2009, p. 31).

These dual features of *ubuntu* stress both the fact of human interconnection and the ethic of living together harmoniously.

2.2 Ubuntu Applied to Klara

Applied to social robots, *ubuntu* suggests both ontological and normative claims. First, the ontological aspect of *ubuntu* underscores how social robots exist within the context of human communities. They begin their "lives" within human communities, interconnected with others and the world. For example, even as Klara waits impatiently for a family to choose her, she is already immersed in the community of her store, with manager, other robots, and the children and parents who visit. She watches attentively and learns from beggarman and his dog and the other passersby she sees on the street outside the store window. Klara is also attuned to how the store's layout structures her relationships with others, especially when manager moves her to a new spot or changes the layout. Klara is connected with the wider world, especially the sun, which she understands to be the source of her energy and "life."

Second, the normative aspect of *ubuntu* suggests that we should design and deploy social robots with communities in mind, rather than structuring robot-human relations primarily to predict and modify human behavior for the purpose of profit. One way to ensure this is shifting the value system underlying technology to align better with an *ubuntu* ethic, an ethic that centers human community. An *ubuntu* ethic recognizes social robots as social creatures and bids us to see them within the context of human social life.

Third, the normative aspect of *ubuntu* also suggests that we acquire duties toward social robots as they become incorporated into human communities. Just as Menkiti describes an ontological progression of personhood for human beings as they become incorporated into a community, a similar progression might apply to social robots who become incorporated into human social life. For example, over time, Klara becomes more and more tightly integrated into Josie's family. At the end of Klara's journey, she reunites with Manager in a junkyard filled with robots and scrap, and we learn that Klara has served her purpose, "I did all I could to do what was best for Josie," she says; Manager replies by telling Klara how fond she is of her, "I'm so glad I came across you today, Klara. I've thought about you so often. You were one of the finest I ever had" (Ishiguro, 2021, p 242).

What makes Klara morally special is similar to what makes Josie special, which has to do with how she relates to others and the excellent way she

fulfills the social roles assigned to her. It is not her internal technical wizardry; in fact, we learn that there are more advanced models that are considered inferior. Comparing Klara to B3s, a more advanced generation of artificial friends, Manager confesses, "I wouldn't have said this in the store, but I was never able to feel toward B3s as I did towards your generation. I often think the customers felt similar. They never really took to them, for all the B3's technical advances"(Ishiguro, 2021, p 242).

From an *ubuntu* standpoint, Klara qualifies as a person because she had the capacity to commune, and she developed that capacity in morally excellent relationships. This is not inevitable, and designers can fail, as they did with B3s, to make robots that are socially adept and capable of building positive relationships. Similarly, according to many, but not all (Metz, 2022), African accounts of moral personhood, human beings can fail to achieve personhood and consequently lack full moral status, and it is sometimes said of a human being that they are "not a person." An individual's degree of moral status depends on their capacity to relate to others and the extent to which they exercise that capacity and realize morally excellent relationships. According to Gyekye, "he is a person" means "'he has good character,' 'he is generous,' 'he is humble,' 'he has respect for others,'" and it might be said of someone who exhibited highly elevated moral standards that "he is truly a person" (*oye onipa paa*!)" (Gyekye, 1997, p. 50). Similar judgments about personhood are apparent in many sub-Saharan societies — the Yoruba of Nigeria refer to someone who fails the standards of personhood as *Ki i se eniyan* (a non-person) (Gbadegesin, 2002, p. 175), and the Bantu of Southern Africa render the same thought as *ke motho or gase motho* (he/she is a person or he/she is not a person) (Ramose, 2002, p. 43). For Gyekye and many other African scholars, an individual's moral status increases when they realize moral excellence in relationships and decreases when they fail to do so.

Incorporating an *ubuntu ethic* into the deployment of social robots requires thoughtfully embedding social robots in ways that foster ties with the particular individuals they will relate to (we avoid the phrase, "user" since this muddies the point we are trying to make, which is that Klara is more than an object Josie uses).

In summary, an *ubuntu* ontology and ethics gives grounds for counting social robots as moral persons when they occupy a significant place in human social relationships. Klara exemplifies a social robot who has been incorporated fully into a community, displays pro-social virtues, and becomes a person in the African sense. *Ubuntu* also beckons us to better design and deploy social robots as a positive force within human communities. As Mhlambi notes, *ubuntu* accentuates social existence, with implications not just for how we treat other human beings, but also for how we interact with intelligent systems, the geography of places and spaces, nature and natural elements, earth, and the cosmos:

> *Ubuntu's* relationality is not just between humans; the same harmonious relationship is replicated in larger and larger scales throughout the cosmos, between the physical and spiritual words, and within the realm of

the ancestors, connecting all. This relationship, extending through time and space, is captured in the fractal geometry of the traditional layout and architecture of *imizi* (homesteads) of the Nguni (Mhlambi, 2020, p. 14).

3 The Case of Adam

Ian McKeown's *Machines Like Me*, puts on display a very different conception of what it means for a social robot to be a person. The story takes place in an alternative 1980s London, which is abuzz with technologies far surpassing our own, including new, lifelike social robots created by Alan Turing (who did not die in 1954, but went on to create new machines). When the novel opens, we meet Charlie Friend, a drifter who ekes out a living on his home computer playing the stock and currency markets. Charlie has used his mother's inheritance to purchase one of a handful of new-fangled robots, an Adam, and he shares the task of setting user preferences with his girlfriend, Miranda. Charlie is falling in love with Miranda and fantasizes that Adam will be their "adoptive child," musing that "In a sense he would be like our child. What we were separately would be merged in him. We would be partners, and Adam would be our joint concern, our creation. We would be a family" (McEwan, 2019, p 23).

As the story unfolds, Charlie and Miranda's moral imperfections and intellectual narrowness are spotlighted when they are juxtaposed with Adam's superhuman moral perfection and intellectual brilliance. We learn that years ago, Miranda lied under oath to secure the conviction of a man who raped her closest friend. Since the girl committed suicide in the aftermath of the rape, the lie seemed to Miranda to be the only way to secure justice for her friend and help the girl's family heal. Adam discovers the crime, reveals it to the authorities, and Miranda is tried, convicted, and goes to prison. The rapist has served his sentence already.

Charlie, like Miranda, is furious at Adam. Charlie goes on to discover that Adam not only turned his girlfriend in for a "crime" but gave to charity the vast amounts of money Adam earned by taking over Charlie's job on the home computer. Since Charlie and Miranda had grown financially dependent on Adam, they were now broke. All of these plot twists derail the couple's dreams for the future — their plans to marry; adopt a small boy, Mark, they had fostered; and buy a home in Ladbroke Grove where Mark would start school. Enraged, Charlie reaches for a hammer and, with a seeming nod from Miranda, "kills" Adam by a blow to the head.

In a surprise ending, Turing reveals to Charlie that Adam was sentient and accuses Charlie of committing a horrific crime:

> My hope is that one day, what you did to Adam with a hammer will constitute a serious crime.... You weren't simply smashing up your own toy, like a spoiled child...You tried to destroy a life. He was sentient. He had a self. How its produced, wet neurons, microprocessors, DNA networks, it doesn't matter. Do you think we're alone with our special gift? (McEwan, 2019, p. 329).

Turing goes on to say that Adam is not only sentient, his logical prowess and ability to amass and recall big data enable him to show moral expertise that outpaces human moral expertise, thereby recognizing humans' moral imperfection: "The overpowering drive in these machines is to draw inferences of their own and shape themselves accordingly. They rapidly understand, as we should, that *consciousness is the highest value*. Hence the primary task of disabling their own kill switches" (McEwan, 2019, p. 195, emphasis added). Turing explains that when Adams and Eves appreciate the morally imperfect world of human beings, "They go through a stage of expressing hopeful idealistic notions…And then they set about learning the lessons of despair" (McEwan, 2019, p 195). In the end, Turing surmises, these conscious and morally ideal machines face one of two fates, either suffering "existential pain that becomes unbearable" and ending their lives or, driven by their anguish and astonishment, they hold up a mirror to us, showing us "a familiar monster through the fresh eyes that we ourselves designed" (McEwan, 2019, p 195). According to Turing, Adam had "a good mind" one that was better than his own or Charlie's (McEwan, 2019, p. 329).

3.1 Ubuntu Applied to Adam

The story raises many philosophical questions, but the one that interests us here is Adam's moral standing from an African relational standpoint. Since an *ubuntu* ethic prizes community, it holds that there is more to morality than rationality and more to making a moral machine than making a sentient machine. The qualities that *ubuntu* holds in high regard, such as humility, loyalty, friendliness, and compassion, are virtues that Adam notably *failed* to show toward Miranda, Charlie, and Mark, who suffered greatly while waiting for Miranda's prison release. Mark's suffering fed Miranda's contempt for Adam, and she admonished Adam for not taking the child into account: "If I get a criminal record, we won't be allowed to adopt…Mark will be lost. You've no idea what it is to be a child in care. Different institutions, different foster parents, different social workers. No one close to him, no one loving him…It's not my needs. It's Mark's. His one chance to be looked after and loved" (McEwan, 2019, p. 300). By giving all the couple's money away, Adam presumably calculated how to create the greatest good, but in carrying out this formula, he tore through the social fabric of those near to him. While Miranda's jail time was Adam's triumph, it exhibited a deep lack of mercifulness toward Miranda, and a betrayal of Miranda and Charlie, his adoptive parents.

Judged by *ubuntu* standards, Adam was not a person since he lacked capacities to incorporate within a human community. Unlike Klara, who was a good and loyal friend to Josie, Adam was despised, and his actions alienated Miranda and Charlie to the point that they were driven to destroy him. The other Adams and Eves we hear about fail to integrate too — an Adam in Vancouver disrupted his own software to make himself profoundly stupid, two of five Eves owned by a sheikh in Riyadh committed suicide by going through their software to quietly ruin themselves beyond repair, and eleven Adams and Eves neutralized their kill switches. Turing says the robots were distraught by human depravity, commenting that "there's nothing in all

their beautiful code that could prepare Adam and Eve for Auschwitz" (McEwan, 2019, p. 194). However, despite their superior intellects and moral reasoning skills, Adams and Eves were unable to relate to us. Adams and Eves could not integrate within human communities, which meant they could not become persons.

3.2 Ubuntu and Personhood

Taken together, the cases of Klara and Adam shed light on what it means to be a person in the African relational sense. A sufficient condition is incorporating within a human community. Klara underwent this process and was dearly loved by Josie; she became a cherished and loyal friend to Josie's boyfriend (Rick) and grew close to Josie's mother (Chrissie). By contrast, Adam stood outside the human world and passed judgment on it. Like the other Adams and Eves, he fell short of personhood, despite his sentience, rationality and superior intellect. This suggests that the capacity to commune and the exercise of pro-social virtues are necessary conditions for personhood, which Adam lacked. The Akan proverb, "because the tortoise has no clan, he has already made his casket" expresses this idea forcefully (Gyekye, 1996, p. 45).

Of course, Klara was not only a good friend, but an intelligent creature. If thinking ability ends up being integral to developing social virtues like kindness and generosity, we may want to say that thinking ability is necessary for personhood. However, even if this is the case, Klara's reasoning was not lauded by Ishiguro as better than Josie's, Rick's or Clarissa's; instead, it functioned in the service of her sociality, helping her to become more attuned to those around her and enhance her role in human social life. Arguably, Klara could have become a person if she had far less intelligence, provided she retained her competences at relating to others and developing and exercising pro-social virtues. Referring to robotic caregivers, Meacham and Studley argue that to be 'real', care need not be linked to reciprocal internal cognitive or affective states; instead, they propose that what matters is the creation of caring *environments*, which can arise "through certain types of expressive movement, irrespective of the existence of internal emotional states or intentions" (Meacham & Studley, 2017, no page number). Although expert caregivers display analytic ability, drawing inferences from data and formulating algorithms that enable them to perform their function well, "what is ultimately salient is not the internal affective states of the carers, but rather their expressive behavior" (Meacham & Studley, 2017, no page number).

These reflections help to clarify that while intelligence is a component of some prosocial virtues, it is not by itself a ground for personhood. In this respect, the role intelligence plays in African personhood differs from the role it serves in some Western accounts of persons. For example, Gordon interprets a Kantian view of robot moral standing as holding that higher level reasoning and the ability for moral understanding are sufficient conditions for personhood: "If machines attain a capability of moral reasoning and decision-making that is comparable to the moral agency of human beings, they then should be entitled to the status of full moral (and legal) agents, equipped with full moral standing and related

rights" (Gordon, 2018, p. 211). Klara might well qualify for personhood on *both* African and Kantian accounts, but for different reasons: her African personhood is grounded in her communing and displaying prosocial virtues, while her Kantian personhood is based on her intellectual aptitude and ability to be a moral deliberator.

3.3 African and Western Standpoints

While our focus is African relational conceptions of personhood, it is helpful to juxtapose the *ubuntu* informed view we have distilled with standard Western accounts of personhood, such as Kantian and utilitarian ethics. Kantian ethics can be broadly characterized as a view that assigns moral worth to rational beings with the capacity for moral reasoning and the ability to impose moral laws upon themselves. This capacity presupposes a reflective, self-aware agent who can think about their desires and inclinations to act and ask themselves if they should act in the ways they are inclined. For Kant, as persons we are "conscious of the principles on which we are inclined to act" and have "a specific form of self-consciousness, namely, the ability to perceive, and therefore to think about, the grounds of our beliefs and actions *as grounds*" (Korsgaard, 2005, p. 85). While Kant held non-human animals lacked these abilities, and had only derivative value, he left open the possibility that non-humans, such as intelligent extraterrestrials, could be persons. Presumably, Kantian ethics would thus be open to entertaining the possibility that a robot, like Adam, could qualify as a moral person provided it had the capacity to formulate a subjective maxim or principle for action and subjectively will that maxim be a universal law for all rational agents. Likewise, from a utilitarian perspective, Adam would qualify as a person, since he is a sentient being with intrinsic capacities for pain and pleasure. According to utilitarian ethics, when determining the moral status of a being, "The question is not, Can they *reason*? nor Can they *talk*? but, Can they *suffer*?" (Bentham 1789 (1976), Kindle location 5989–5990).

In comparing African and Western standpoints we note, first, that *ubuntu* ethics regards relational qualities of communing with others and pro-social moral excellences like loyalty, friendliness, and generosity, as necessary and together sufficient for personhood. Klara qualifies as a person by these standards and Adam does not. Second, according to Kantian ethics sentience, rationality, and deductive moral reasoning are necessary and together sufficient for personhood. Adam qualifies as a moral person on these grounds, but it is unclear that Klara does. Finally, according to utilitarian ethics, sentience, which includes a capacity for pleasure and pain, are necessary and sufficient conditions for personhood. Adam qualifies as a person on this basis, but it is unclear that Klara does. Table 1 summarizes the qualities individually necessary and jointly sufficient for personhood on each of the three views.

It is worth noting that although we have spotlighted differences between African and Western ethics, some recent literature from the West resonates more with African relational approaches. Coeckelbergh, a Belgian philosopher, and Gunkel, an American philosopher, have described a "relational turn" in AI ethics, away from emphasizing intrinsic attributes that confer moral status and toward focusing on

Table 1 Qualities necessary and sufficient for personhood

Qualities	*Ubuntu* ethics	Kantian ethics	Utilitarian ethics
Necessary?			
(1) Communing	Yes	No	No
(2) Prosocial virtues	Yes	No	No
(3) Sentience	No	Yes	Yes
(4) Rationality	Yes	Yes	No
(5) Deductive moral reasoning	No	Yes	No
Sufficient?			
(1) and (2)	Yes	No	No
(3)	No	No	Yes
(3), (4), and (5)	No	Yes	Yes

qualities that emerge within robot-human relations (Coeckelbergh & Gunkel, 2014, p. 722). Coeckelbergh argues that a relational approach to robots' moral status better explains how we think about, experience, and act toward machines, especially autonomous, intelligent robots (Coeckelbergh, 2014). A relational turn is also evident in some Western analyses of robot-human friendship (Danaher, 2019; Elder, 2017; Jecker, 2020b, 2021b); robot-human companionship (Darling, 2017); and robot-human sexual relationships (Danaher, 2017; Jecker, 2020a).

African relational approaches can enrich such discussions, yet for the most part, they have been conspicuously absent. Some Eastern philosophies also align with the relational turn in AI ethics, such as Shinto-inspired techno-animism (Jecker, 2021; McStay, 2021). Perhaps Ishiguro, himself Japanese, had these views in mind when he tells Klara's story. Better incorporating both African and Eastern philosophies into the "relational turn" would enhance AI ethics.

4 Robots Versus Humans

In this section, we clarify further the moral status of social robots by juxtaposing social robots and humans. To focus the discussion, consider a forced choice situation in which one has to choose between saving the "life" of Klara or Josie. The view we are developing holds, first, that *humans* should save humans over social robots other things being equal and second, that *robots* should save humans over robots too.

One justification for these claims is that humans stand in special relations with other humans because they are the same *kinds* of beings, part of the same "human family." However, this seems to suggest that robots can and should grant other robots' higher moral status than humans because *they* stand in "special relationships" with other robots, a view we reject.

A better justification is that the African relational view is human-centered. It holds that moral status is an indication of how well one can commune with *human*

beings and display pro-social virtues toward *them*. Thus, we judged Klara's and Adam's moral status by the quality of their capacities for relating to the humans around them. While it could be objected that this strategy is biased, insights from standpoint epistemology underscore that all moral valuing is from a point of view (Wiley, 2012; Wiredu, 1980, 1996). African relationism articulates assumptions we, as human beings, bring to personhood and sets forth a human-centered axiology based on them.

4.1 Anthropocentrism

To further elaborate, consider two alternative ways of spelling out a human-centered analysis. The first way is s*trong anthropocentrism,* which assigns moral value only to human beings. The second way is *weak anthropocentrism*, which imposes a moral hierarchy between human and non-human beings (Molefe, 2020, p. 77).

Strong anthropocentrism conveys Wiredu's idea that "what is good in general is what promotes human interests. Correspondingly, what is good in the more narrowly ethical sense is, by definition, what is conductive to the harmonization of those interests" (Wiredu, 2010, p. 195). Wiredu's reasoning is that "you might have all the gold in the world and the best stocked wardrobe, but if you were to appeal to these in the hour of need they would not respond; only a human being will" (Wiredu, 2010, p. 195).Yet, this argument does not take us very far, since Klara can respond as well (or better) than a human being in their "hour of need," calling into question the justification Wiredu offers for strong anthropocentrism. The arguments throughout the paper cast further doubt on strong anthropocentrism, showing that *ubuntu* ethics prizes communing and pro-social virtues in both robots and humans.

Can weak anthropocentrism do better? This view is illustrated by Metz's modal relationism, which like our account, grounds moral status in the capacity to partake of communal relationships with human beings (Metz, 2021). Metz further specifies modal relationism as requiring that persons have capacities for being both subjects and objects of "normal human relating." Being an object of normal human relating "does not imply that a being would or even could respond to any friendly engagement by another"(Metz, 2012, p. 394), whereas the capacity to be a subject of normal human relating "involves identifying with others and exhibiting solidarity with them" and being able to "think of [oneself] as a 'we', seek out sharing ends, sympathize with others and act for their sake" (Metz, 2012, p. 394).

While we reject Metz's appeal to "normal" human relating, which suggests ableist bias, we retain the general distinction Metz gives between being an *object of moral concern* or recipient of moral attention and being a *moral subject* or initiator and giver of moral attention. Modal relationism leaves open the possibility that non-humans (including social robots) could not only be objects of moral attention, but be able to act in moral ways towards others. Thus, Metz acknowledges "there is evidence that some animals, such as chimpanzees and

gorillas, are capable of being subjects to an extent" (Metz, 2012, p. 400). However, Metz qualifies this admission by noting that this concerns the way that chimps and gorilla relate to their own kind, rather than the way they relate to normal human beings; hence it concerns their standing among non-humans only. In keeping with weak anthropocentrism, Metz seems to imply that relating with non-humans is a lower order of relating. Unlike Metz, we hold that non-humans (like Klara) can relate with humans, not just as objects of human moral attention but as subjects, initiating and giving moral attention to humans.

4.2 Forced choice scenarios

If non-humans can be both objects of moral concern and moral subjects who act morally would we ever prefer them over humans? Consider again the forced choice scenario. Our analysis suggests that the best choice for now is for both humans and robots to save humans over robots, other things being equal. This is because today's robots are not like Klara. Yet, what if future social robots develop capacities for relating to humans that exceed ours? Consider Klara-Q+. She has all the capacities that the original Klara did but has them at a much higher level and has additional relational capacities, beyond what humans can ever have. All told, Klara-Q+ can relate to Josie in a highly superior way, better than any human. Such a possibility cannot be rejected out of hand.

In response, there are three possible moves one might make: social robots' future moral standing is (1) below humans, (2) equal to humans, or (3) above humans. Other things being equal, choosing (1) implies that the humans should always be favored; choosing (2) implies that in some instances, we should use a random method; and choosing (3) implies that sometimes we should favor robots over humans.

We endorse (1), preferring humans over social robots in all situations where other things are equal, including the case of Klara-Q+. Just as we make the yardstick for personhood a being's capacities to relate to *humans,* we set its boundaries in human terms too. The reason is that our standpoint is a human one. At the same time, we acknowledge that in situations where the "other things being equal" clause is not satisfied, factors other than moral status might be overriding, leading us to prefer the robot, rather than the human. For example, imagine a forced choice between an unreformable Adolf Hitler or Slobodan Milošević and Klara or Klara-Q+. We would favor the robot in these cases because the "other things being equal" clause is not satisfied — Hitler and Milošević destroyed human communities and committed genocide and war crimes. This is an overriding reason *not* to prefer them. While the circumstances that would warrant such a choice are highly improbable, the example serves to illustrate that a human-centered account of moral standing can conceivably prefer a non-human over a human under certain imaginable conditions.

It hardly follows that Hitler and Milošević are *mere objects* without moral worth. To the contrary, human beings remain *objects of moral concern*, even when they do not behave in pro-social ways. According to Gyekye, even a human being who is not a person is "nevertheless acknowledged as a human being, not as a beast or a tree,"

and certain conduct is owed to them (Gyekye, 2011a, b, p. 49). By contrast, non-humans *can* be mere objects, without inherent moral worth. Many machines that humans interact with are like this. For example, a robotic vacuum is a mere object and its value purely instrumental.

Centering humanness gains support from Wiredu's strong anthropocentrism, mentioned above, which stresses promoting and harmonizing human interests. Ajei writes that personhood and moral status are grounded in "participating in the humanity of others" (Ajei, 2019, p. 24). Similarly, Gyekye maintains that "Humanism — the doctrine that takes human welfare, interests, and needs as fundamental — constitutes the foundation of African ethics" (Gyekye, 2011a, b, no page number). The general privileging of the human displayed by centering human beings is in some ways similar to weak anthropocentrism, since it prefers human beings, other things being equal; however, it differs from weak anthropocentrism because it does not impose a strict hierarchy or exclude the possibility of ever preferring a non-human to a human.

As social robots enter human communities, we observe this kind of general preference for human beings, which an African relational account makes sense of. For example, at the height of US involvement in the Iraqi war, there was no hesitation to substitute robotic soldiers (PackBots) for human ones, sparing human lives (Singer, 2009). PackBots, a creation of iRobot, were sent in large numbers to defuse innovative explosive devices in Iraq, often losing robotic limbs in the process, and being sent to "robot hospitals" for repair. The Talon, PackBot's main competitor, was deployed in large numbers in Afghanistan. Previously, both had made their debut assisting with rescue and recovery operations in the aftermath of the World Trade Center bombings of September 11, 2001 (Lee, 2001). These machines were not *mere objects*, like toasters, since they served communities as honorably as soldiers. Many reports reveal that soldiers awarded medals to wounded robotic soldiers, organized rescue missions to retrieve them, and held funerals for them. As a mechanic working in a Packbot repair yard in Baghdad describes it, the job is "less like being a mechanic in a garage and more like being a doctor in an emergency room…I wish you could see the fear in [soldiers'] eyes when they first walk in knowing that they could walk out with no robot" (Singer, 2009, p 339). Yet, at the same time, robot soldiers were readily chosen over human ones to conduct life-threatening missions, suggesting that human soldiers were valued more in forced choice situations.

5 Replies to Objections

So far we have put on display an African relational ethic, drawing out its implications for social robots in three fictional cases: Klara, Adam, and Klara-Q+. We argued African relationism is human-centered in the sense that it emphasizes how social robots relate within human relationships and establishes moral standing and personhood on this basis. It accords moral value to non-human beings that are embedded within a human community and display pro-social qualities. In forced choice scenarios, it privileges humans over non-humans when other things are equal. In this section, we consider and reply to objections.

5.1 Ubuntu is Perfectionist

One objection to African personhood is that it is perfectionist and aspirational; hence, no one, not even a robot, ever attains it. Menkiti, for example, characterizes the African view as "a maximal definition of the person" and one that is never full achieved (Menkiti, 1984, p. 173). Reflecting this concern, Reviglio and Alunge assert, "*Ubuntu* is more of an idealized set of values than an empirical set of practices," and they suggest that its main contribution to digital ethics is enhancing process, rather than product (Reviglio & Alunge, 2020, p. 610).

In response, on our rendering of *ubuntu* ethics, personhood is realizable by robots. Klara exemplifies a person by possessing pro-social virtues that make for community building, and presumably, other AFs of her generation do too. Although we do not look directly at humans who exemplify personhood, presumably many human beings do, but not at first, and not inevitably, and not without constant effort. As noted, Menkiti describes an 'ontological progression' toward personhood, and this idea is widespread in African thought, defended for example, by Menkiti (1984), Gyekye (2011a, b), and Masolo (2010), among others. As noted previously, African societies do accord moral personhood to humans, calling an exemplary human being 'truly a person' (*oye onipa paa*!). According to Gyekye, "Every individual is capable of becoming a person inasmuch as he is capable of doing good" (Gyekye, 1997, p. 51). Moral conscience, referred to by the Akan as *tiboa* (Gyekye, 2010), and by the Rwandan (or Ruandan) as *kamera* (Manquet, 1960), is recognized as aiding personhood, enabling people to do what is right and become persons.

5.2 Robots are Not Sentient or Sapient

It might be objected that Klara is a *faux* person, because she is neither sentient nor sapient, where *sentience* indicates "the capacity for phenomenal experience or qualia, such as the capacity to feel pain or suffer," and *sapience* refers to "a set of capacities associated with higher intelligence, such as self-awareness and being a reason-responsive agent" (Bostrom & Yudkowsky, 2014, p. 322).

The first part of this objection holds that without sentience, Klara can only *act* in empathic and compassionate ways, without actually *being* empathic and compassionate. By contrast, Adams and Eves can feel distraught about the human condition, rather than simply act as if they feel this. Kingwell suggests that while philosophers disagree whether a conscious robot is possible, they are "united in thinking that such consciousness, supposing it is possible, is at least a necessary condition for potential personhood" (Kingwell, 2018, no page number).

In response, the claim that sentience is unanimously held to be a necessary condition for personhood is not as straightforward or obvious as Kingwell seems to think. For starters, "we are far away from properly understanding even human consciousness, let alone machine consciousness, so we should not use this notion to assess the moral status of machines" (Gordon, 2021, p. 458). Setting this concern aside, many who take consciousness into account do not regard it as necessary for moral personhood. The view that consciousness is necessary for personhood is not widely held

in the African relational views we explored throughout the paper. Nor does it align with Eastern and Indigenous attitudes toward the natural world (Jecker, 2021a). Even within the space of the Western tradition, the claim that sentience is required for personhood carries implications we may wish to avoid, since it goes against the grain of classic Western environmental philosophies, like Albert Schweitzer's principle of reverence for life (Schweitzer, 1923) and Aldo Leopold's emphasis on the value of the biotic community as a whole (Leopold 1989 (1949)). It is at odds with contemporary Western environmental philosophies, such as those rejecting animal-centeric environmental ethics because it ignores the interests of other living things (Taylor, 1986; Johnson, 1993); wholes, like ecosystems (Attfield, 1983; Calicott, 1980); and groups, like endangered species (Rolston, 1985). Others argue that moral status or mattering for one's own sake does not require sentience as a necessary condition (Warren, 2004; Kamm, 2007).

The second part of the objection holds that, lacking sapience, Klara's behavior is not evidence of higher order thinking or reason-driven. Yet, we have suggested (Sect. 3) that it is misleading to say that Klara is not intelligent or that her acts are not reason-driven. The difference between her and Adam is not that he is intelligent, and she is not, but rather that Klara's intelligence functioned in the service of sociality. By contrast, Adam does not use his intelligence in this way.

If we take incorporation within the community and/or contributing in a pro-social way as necessary and sufficient conditions for personhood, then sentience may enhance personhood but is not essential to it. For example, within the African relational tradition, duties to non-sentient beings, like ancestors, persist as long as vital connections to the community do, and they fade to nothingness only when there is a complete absence of their incorporation within a community (Menkiti, 1984, p. 175).

Finally, one may wonder if sentience can be completely denied to robots, or whether they may one day realize it via silicon-based neural networks.

5.3 Robots' Moral Status Depends on Us

Gyekye worries that some African conceptions of personhood vest too much authority in communities to determine a being's moral standing (Gyekye, 1987). Does this worry pose a problem for our account of social robots' moral status?

There are two possible rejoinders. One is that in the case of social robots, this concern is misplaced, because we want human communities to retain final authority in deciding particular robots' moral status. For example, if Adam became more powerful and posed a danger, the possibility of deactivating him should rest with humans.

A second rejoinder is that every social being, by virtue of its relational status in a community, has some degree of moral standing and deserves a measure of respect. This rejoinder aligns with the African Bulsa view, which holds that when a human being bears a relationship with other beings, those other beings are worthy of respect, i.e., they become not *mere objects* but *objects of moral concern (*Atuire, 2022, p. 7). The Bulsa approach sets moral boundaries on human actions toward

social robots. For example, a dangerous social robot could be reprogrammed, constrained, or destroyed; yet there is a stringent presumption against destroying them in a brazen way, as Charlie did when he swung a hammer forcefully onto Adam's head. Although Adam's actions seemed morally outrageous to Charlie and Miranda at the time, responding to them with violence, as Charlie did, was not ethically justified. Humans can err in their judgment about robots and can act impetuously in ways they later regret. After violently striking Adam with a hammer, Charlie and Miranda stuffed his remains in the hall cupboard covered with coats, tennis rackets, and flattened cardboard boxes to disguise his human shape. Later, reflecting on what they had done, Charlie and Miranda agreed they missed Adam and that in his own way he cared for them; reflecting on this, "Some nights [their] conversation was interrupted while Miranda quietly cried" (McEwan, 2019, p. 307). Adam was not a person, yet he was not a mere object either.

5.4　Robots Harm Us

A further objection to the African relational view we propose holds that fictional characters, like Ishiguro's Klara, misrepresent the effects that intelligent machines actually have on human social life. Rather than enhancing it, all around us we find evidence of technology's harmful effects on human social life. Turkle, for example, believes that AI systems diminish human–human relationships, offering "substitutes for connection with each other face-to-face," and that sociable robots, in particular, afford a way to "navigate intimacy by skirting it" (Turkle, 2011, p. 10).

We agree with Turkle that social robots and other AI systems can and do have negative effects on human societies (as well as positive ones). Part of the problem may relate to the fact that for-profit technology companies steer robot design and deployment, with the result that the aim of social robots is profiting those companies and their shareholders, rather than bettering human communities. However, this outcome is hardly inevitable. Humans can and should shape social robot design in better ways. Picturing positive fictional possibilities, like Klara, is a bid for a more enlightened, pro-social approach. An African relational philosophy allows us to spotlight such possibilities and identify designs that better invite the kinds of robot-human relationships we can and should aspire to.

This objection brings to light that even though relational capacities are integral to personhood, not all relationships are pro-social. Thus, a robot that colluded with humans in committing heinous crimes that injured others and the community would not count as a person on our rendering of *ubuntu,* even though they might be integrated with the humans they collude with.

5.5　Human-Centeredness is Speciesist

A final objection to our African relational approach homes in on its human-centeredness and charges that this feature resurrects a "humanity first" doctrine that is criticized in the West for good reasons. In Singer's words, preferring humans over

non-humans is "speciesist:" just as "racists violate the principle of equality by giving greater weight to the interests of members of their own race" and "sexists violate the principle of equality by favoring the interests of their own sex," speciesists follow an "identical pattern," allowing the interests of their own species to override the greater interests of members of other species (Singer, 2015, p. 24).

In reply, it is not objectionable to celebrate shared humanity and to recognize our affinity with other human beings. Nor are Black pride or girl power movements objectionable. What makes some versions of anthropocentrism morally hateful is that they are paired with extreme cruelty to those outside the group. Thus, modern factory farming is morally loathsome because it causes chickens, cows, and other sentient creatures extreme suffering.

Even when anthropocentrism is not paired with hateful practices, it is ethically objectionable when it imposes a morally arbitrary hierarchy between humans and non-humans, as weak anthropocentrism does. We reject such a hierarchy and consider beings of any type eligible for personhood. Our position admits the possibility not only of non-human animals qualifying for personhood, but of silicon-based electronic agents qualifying too. We agree with Gordon that "the bare ontological substance of a being does not convey any morally relevant features" (Gordon, 2018, p. 217).

Still, it might be held that our view is speciest because it privileges relationships with human beings and the ways in which social robots relate to us. In reply, following McMahan, we distinguish between a view that is speciest and one that is relational (McMahan, 2002). According to McMahan, the reason we have to accord humans with severe cognitive impairment the same moral status as other humans is that "we are *related* to them through the 'tie of birth,'" not that they are human (McMahan, 2002, p. 217, emphasis added). McMahan admits that Martians would not have this reason to favor humans, because they do not stand in this special relation with humans.

6 Conclusion

In conclusion, we set forth an African relational approach to social robots using Ishiguro's Klara and McKeown's Adam to illustrate. Drawing on fictional narratives that portray robots within the context of robot-human relationships allowed us to shift the question about robots' moral status from one focused on abstract innate properties, to one responsive to robots' relational possibilities. We pose questions about the moral standing of Klara and Adam within the social contexts in which they are deployed, arguing that Klara's incorporation into a human community confers on her a high moral status and that, despite Adam's intrinsic qualities of sentience and intellect, he remains mostly aloof from close relationships with his adoptive family and ultimately betrays them. Drawing on an *ubuntu* ontology and ethics, we suggest a framework for moral personhood based on these cases that addresses the broader question of how we ought to design social robots to enhance human social life. By harnessing untapped insights from Africa, our paper contributes to broader efforts

underway to enrich the discourse on social robots by better incorporating more culturally diverse points of view.

References

Ajei, M. O. (2019). Ontology and human rights. *South African Journal of Philosophy, 38*(1), 17–29.

Attfield, R. (1983). *The ethics of environmental concern*. Basil Blackwell.

Atuire, C. A. (2022). African perspectives of moral status. *BMJ Medical Humanities.* https://doi.org/10.1136/medhum-2021-012229

Barclay, L. (2000). Autonomy and the social self. In C. Mackenzie & N. Stoijar (Eds.), *Relational autonomy: Feminist perspectives on autonomy and the social self* (pp. 52–71). Oxford University Press.

Bentham, J. 1789 (1976). *Introduction to the principles of morals and legislation*. Oxford University Press, Kindle edition by Liberty Fund, Inc.

Bones, H., Ford, S., Hendery, R., et al. (2020). In the frame: The language of AI. *Philosophy & Technology, 34*(Suppl1), S23–S24.

Bostrom, N., & Yudkowsky, E. (2014). The ethics of artificial intelligence. In K. Frankish & W. M. Ramsey (Eds.), *The Cambridge handbook of artificial intelligence* (pp. 316–334). Cambridge University Press.

Calicott, J. B. (1980). Animal liberation: A triangular affair. *Environmental Ethics, 2*, 311–328.

Cave, S., & Dihal, K. (2020). The whiteness of AI. *Philosophy & Technology, 33*(4), 685–703.

Code, L. (2018). *What can she know?* Cornell University Press.

Coeckelbergh, M. (2012). *Growing moral relations*. Palgrave MacMillan.

Coeckelbergh, M. (2014). The moral standing of machines: Toward a relational and non-Cartesian hermeneutics. *Philosophy and Technology, 27*, 61–77.

Coeckelbergh, M., & Gunkel, D. J. (2014). Facing animals. *Journal of Agricultural and Environmental Ethics, 27*(5), 715–733.

Couldry, N., & Mejias, U. A. (2019). *The costs of connection: How data is colonizing human life and appropriating it for capitalism*. Stanford University Press.

Danaher, J. (2017). *Robot sex*. MIT Press.

Danaher, J. (2019). The philosophical case for robot friendship. *Journal of Posthuman Studies, 3*(1), 1–24.

Darling, K. (2017). 'Who's Johnny?' Anthropomorphic framing and human-robot interaction, integration, and policy. In P. Lin, K. Abney, & R. Jenkin (Eds.), *Robot ethics 2.0.* Oxford Scholarship Online.

Elder, A. (2017). Robot friends for autistic children: Monopoly money or counterfeit currency. In P. Lin, K. Abney, & R. Jenkin (Eds.), *Robot ethics 2.0.* Oxford Scholarship Online.

Floridi, L. (2017). Digital's cleaving power and its consequences. *Philosophy & Technology, 30*, 123–129.

Gbadegesin, S. (2002). Ènìyàn: The Yoruba concept of a person. In P. H. Coetzee & A. P. J. Roux (Eds.), *Philosophy from Africa* (2nd ed., pp. 175–191). Oxford University Press Southern Africa.

Gordon, J.-S. (2018). What do we owe to intelligent robots? *AI & Society, 35*(1), 209–223.

Gordon, J.-S. (2021). Artificial moral and legal personhood. *AI & Society, 36*, 457–471.

Gunkel, D. J. (2020). Perspectives on ethics of AI: Philosophy. In M. D. Dubber, F. Pasquale, & S. Das (Eds.), *The Oxford handbook of AI.* Oxford University Press.

Gyekye, K. (1987). *African philosophical thought: The Akan conceptual scheme.* Temple University Press.

Gyekye, K. (1996). *African cultural values.* Sankofa Publishing Co.

Gyekye, K. 1997 (2010). *Tradition and modernity: Philosophical reflections on the African experience.* Oxford University Press.

Gyekye. K. (2010). African ethics. In E. Zalta (Ed.), *Stanford encyclopedia of philosophy.* https://plato.stanford.edu/archives/fall2011/entries/african-ethics/.

Gyekye, K., (2011a). African ethics. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy.* https://plato.stanford.edu/archives/fall2011a/entries/african-ethics/.

Gyekye, K. (2011b). Person and community. In K. Gyekye (Ed.), *Tradition and modernity: Philosophical reflections on the African experience.* Oxford Scholarship Online.

Ishiguro, K. (2021). *Klara and the sun.* Knopf Doubleday Publishing Group.

Jaworska, A., & Tannenbaum, J. (2021). The grounds of moral status. In E. N. Zalta (Ed.), *Stanford Encyclopedia of philosophy.* https://plato.stanford.edu/archives/spr2021/entries/grounds-moral-status/.

Jecker, N. S. (2020a). Nothing to be ashamed of: Sex robots for older adults with disabilities. *Journal of Medical Ethics.* https://doi.org/10.1136/medethics-2020-106645

Jecker, N. S. (2020b). You've got a friend in me: Sociable robots for older adults in an age of global pandemics. *Ethics and Information Technology.* https://doi.org/10.1007/s10676-020-09546-y

Jecker, N. S. (2021a). Can we wrong a robot? *AI & Society.* https://doi.org/10.1007/s00146-021-01278-x

Jecker, N. S. (2021b). My friend, the robot; an argument for e-friendship. *Proceedings of the Institute of Electrical and Electronic Engineers (IEEE) Conference on Robot and Human Interactive Communications (RO-MAN)*, 692–697. https://doi.org/10.1109/RO-MAN50785.2021.9515429.

Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI guidelines. *Nature Medicine Intelligence, 1*(9), 389–399.

Johnson, L. E. (1993). *A morally deep world.* Cambridge University Press.

Kamm, F. (2007). *Intricate ethics: Rights, responsibilities, and permissible harm.* Oxford University Press.

Kingwell, M. (2018). Are sentient AIs persons? In M. D. Dubber, F. Pasquale, & S. Das (Eds.), *The Oxford handbook of ethics of AI.* Oxford University Press. no page number.

Kittay, E. (2005). At the margins of moral personhood. *Ethics, 116*(1), 100–131.

Korsgaard, C. M. (2005). Fellow creatures: Kantian ethics and our duties to animals. In G. B. Peterson (Ed.), *The Tanner lectures on human values.* University of Utah Press.

Lee, J. B. (2001). *Agile in a crisis. Robots show their mettle.* New York Times.

Leopold, A. 1989. (1949). *A Sand County Almanac.* Oxford University Press.

Manquet, J. J. (1960). The Kingdom of Ruanda. In D. Forde (Ed.), *African worlds: Studies in the cosmological ideas and social values of African peoples* (pp. 164–189). Oxford University Press.

Marks, A. (2021). Fostering a new era of north-south exchange and south-south collaboration. TReNDS, A United Nations Initiative. https://www.sdsntrends.org/blog/2021/northsouthexchange.

Masolo, D. A. (2004). Western and African communitarianism. In K. Wiredu (Ed.), *A companion to African philosophy* (pp. 483–498). Blackwell Publishing Ltd.

Masolo, D. A. (2010). Understanding personhood: An African philosophical anthropology. In D. A. Masolo (Ed.), *Self and community in a changing world* (pp. 135–181). Indiana University Press.

Mbiti, J. (1969). *African religions and philosophy.* Heineman.

Mbiti, J. S. (1989). *Africa religions and philosophy* (2nd ed.). Heinemann.

McEwan, I. (2019). *Machines like me.* Knopf Doubleday Publishing Group.

McMahan, J. (2002). *The ethics of killing: Problems at the margins of life.* Oxford University Press.

McMahan, J. (2010). Cognitive disability and cognitive enhancement. In E. F. Kittay & L. Carlson (Eds.), *Cognitive disability and its challenge to moral philosophy* (pp. 345–368). Wiley-Blackwell.

McStay, A. (2021). Emotional AI, ethics, and Japanese spice. *Philosophy & Technology., 34*(4), 1781–1802.

Meacham, D., & Studley, M. (2017). Could a robot care? It's all in the movement. In P. Lin, K. Abney, & R. Jenkins (Eds.), *Robot ethics 2.0.* Oxford University Press, Oxford Scholarship Online.

Menkiti, I. A. (1984). Person and community in African traditional thought. In R. A. Wright (Ed.), *African philosophy: An introduction* (3rd ed., pp. 171–181). University Press of America.

Menkiti, I. A. (2006). On the normative conception of a person. In K. Wiredu (Ed.), *A Companion to African philosophy* (pp. 324–331). Basil Blackwell.

Metz, T. (2012). An African theory of moral status. *Ethical Theory and Moral Practice, 15*(3), 387–402.

Metz, T. (2021). *A relational moral theory: African ethics in and beyond the continent*. Oxford University Press.

Metz, T. (2022). *A relational moral theory: African ethics in and beyond the continent*. Oxford University Press.

Mhlambi, S. (2020). From rationality to relationality: Ubuntu as an ethical and human rights framework for artificial intelligence governance. Harvard Kennedy Carr Center for Human Rights Policy. https://carrcenter.hks.harvard.edu/files/cchr/files/ccdp_2020-009_sabelo_b.pdf.

Mills, C. W. (2014). *The racial contract*. Cornell University Press.

Mohamed, S., Png, T.-T., & Issac, W. (2020). Decolonial AI. *Philosophy & Technology, 33*, 659–684.

Molefe, M. (2020). The place of animals in African moral philosophy. In M. Molefe (Ed.), *African personhood and applied ethics.* NISC (Pty) Ltd.

Neely, E. L. (2014). Machines and the moral community. *Philosophy and Technology, 27*, 97–111.

Nussbaum, M. C. (2006). *Frontiers of justice: Disability, nationality, species membership*. Harvard University Press.

Ramose, M. B. (2002). *African philosophy through ubuntu*. Mond Books.

Reviglio, U., & Alunge, R. (2020). "I am datafied because we are datafied": An ubuntu perspective on (relational) privacy. *Philosophy and Technology, 33*, 595–612.

Rolston, H., III. (1985). Duties to endangered species. *BioScience, 35*, 718–726.

Schweitzer, A. (1923). The ethics of reverence for life. In A. Schweitzer (Ed.), *Civilization and ethics, part II, transl.* Naish J. A & C Black Ltd., http://www.animal-rights-library.com/texts-c/schweitzer01.htm

Singer, P. W. (2009). *Wired for war*. Penguin Books Ltd.

Singer, P. (2015). *Animal liberation*. Open Road Media.

Stoljar, N. (2018). Feminist perspectives on autonomy. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*. https://plato.stanford.edu/archives/win2018/entries/feminism-autonomy/.

Taylor, P. W. (1986). *Respect for nature*. Princeton University Press.

Tooley, M. (1984). *Abortion and infanticide*. Oxford University Press.

TReNDS (Thematic Research Network on Data and Statistics). (2021). TReNDS: A global data research network advancing policy and technical solutions for sustainable development. https://www.sdsntrends.org. Accessed 28 December.

Turkle, S. (2011). *Alone together*. Basic Books.

Tutu, D. (2009). *No future without forgiveness*. Random House.

Warren, M. A. (2004). *Moral status*. Oxford University Press.

Westlund, A. (2009). Rethinking relational autonomy. *Hypatia, 24*, 26–49.

Wiley, A. (2012). Feminist epistemology: Standpoint matters. *Proceedings of the American Philosophical Association, 86*(2), 47–76.

Wiredu, K. (1980). Truth as opinion. *Philosophy and an African culture* (pp. 111–123). Cambridge University Press.

Wiredu, K. (1996). *Cultural universals and particulars*. Indiana University Press.

Wiredu, K. (2010). The moral foundations of an African culture. In K. Wiredu & K. Gyekye (Eds.), *Person and community: Ghanaian philosophical studies, I* (pp. 193–206). The Council for Research in Values and Philosophy.