

Do robots have race?

Forthcoming, *Robotics and Automation Magazine* (accepted 7 June, 2019)

Professor Robert Sparrow
Department of Philosophy
School of Philosophical, Historical and International Studies
Monash University
ARC Centre of Excellence for Electromaterials Science
Monash University
Victoria 3800
Australia.

(WORKING PAPER ONLY)

Abstract:

In a paper presented at HRI 2018, Bartneck and colleagues argued that people attribute race to humanoid robots. Our paper joined a (small) number of other papers in providing evidence that people attribute racial and/or ethnic identities to robots. In this essay, I want to offer some observations about what is at stake when we ask “do robots have race?” Drawing on resources from the philosophy of race, cultural studies, and media ethics, which may not be familiar to many readers of *Robotics & Automation Magazine*, I will argue that the discovery that robots have race that would pose unique ethical and political challenges to the project of building humanoid social robots given the historical associations between robots and slaves. Performing experiments to determine whether robots have race, informed by a proper understanding of the nature of race according to the latest scholarship, should therefore be a priority for the Human-Robot Interaction community.

Keywords: Robotics; Ethics; Robots; Race; Racism; Diversity; HRI.

Do robots have race?

In a paper presented at HRI 2018, Bartneck and colleagues argued that people attribute race to humanoid robots [1]. The experiments we conducted suggested that people are willing to identify robots as being of different races when asked and that, at least when confronted with the shooter bias test [2], people also behave in ways that suggest that they attribute race to robots. This paper joins a (small) number of other papers in providing evidence that people attribute racial and/or ethnic identities to robots [3, 4]. While the methods we used in our research have come in for some criticism [5], the paper has at least highlighted this question as a topic for future research.

In this essay, I want to offer some observations about what is at stake when we ask “do robots have race?” Drawing on resources from the philosophy of race, cultural studies, and media ethics, which may not be familiar to many readers of *Robotics & Automation Magazine*, I will argue that the discovery that people attribute race to robots would pose unique ethical and political challenges to the project of building humanoid social robots given the historical associations between robots and slaves. Performing experiments to determine whether robots have race, informed by a proper understanding of the nature of race as a “social construction” according to the latest scholarship, should therefore be a priority for the Human-Robot Interaction (HRI) community.

Social robots as “machines with meaning”

The idea that people relate to robots using the same set of evolved responses and social schemas that they use to relate to people and animals has played an important role in the development of the field of social robotics. The mere fact that an object moves without being subject to external force prompts people to treat it as if it were alive. When robots “smile”, people respond positively to them

and think that they are “happy”. When a robot extends its “arm” and “points” to something people look where it is pointing. The capacity of robots to mobilise these sorts of responses is precisely what has led engineers to believe that “social” robots could allow human beings and robots to work together more effectively: the “media equation” [6] seems especially compelling when it comes to *embodied* machines.

Robots that can be social in this way are, we might say “machines with meaning”. Machines can communicate with — enter into social life with — people only if we can understand what they mean. This is most evident when robots speak but it also applies to their gestures and bodily affect. Robots can only succeed in looking sad, for instance, to the extent that we recognise sadness in their bodies and faces.

The bodies of robots also convey meaning by virtue of their nature and appearance. We can see this, for instance, when people attribute species to some robots. Aibo is very clearly a robot dog, while Paro is a robot seal, and JustoCat a robot cat. The bodies of these robots refer to the bodies of animals with the result that people understand the robots as mechanical versions of these animals.

Sometimes, as with Hiroshi Ishiguro’s “Geminoid”, *humanoid* robots represent particular people. However, some humanoid robots seem to represent *types* of people. For instance, most sex robots are clearly female or (occasionally) male: they represent a woman (or a man) without necessarily representing any particular woman (or man).¹ One question that has received some attention in the literature is whether people attribute sex to humanoid robots more generally: a number of results suggest that they do [8-11].

¹ For an investigation of the social meanings communicated by choices in the design of sex robots, see [7].

Could robots have race?

If robots may sometimes be distinguished by species and sex, it is at least possible that some robots may also be distinguished by race. That is to say, it is possible that people attribute race to at least some robots. Indeed, as we have seen, there is some evidence that they do.

One reason why, I suspect, engineers appear to have not thought very much about the “race” of robots is because engineers relate to robots as machines, while we typically think of race as something that — if it exists at all — is possessed only by human beings.² If one is tempted by genetic or genealogical accounts of race then it seems straightforward that robots, who have neither genes nor parents nor grandparents, could not have race.

However, contemporary anthropology, race studies, and social theory have moved away from accounts of race as being “in the body” of individuals [14]. Rather, race is now recognised to be a social relation. To have race is to be recognised as, and to recognise oneself as, being a member of a particular racialised group [15]. What races there are, and who counts as which race, vary from place to place and from time to time [16]. Race, then, is “socially constructed” [17, 18]. While biological features, including skin colour, hair type, and facial physiognomy, as well as contextual features, such as accent, social status, and profession, play a role in determining which individuals are held to be members of which races in which contexts, the sorting is done by people on the basis of historically contingent ideas about which races exist and how they are to be distinguished. In any given social context, though, race tracks real social differences [16, 19].

² Two recent papers [12, 13] have drawn attention to the existence of “bias”, including racial bias, in the data about human behaviour that is used to train AIs, for the behaviour of robots and to the ethical issues this raises for engineers. However, while undoubtedly important, this is a separate issue to the question as to whether people may be biased in their treatment of robots depending on the “race” of the robot.

An understanding of race that is informed by the state-of-the-art of research in the social sciences therefore allows for the possibility that robots might have race. Robots would have race if people *treated* them as though they had race. That is, if people identified them as being White, Black, Asian, et cetera, and if the differences in appearance associated with these identifications shaped people's responses to, and behaviour around, robots in ways relevantly analogous to the ways in which race shapes relations between human beings. While not my focus here, it's also clear that if people treat robots as though they have race then robots that learn from social interactions will eventually come to behave differently depending on whether they are identified as being Black or White (for instance). By tracking differences in people's expectations and responses to people (and robots) of different races, sufficiently sophisticated robots may even come to recognise and affirm their own racial identity.

The whiteness of social robots

The idea that robots might have race should be a cause of some anxiety for the social robotics community — and especially for those engineers building humanoid robots — given that the robots being built by this community overwhelmingly tend to have white surfaces and, thus — if they do have race — are most plausibly thought of as being White.³ For instance, if one Google-image searches “humanoid robot” or “android,” the results that are produced are overwhelmingly pictures of robots with predominantly white surfaces. Admittedly, some of these images originate from science-fiction or the media but the glossy white plastic surfaces of the media stereotypes of robots both reflect and shape the real-world products built by engineers.

³ It is worth noting that not all engineers think that it is a good idea to make robots that look human and that there exists a significant body of opinion that is critical of the deception involved in designing robots that encourage people to anthropomorphise them. See, for instance, [20] and [21].

A slightly more scientific measure is provided by the Anthropomorphic roBOT (ABOT) database [22]. Of the 125 robots included in this dataset at the time of writing that achieve a “human-likeness” score of 25 or more (which excludes only the most obviously non-humanoid robots), 61 are entirely white (or predominantly white, having only grey or black actuators), 21 are predominantly white with a small splash of another colour (other than grey or black actuators), eight are mostly white but a significant portion of their surface is another colour (other than grey or black actuators), sixteen are predominantly silver, fourteen are predominantly blue, red, green, yellow, or another bright colour, two are predominantly grey, and only three have predominantly black or brown surfaces.⁴

ABOT includes many robots being built in university robotics laboratories as student projects, most of which are destined never to go into production. So, a final piece of data about the aesthetics of social robots is provided by the appearance of those robots that are marketed to the public and/or are widely used in HRI research, which includes Pepper, Nao, Buddy, and Jibo. All these robots have predominantly white surfaces.

Thus, while it is difficult to put a precise figure on the percentage of robots that have predominantly white surfaces, it is clear that the vast majority of humanoid robots being built today are white and therefore, perhaps, White.

There are a number of exceptions to this rule but these serve mainly to show that at least some robots *do* have race.

First, there exist a small number of robots that are modelled on individuals. Famously, Hiroshi Ishiguro manufactured his “Geminoid” robots in the image of himself and his partner. “Bina48”, is an

⁴ A number of the silver or brightly coloured robots also have faces that are white. Five of the robots with white skin tones have hair or facial features that make it clear that they are supposed to be understood as Asian.

animatronic talking head, which Martine Rottblatt commissioned and modelled on her partner, Bina, who is African-American. Insofar as the individuals upon which the robots are modelled have race, these robots must also have the *same* race otherwise they would not succeed in representing these individuals.

Second, there are robots that are built in the service of national pride, to show that local engineers are participating at the forefront of research. For instance, researchers at the Interactive Robots and Media Laboratory at United Arab Emirates University built a robot which they named Ibn Sina, after an 11th Century Muslim philosopher. Video footage of this robot reveals it to be wearing a turban, have olive “skin”, and displaying a long grey beard. These features seem intended to provide the robot with a non-white race for the sake of signifying its national origins. Some robots built in Japan arguably also play the same political role. Again, these robots couldn’t succeed in appearing Arabic or Japanese unless people recognise them as having race.

There is a third exception to the rule that humanoid robots are white, which very clearly shows that some robots *do* have race. “Rastus the mechanical man” was a robot built, and toured around the United States, by Westinghouse Corporation in the 1930s [23]. Rastus had black rubberised skin, “Negroid” features, and was dressed in cotton overalls, in order to appear like a (racist) stereotype of a Black sharecropper. “He” was used in performances that involved shooting an apple off his head. Text underneath a photo of Rastus in a newspaper report at the time describes Rastus as a “Realistic Mechanical Negro” [24]. If the black surface of this robot could succeed in making it Black, this suggests — although it does not prove — that the white surfaces of Pepper, Nao, and Buddy, may mean that they are, by contrast, White.⁵

⁵ Another significant exception to the rule that humanoid robots are white are those robots that are a gleaming silver or gold colour. While I do not have space to make the argument at length here, it’s worth noting

The race history of robots

Robots featured in science-fiction long before they began to emerge from the laboratories of engineers. Indeed, it is fair to suggest that the expectations of the public — and thus the research agendas of engineers — were determined by this history of representations in books and films. For this reason, it is worth spending a little time discussing the cultural legacy of these representations, which complexifies the question of the race of robots.

Interestingly, there is a great deal of historical evidence that robots have always been understood as having race. As is well known, the very word robot means “worker” or “slave” in Czech. Capek, who coined the word, was arguably more concerned about class relations, and the dehumanising implications of modern technology, than he was about the slave trade. However, the cultural resonances of robots quickly included the institution and history of slavery, which became especially pronounced when the project of robotics was embraced in the United States. The popularity of stories about robot uprisings in science fiction all too clearly represents the working through of social and political anxieties about slavery and colonialism [26, 27]. As Louis Chude-Sokei [28] has argued, the debate in literature and popular culture about whether robots are — or could become — sentient parallels, with eerie precision, the historical debate about whether Africans might have souls. The

that it is possible — even likely — that these robots are still racially coded as White. In European and North American contexts (and perhaps elsewhere), non-White peoples are marked by their skin colours but to be White is just to be “normal”. Moreover, Whiteness is a status that may be awarded or taken away: to be White is to be included in an imagined community defined by association with a set of purported virtues such as purity and decency [25]. Silver or gold surfaces need not necessarily disqualify a robot of Whiteness if it is seen to possess such virtues.

history of robots as literary and filmic devices to talk about slaves and colonised peoples therefore strongly suggests that the idea of a robot is racially inflected as Black or Brown [29, 30].⁶

The racial politics that robots have by virtue of being slaves, though, are in tension with another way in which science-fiction representations of the future are colour-coded.⁷ The predominant aesthetic of “classic” science-fiction valorises smooth white and silver curves [33]. The glossy surfaces of this style gesture towards a manufactured environment devoid of all dirt or corruption. To be frank, I suspect that the cultural hegemony of this aesthetic is why engineers build the white robots that they do. However, this aesthetic has political content: the dirt or darkness that is excluded from this future represents people of colour. The main stream of science-fiction has almost always imagined the people living in the future as White: if people of colour appear in these visions it is usually as a threat to the harmony of White society [33].⁸

The visual race politics of contemporary robots, then, are at odds with the historical race politics of robots. Culturally speaking, we might say, modern robots have White bodies but servile (read “Black”) souls. Which of these cultural preconceptions might prove more influential in determining the race of any particular robot, in any particular context, is an empirical question.

Why it matters

⁶ One of the difficulties in writing about race is that the terms used to discuss race in any given context are themselves the products of, and inflected by, racism. My use of the terms White, Black, and Brown to refer to racial identifications should in no way be read to imply that I believe that these are biological categories.

⁷ It is highly striking, in this context, that neither of two papers that explicitly argue that we should conceptualise robots as slaves [31], [32] discusses the real world history of slavery.

⁸ Later representations of the future in science-fiction allowed for the presence of people of colour as full members of multicultural societies in which “race did not matter”. However, the price of this “transcending” of race was that non-white citizens became honorary Whites in a social milieu that remained fundamentally White [33].

One reason why it would matter if robots are revealed to have race is the fact that, as I observed above, because the vast majority of humanoid robots have white surfaces, the vast majority of humanoid robots may turn out to be White. If it turns out that engineers building humanoid robots are, with very few exceptions, building robots that are racially coded as White, this would invite an accusation of — at least implicit — racism. The assumption that the “default” human body is a White body is a racist one. The majority of human bodies *aren't* white, in fact, and there is no need to think that human beings have any standard skin colour. Granted that, as I observed above, most engineers may have chosen white surfaces for their robots solely because they wanted them to look futuristic, a failure to notice that this meant that they were building White robots would still invite criticism. Perhaps more importantly, a lack of diversity in robots should be expected to have the same negative social consequences, including consequences for equity, as a lack of diversity in other forms of media. If members of minorities don't see themselves represented in public life it becomes harder for them to imagine themselves as equal citizens and especially as people in positions of social authority. A lack of diversity in the media also contributes to pernicious social stereotyping more generally. Regardless of the intentions — or the moral blameworthiness — of their designers, building only white robots would exacerbate racism.

The implications of robots having race will be different if it turns out that, because of their role as servants, white robots are nevertheless identified as Black or Brown. In this case, the lack of diversity in the robots that engineers are building will not be so problematic given that, at one level at least, this will expand the representation of people of colour in the media more generally. Yet the politics of having all those artefacts built for service roles being understood to be Black or Brown is even more troubling. Admittedly, should this occur, it will be the result of pre-existing cultural prejudices. Nevertheless,

replicating and reinforcing these by developing a new class of artefacts that are understood to represent Black or Brown servants seems deeply problematic.⁹

Another reason why it would matter if robots have race is that it would have implications for their performance. The race of robots will matter because race matters in contemporary societies. In particular, race shapes how people relate to each other. If robots have race, then how people relate to a robot is likely to be shaped by its race. If we imagine a White medical robot placed into a practice that serves mostly African-American patients, for instance, we might expect the results to be quite different than those produced by a robot that was racially coded as Black. Conversely, a Black robot in a medical practice serving Whites in the rural south of the United States might not produce the social rapport its designers hope for. Unless engineers are aware of these dynamics and consider the race of the robots they are building, they may struggle to facilitate the human-robot relationships they intend.

Revealing the race of robots

Whether or not – or, perhaps, which - robots have race is ultimately an empirical question, which I cannot hope to settle in an essay of this sort. However, I hope my discussion has served to clarify the nature of the question that needs to be investigated. To settle the question of whether robots have race we will need to investigate whether people respond to variations in the surface colour, and other features, of humanoid robots in the same way that they do to similar variations that are held to be markers of race in people. If such variations trigger racialised responses, for instance, that will be strong evidence that people are attributing race to robots.

⁹ This also suggests that — should white robots be racially coded as White — building robots with darker surfaces might be similarly problematic.

A starting point to determine whether people recognise race in robots would be to ask them about the race of robots with different surface colours [1]. Concerns about priming will inevitably arise here [5], even though they are slightly misplaced given that at least part of what needs to be investigated when we seek to determine whether people attribute race to robots is precisely the conscious attitudes of those who encounter robots, and that respondents will presumably have the option of answering “none” or “does not apply”. Nevertheless, ideally one would try to research the attitudes of the community without the community becoming aware of the researchers’ interests. For instance, one might include a question about race in the context of other questions about the traits that people attribute to robots. Alternatively, one might ask open questions inviting research participants to describe various robots, with different surface colours, and see if they refer to race in their answers.

Other cues that might influence the attribution of race to robots include their behaviour, accent and mode of speech (if any), name [34], “facial” physiognomy, intended function (for instance, sweeping the floor versus providing medical advice), and social context. There is a rich vein of research topics that would reward investigation here.

Note that it is not necessary to show that everyone agrees on the race of particular robots in order to show that robots have race. It will be enough to show that most people do attribute race to robots and that there is a significant amount of consensus on the race of robots with particular appearances. The fact that two people might disagree on whether a third person was White or Asian typically doesn’t lead us to the belief that the third person has no race. Nor are we swayed in this belief when we encounter people who claim not to recognise race and insist that “we are all just human beings.”

Importantly, even if it turns out that people are reluctant to consciously attribute race to robots, this will not be the end of the matter. It is not uncommon, for instance, for those people who say, apparently sincerely, that they don’t care about race, or even that they don’t notice race, nevertheless

to be influenced by race when it comes to their behaviour and/or attitudes towards other people [35, 36]. For that matter, it is also possible that even though people are willing to attribute race to robots their attributions make no other difference to their behaviour.

It will also be necessary, then, to examine how altering the surface colours (and perhaps other racial codings) of robots affect how people respond to and behave around them regardless of what those people say about the robots' race. The HRI research community would well advised to look to the large literature in social psychology on race [37] for experimental protocols to detect when people are responding to race in each other and consider which might most easily be adapted to detect racialised responses to robots. For instance, researchers might investigate how altering racial cues in robots affects what other characteristics people attribute to them. They might also explore how racial congruence (and racial difference) between robots and their interlocutors shapes how inclined people are to trust them and/or how they speak to them.

One thing we already know about race from the social sciences is that the salience of race may vary in different contexts, which means that experiments performed using undergraduate students in laboratory conditions, or with participants on Mechanical Turk, may not be a reliable means of examining these questions as they will arise in practice. As much as possible, then, it would be best to perform experiments with real robots "in the wild": that is, in the sorts of settings in which the HRI community imagines robots appearing in the future. It is also possible that people will change their behaviour around robots in ways that depend upon the race of the robot *and* on the race of the person. It would therefore be advisable to repeat any experiments with different cohorts of participants, with different racial identities in order to control for the impact of the race of the participants. Again, insofar as this is a feature of the way race works in human beings, discovering that responses to different

surface colours of robots were shaped by the racial identification of the person interacting with the robot would itself be further evidence that robots have race.

The setup for both these sorts of experiments will be complicated by widespread disagreement about whether races exist at all, how many races there are, how they should be named, and what features they possess. Answers to all these questions are likely to be both socially contextual and intensely controversial. For this reason, it would benefit the advance of knowledge in this area if those called upon to referee such papers did not insist that the authors adopt the particular racial categories, and names for races, that are privileged in the referee's community. The question as to whether robots are placed in *any* racial category is more interesting than the boundaries of the categories into which they are placed in any given experiment. If the same robots are held to have different races in different contexts, or if the racial categories into which robots with particular surfaces are placed changes from location to location, this would not distinguish them from human beings.

Finally, it will be important that experiments have relevant controls. That is to say, it will be important to check that any behavioural changes that are thought to be a response to the race of the robot are not responses to cues that we would be reluctant to hold were constitutive of race in other contexts. For instance, one possibility is that *any* object with a surface colour close to any of those that people typically associate with a particular race might prompt observable changes in behaviour. Notice, however, that if this phenomenon proves to be significant in shaping encounters between robots and human beings this will itself be enough to suggest that engineers need to think more about the colour of the surfaces of the robots they build.

If some humanoid robots are revealed to have race in at least some contexts, a further question will arise about just how far attribution of race extends. Do people attribute race to robots that are blue or yellow, for instance? Do they attribute race to robots that stray from the human body schema or that

aren't humanoid at all? A related question with practical import is whether it is possible to build a humanoid robot that does *not* have race. The answers to these questions will be of vital interest to engineers who might wish to build social robots without having to deal with the ethical and political complexities introduced when robots do have race.

Robots, race, and HRI

Engineers are accustomed to the struggle of bending the material world to their will and are increasingly skilled at doing so. The project of building social robots, though, requires that engineers consider the meaning that robots — and their actions — have for other people, a task that is arguably less familiar and for which engineers' training does not always prepare them. As I have argued here, the possibility that one of the meanings conveyed by robots might be "race" problematizes the largely unexamined assumption that robots should have white surfaces. Moreover, the history of slavery and the historical associations between robots and slaves mean that the discovery that robots have race would pose a unique ethical and political challenge to the project of building humanoid social robots. If humanoid robots are understood to be White, this risks reinforcing racism; if they are understood to be Black or Brown, building robots for service roles will also reinforce racism. My hope is that, by clarifying the nature of the question "do robots have race" and suggesting some means by which it might be answered, this paper may assist the engineering profession to navigate these troubled ethical and political waters in order to better realise the potential of social robotics.

Acknowledgements

The research for this paper was supported under the Australian Research Council's Centres of Excellence funding scheme (project CE140100012). The views expressed herein are those of the author and are not necessarily those of the Australian Research Council. I would like to thank Max Cappuccio for the invitation to present at the *2nd Joint UAE Symposium on Social Robotics*, in 2016, which first provoked me to work on this topic. Some of the research for this paper was conducted while I was a Visiting Fellow in the Department of Philosophy at Carnegie Mellon University: I would like to thank David Danks for the opportunity to visit CMU. Conversations and correspondence with Christoph Bartneck and Friederike Eysel, in the course of our collaboration on an earlier manuscript, and with Max Cappuccio, Louis Chude-Sokei, Ayanna Howard, Ole Koksvik, Sascha Morrell, Illah Nourbakhsh, and Aaron Steinfeld and have influenced my thinking on the race politics of robots. I am grateful to Ole Koksvik and Mark Howard for assistance with bibliographic research and with preparing the manuscript for publication.

Bibliography

- [1] C. Bartneck *et al.*, "Robots And Racism," *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction - HRI '18*, pp. 196-204, 2018.
- [2] J. Correll, B. Park, C. M. Judd, and B. Wittenbrink, "The Police Officer's Dilemma: Using Ethnicity to Disambiguate Potentially Threatening Individuals," *Journal of Personality and Social Psychology*, vol. 83, no. 6, pp. 1314-1329, 2002.
- [3] F. Eysel and S. Loughnan, "'It Don't Matter If You're Black or White'?" Effects of Robot Appearance and User Prejudice on Evaluations of a Newly Developed Robot Companion," in *International Conference on Social Robotics*, 2013, pp. 422-431.

- [4] M. Strait, A. S. Ramos, V. Contreras, and N. Garcia, "Robots Racialized in the Likeness of Marginalized Social Identities are Subject to Greater Dehumanization than those racialized as White," in *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2018, pp. 452-457: IEEE.
- [5] T. Ogunyale, D. A. Bryant, and A. Howard, "Does Removing Stereotype Priming Remove Bias? A Pilot Human-Robot Interaction Study," in *5th Workshop on Fairness, Accountability, and Transparency in Machine Learning*, 2018, pp. 3-7.
- [6] B. Reeves and C. Nass, *The Media Equation How People Treat Computers, Television, and New Media Like Real People and Places*. CUP, 1998.
- [7] R. Sparrow, "Robots, Rape, and Representation," *International Journal of Social Robotics*, vol. 9, no. 4, pp. 465-477, 2017.
- [8] J. Bernotat, F. Eyssel, and J. Sachse, "Shape It – The Influence of Robot Body Shape on Gender Perception in Robots," *International Conference on Social Robotics 2017*, pp. 75-84, 2017.
- [9] F. Eyssel and F. Hegel, "(S)he's Got the Look: Gender Stereotyping of Robots," *Journal of Applied Social Psychology*, vol. 42, no. 9, pp. 2213-2230, 2012.
- [10] J. Otterbacher and M. Talias, "S/he's too Warm/Agentic!: The Influence of Gender on Uncanny Reactions to Robots," *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 214-223, 2017.
- [11] M. Siegel, C. Breazeal, and M. I. Norton, "Persuasive Robotics: The Influence of Robot Gender on Human Behavior," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 2563-2568.
- [12] A. Howard and J. Borenstein, "Hacking the Human Bias in Robotics," *ACM Transactions on Human-Robot Interaction (THRI)*, vol. 7, no. 1, article 3, 2018.

- [13] A. Howard and J. Borenstein, "The Ugly Truth About Ourselves and Our Robot Creations: The Problem of Bias and Social Inequity," *Science and Engineering Ethics*, vol. 24, no. 5, pp. 1521-1536, 2018.
- [14] A. R. Templeton, "Biological Races in Humans," *Studies in History and Philosophy of Science Part C :Studies in History and Philosophy of Biological and Biomedical Sciences*, vol. 44, no. 3, pp. 262-271, 2013.
- [15] N. A. Denton and G. D. Deane, "Researching Race and Ethnicity: Methodological Issues," in *The SAGE Handbook of Race and Ethnic Studies*, P. H. Collins and J. Solomos, Eds., 2010.
- [16] M. Root, "How We Divide the World," *Philosophy of Science. Supplement. Proceedings of the 1998 Biennial Meetings of the Philosophy of Science Association. Part II: Symposia Papers.*, vol. 67, pp. S268-S639, 2018.
- [17] K. A. Appiah, "Race, Culture, Identity: Misunderstood Connections," in *Color Conscious: The Political Morality of Race*, K. A. Appiah and A. Gutmann, Eds., 1996, pp. 30-105.
- [18] N. Zack, *Philosophy of Science and Race*. Routledge, 2003, pp. 87-88.
- [19] A. Smedley and B. D. Smedley, "Race as Biology is Fiction, Racism as a Social Problem is Real: Anthropological and Historical Perspectives on the Social Construction of Race," *American Psychologist*, vol. 60, no. 1, pp. 16-26, 2005.
- [20] M. Boden *et al.*, "Principles of Robotics: Regulating Robots in the Real World," *Connection Science*, vol. 29, no. 2, pp. 124-129, 2017.
- [21] R. Sparrow and L. Sparrow, "In the Hands of Machines? The Future of Aged Care," *Minds and Machines*, vol. 16, no. 2, pp. 141-161, 2006.
- [22] E. Phillips, X. Zhao, D. Ullman, and B. F. Malle, "What is Human-like?: Decomposing Robots' Human-like Appearance Using the Anthropomorphic roBOT (ABOT) Database," *HRI '18: 2018 ACM/IEEE International Conference on Human-Robot Interaction*, vol. 18, pp. 105-113, 2018.

- [23] B. Johnson, "Energy Slaves: Carbon Technologies, Climate Change, and the Stratified History of the Fossil Economy," *American Quarterly*, vol. 68, no. 4, pp. 955-979, 2016.
- [24] "Mechanical Men That Excel Any Human Being," in *San Antonio Light*, ed, 1931, p. 48.
- [25] D. R. Roediger, *Working Toward Whiteness: How America's Immigrants Became White: The Strange Journey from Ellis Island to the Suburbs*. Basic Books, 2006.
- [26] S. Morrell, "Zombies, Robots, Race, and Modern Labour," *Affirmations: of the modern*, vol. 2, no. 2, pp. 101-134, 2015.
- [27] J. Rieder, *Colonialism and the Emergence of Science Fiction*. Wesleyan University Press, 2012.
- [28] L. Chude-Sokei, *The Sound of Culture: Diaspora and Black Technopoetics*. Wesleyan University Press, 2016.
- [29] G. J. Hampton, *Imagining Slaves and Robots in Literature, Film, and Popular Culture: Reinventing Yesterday's Slave With Tomorrow's Robot*. London: Lexington Books, 2015.
- [30] D. Kakoudaki, *Anatomy of a Robot: Literature, Cinema, and the Cultural Work of Artificial People*. Rutgers University Press, 2014.
- [31] J. J. Bryson, "Robots Should be Slaves," in *Close Engagements With Artificial Companions: Key Social, Psychological, Ethical and Design Issues*, Y. Wilks, Ed.: John Benjamins Publishing, 2010, pp. 63-74.
- [32] S. Peterson, "Designing People to Serve " in *Robot Ethics: The Ethical and Social Implications of Robotics*, P. Lin, G. A. Bekey, and K. Abney, Eds. Cambridge: MIT Press, 2011, pp. 283-298.
- [33] D. Bernardi, *Star Trek and History: Race-ing Toward a White Future*. Rutgers University Press, 1999.
- [34] F. Eyssel and D. Kuchenbrandt, "Social Categorization of Social Robots: Anthropomorphism as a Function of Robot Group Membership," *British Journal of Social Psychology*, vol. 51, no. 4, pp. 724-731, 2012.

- [35] B. K. Payne, J. A. Krosnick, J. Pasek, Y. Lelkes, O. Akhtar, and T. Tompson, "Implicit and Explicit Prejudice in the 2008 American Presidential Election," *Journal of Experimental Social Psychology*, vol. 46, no. 2, pp. 367-374, 2010.
- [36] D. A. Stanley, P. Sokol-Hessner, M. R. Banaji, E. A. Phelps, and D. Kahneman, "Implicit Race Attitudes Predict Trustworthiness Judgments and Economic Trust Decisions," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 108, pp. 7710-7715, 2011.
- [37] J. A. Richeson and S. R. Sommers, "Toward a Social Psychology of Race and Race Relations for the Twenty-First Century," *Annual Review of Psychology*, vol. 67, pp. 439-463, 2016.