



REVIEW ARTICLE

Machine morality, moral progress, and the looming environmental disaster

Ben Kenward¹ | Thomas Sinclair² ¹Centre for Psychological Research, Oxford Brookes University, Oxford, UK²Wadham College, University of Oxford, Oxford, UK**Correspondence**

Ben Kenward, Centre for Psychological Research, Oxford Brookes University, Oxford, UK.

Email: benkenward@brookes.ac.uk**Abstract**

The creation of artificial moral systems requires making difficult choices about which of varying human value sets should be instantiated. The industry-standard approach is to seek and encode moral consensus. Here the authors' argue, based on evidence from empirical psychology, that encoding current moral consensus risks reinforcing current norms, and thus inhibiting moral progress. However, so do efforts to encode progressive norms. Machine ethics is thus caught between a rock and a hard place. The problem is particularly acute when progress beyond prevailing moral norms is particularly urgent, as is currently the case due to the inadequacy of prevailing moral norms in the face of the climate and ecological crisis.

1 | MACHINE MORALITY

It has been argued that the development of machines programmed to make moral decisions risks atrophying human moral reasoning capacities, as humans come to outsource to the machines responsibility for decisions in which moral values are at stake [1, 2]. In this article, we identify a further concern that moral machines risk impeding moral progress. Though our argument is general, the concern is particularly acute in the context of the urgent need for such progress at the present moment, in which prevailing moral norms are inadequate to address the ongoing environmental crises of climate breakdown [3] and ecological collapse [4] and the major threat to human life they pose.

We define morality as a set of values, in light of which those holding the values regard some types of actions as good (to be striven after) and others as bad (to be avoided) [5]. Moral behaviour is defined as behaviour motivated by or designed to realise such values. We further define moral machines as those engaging in computational operations involving representations of values that are intended to reflect morality in order to yield moral behaviour, either on the part of the machines themselves or on the part of humans to whom the machines issue guidance. By these definitions, a fire alarm is not a moral machine, even though it tends to produce outcomes that humans approve of, because it involves no internal representations of value. By contrast, a self-driving vehicle that cannot avoid a collision but

makes a decision on whether to collide with three adults or two children [6] is a moral machine because it must operate on values derived from human morality to make the decision. Although these definitions may be challenged in various ways, they are sufficient for our argument here.

Encoding moral values in the machines we design is increasingly a necessity to the extent that these machines are built to complete complex tasks involving interactions with humans but without human supervision, as in the case of self-driving cars and other autonomous robots. Engineers are usually regarded as obliged to programme moral behaviour in such machines [2, 7–10]. But moral machines may also be designed for decision-support roles, where it is argued that morally salient decisions that are currently normally made by humans can be made better or more efficiently by artificial systems [2, 8, 9, 11]. For example, longer prison sentences are thought to be appropriate in the case of criminals who are more likely to reoffend, but algorithms are (in some cases) better than human judges at predicting reoffending rates [12, 13]. In more private contexts, it has been argued that artificial systems could assist humans in living up to moral standards by providing moral advice [14–18].

A key question facing engineers building moral machines is this: which morality should they program? Which moral values in particular should be reflected in the machines' operations? Given the diversity of human moralities, the answer is not

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2021 The Authors. *Cognitive Computation and Systems* published by John Wiley & Sons Ltd on behalf of The Institution of Engineering and Technology and Shenzhen University.

obvious. One approach is implied by the definition of machine ethics offered by the relevant international industrial body, according to which ethical machines are those whose behaviour is based on human moral values taken to be universal, such as human rights [2, 9]. This suggests that moral machine engineers are responding to the problem of diverse human moralities by attempting to identify universal values and encode these (though some more nuanced approaches have also been proposed [6, 19–22]). However, universal values may be harder to identify than it seems. Even human rights are not uncontroversial [21, 23]. Similarly, the consensus on the importance of fairness, another value that is emphasised in many of the available sets of ethical guidelines for artificial intelligence, is stronger amongst demographics typical for engineers than for other humans [24]. Perhaps more importantly, even if minimal principles such as human rights were universally accepted, such principles are not sufficiently comprehensive or specific to guide all the decisions that machine ethicists want machines to be able to make [6]. No consensus account of human rights determines how a self-driving car should choose in a conflict between protecting the life of its passenger and protecting the life of a pedestrian, for example, or how an AI system advising on the distribution of scarce medical resources should weigh the lives of different people with different diseases.

Faced with these challenges, as well as the distinct concern that the systems they build might not be trusted by humans, engineers have aimed to design machines that instantiate the values of their primary users as much as possible [9, 24, 25]. Indeed, research on human trust of artificial decisions indicates that reliability, defined as the consistent production of expected decisions, is key [26]. Such research has focussed primarily on perception of competence to make non-moral judgements. But a small body of work on human responses to artificial systems that are explicitly presented as making moral decisions confirms that humans tend to prefer such systems when they make judgements in line with their own judgements and those of other humans [27–29]. Further, trust in a given system tends to increase over time when people experience judgements in alignment with their expectations [26]. Especially given current low levels of trust in complex automated systems [30], it is unsurprising that organisations hoping for market success embed values in their machines that they take to reflect those of users.

2 | MACHINE MORALITY AND MORAL PROGRESS

Most moral machines designed to perform complex tasks or provide decision-making assistance to humans will, then, reflect answers to the key question that go beyond consensus-reflecting commitments to minimal moral values such as human rights. This obviously raises concerns about the legitimacy of any particular machine morality with respect to the people with whom a given machine will interact or

whose lives will be affected by the decisions it informs. These concerns are heightened by recent studies suggesting that algorithms trained or programed to replicate human decision making may instantiate existing human biases, effectively reflecting the distinctive moral failings of a particular group as well as the values it profess [31, 32]. Job applicant screening algorithms have been shown to discriminate against women, for example, reflecting a sexism that is at odds with the professed values of those whose behaviour the algorithm was designed to replicate [33].

However, our focus here is on the risk not that moral machines will perpetuate commonly occurring violations of prevailing moral norms, but that they will calcify those norms and the values they reflect, impeding the ordinary mechanisms of moral progress. If it were possible to restrict machine morality to minimal, universal values such as human rights, this might not seem a particularly grave concern (though we note, again, that human rights are not uncontroversial). But, as we argued above, no such restriction is realistic for the kinds of moral machines that are now being built. These machines reflect values that go beyond the minimal and universal. Even at the abstract level, there is reason to think that progress beyond these values may be desirable. If moral machines risk calcifying them, that should be cause for concern in itself. Moreover, as we will go on to argue, there is reason to think that progress beyond the values likely to be calcified by the moral machines being designed and built today is urgently needed. The cause for concern is correspondingly greater.

Before we turn to the way in which moral machines risk impeding mechanisms of moral progress, we should explain what we mean by ‘moral progress’. We have said that it is a challenge for engineers to select any particular morality from among the diversity of human moralities, and we have argued that minimal moral values do not provide a sufficiently comprehensive basis for overcoming that challenge. It might be thought that we face more or less the same problem in identifying values on the basis of which to formulate a criterion of moral progress, on the basis of which to argue that moral machines risk impeding it. If we could overcome that problem by identifying the necessary values, then so, it would seem, could the engineers of moral machines.

In response, we do not rest our argument on a complete set of values on the basis of which to formulate criteria of moral progress. If we could plausibly identify any such values, that would indeed appear to answer the key question and undercut the concern about value calcification. Instead, we suggest that the following minimal principle of moral progress is plausible enough to justify the concerns that we raise:

Minimal Principle of Moral Progress

Moral progress obtains if the values internal to a morality change in ways that alter participant behaviour in ways that substantially reduce the likelihood of catastrophic, large-scale suffering without a corresponding increase in human rights violations.

The intuitive idea underlying the Minimal Principle is that all other things equal, a development in a set of values represents an improvement if the community whose behaviour it regulates is steered away from catastrophic outcomes such as nuclear war or severe environmental collapse. No doubt refinements could be made to the Minimal Principle that would improve the degree to which it captures this underlying idea, but we take it that the point is clear. Note that neither the Minimal Principle as we have formulated nor any principle likely to capture the underlying idea implies any determinate set of progressive values that might provide the basis for an answer to the key question facing moral machine engineers. In particular, the Minimal Principle does not imply act utilitarian values, understood as those favouring individual acts that tend to maximise aggregate utility or welfare and disfavouring individual acts that tend to do anything else [34], since such values do not provide clear protection against human rights violations [35, 36]. Also, note that the Minimal Principle offers a sufficient condition of moral progress, not a necessary condition. We do not deny that there can be moral progress that the Minimal Principle does not provide grounds for classifying as such.

3 | MACHINE MORALITY IS LIKELY TO IMPEDE MORAL PROGRESS

We have argued that because people prefer machines whose behaviour chimes with their own preferences, then machines that make moral decisions will tend to be designed to reflect the dominant behavioural moral status quo. When machines reflect the status quo, then to demonstrate that they will impede moral progress, it is sufficient to demonstrate that people will sometimes be influenced by these machines when they might instead have been influenced by progressive human voices. (For the response that moral machines need not reflect the status quo, see the next section.)

Individuals already frequently respond to moral choices by relying on external expertise such as religious authorities. In medical dilemmas, for example, even non-religious Jews sometimes ask rabbis to make decisions on their behalf in a manner that has been described as moral “outsourcing” [37, 38]. Similarly, automatic intelligent systems are valued by humans because they save us the effort and worry of making decisions for ourselves. People can easily come to heavily rely on complex automated systems, to the extent that their own motivation and ability to make relevant decisions is reduced [1, 2, 39]. A well-known example of this is the phenomenon of transportation accidents due to over-reliance on auto-pilots [40]. When automatic systems consistently make decisions that appear correct, human motivation and ability to question the systems are reduced.

It has been suggested that over-reliance on algorithmic systems in the context of migration (as used in Canada) already represents a form of moral outsourcing that is reinforcing existing norms about the values of different types of migrants [41, 42]. As the use of such systems becomes the norm, they

have the potential to become the type of established cultural institution that places inertia on cultural evolution (as the sociologist Bourdieu discussed [43]).

Legal scholars have discussed the way in which formalising social rules can result in the perpetuation of practices which later generations come to regard as suboptimal but are difficult to change because of inertia within the system. Deakin [44] gives examples of ‘frozen accidents’ in English employment law, where laws once regarded as appropriate to govern the ‘master-servant’ relationship are no longer optimal for governing the modern employer-employee relationship but are still in place. Deakin argues that this is in part because of a tendency for new legal frameworks to rely on the older ones in ways that make older ones hard to change, and in part because of a tendency to rationalise to defend the status quo (a well-known and wide-spread psychological phenomenon [45]). Given the functional similarity between machine morality and cultural institutions such as the law—they are both systems intended to issue relatively objective, reliable, and definitive judgements as to appropriate behaviour—it is reasonable to assume machine morality could also lead to similar ‘frozen accidents’ because of similar institutional and psychological processes.

Perhaps the strongest motivation for becoming comfortable with moral outsourcing to machines is that it offers a further psychological defence against the aversive state of cognitive dissonance that is experienced when our actions are not congruent with our explicit values [46, 47]. Moral behaviour is an expression of a combination of implicit and explicit values. Implicit values are strongly internalised and influence behaviour without the necessity for conscious thought. Explicit values are those which we claim to abide by. Because talk is cheap, our explicit values tend to be more progressive [48]. Morally relevant behaviour is strongly influenced by implicitly held values, however [49–51]. For example, treatment of outgroups is at least as strongly influenced by implicitly held attitudes as by explicit attitudes ([52], for a large-scale meta-analysis see [53]). Reference to moral advice from machines, if it is based on existing behavioural norms (or encoded reflections of them such as laws), may therefore help us to justify behaviour that is not in line with our explicitly claimed values. Consultation with such moral machines has the potential to fulfil the same moral function as token moral efforts such as ‘ethical’ consumption, which tends to make us feel good in a way which licences negative behaviour even outweighing the token positive acts [54, 55].

The history of moral progress is the history of hard-fought struggles by those who succeed in persuading the majority to adopt an initially minority view [56–58]. These social movements successfully utilise a broad diversity of inter-human tactics, from reasoned persuasion to example-setting to emotive acts of self-sacrifice [59, 60]. To the extent that outsourcing morality to machines places some decision making outside the realms of these processes, rendering us less susceptible to influence from progressive human values, moral machines represent a risk to moral progress.

4 | WHY ENCODING PROGRESSIVE VALUES IS A PROBLEMATIC SOLUTION

It might be objected that moral machine engineers can, in recognition of the concerns we have just identified, simply programme the machines they build with more progressive values, and thereby address the concerns [14–18]. However, this is a problematic proposal. One reason for this is psychological and empirical; the other one is political and philosophical.

The political and philosophical reason is as follows. As we noted above in our discussion of the Minimal Principle of Moral Progress, no determinate set of progressive moral values is implied by that principle. By contrast, any particular choice of progressive moral values, even if a morality instantiating these would indeed constitute progress according to the Minimal Principle, faces problems of legitimacy. Efforts to reflect prevailing moral norms, even if any particular choice inevitably involves privileging some contemporary moralities over others, can at least claim a sort of democratic legitimacy in virtue of that effort. By contrast, an engineer's choice of one from among the many possible sets of progressive moral values that would constitute progress according to the Minimal Principle must inevitably reflect that engineer's particular outlook and reasoning, or at best the outlook and reasoning of the select group involved in making the choice. In brief, there is no comparably legitimate substitute for the mechanisms of human moral progress that such a choice would short-circuit.

The psychological and empirical reason for thinking that the concerns we have raised cannot easily be addressed simply by programming machines with more progressive values is that machines programmed with such values risk being rejected by the humans with whom they interact. There are reasons to expect that humans will react negatively in interactions with such machines even if they agree with the idea in principle. The first has to do with the phenomenon of 'do-gooder derogation': individuals tend to react negatively to those who they perceive as promoting exemplary moral behaviour. Do-gooder derogation is in large part motivated by the desire to maintain a positive self-image in the face of potential moral criticism from others, which can be achieved by discounting the value of others and their opinions [61]. Circumstances tending to prompt do-gooder derogation include: (1) that the derogator specifically desires to act in a way that is contrary to moral advice [62, 63]; (2) that the derogator perceives the other as deliberately occupying a moral high ground [64]; and (3) that the derogator perceives the other as acting hypocritically [65]. We are unaware of any research which directly examines how these effects might apply in the case of decisions taken or advice given by moral machines. However, we note that machines designed to reflect progressive moral values may well be perceived as occupying the moral high ground by design, will offer advice which humans are tempted not to follow, and may (as machines) not be subject to their own prescriptions. It is therefore reasonable to predict that machines offering moral advice based on values more advanced than those of their advisees are subject to serious risk of do-gooder derogation.

The second reason to expect that humans will react negatively to machines programmed with progressive values is the phenomenon of 'outgroup derogation', the most relevant aspect of which involves individuals' discounting the attitudes of others whom they perceive as having different social identities to themselves [66, 67]. Non-human agents are in some sense the ultimate outgroup, and indeed humans display strong negative responses towards non-human agents when these appear or behave in ways that are similar but not indistinguishable from human appearances and behaviours (the uncanny valley phenomenon [68]). Interaction with non-human systems that act in some ways similarly to humans (by offering moral advice) may therefore activate processes of outgroup derogation. Outgroup derogation can be so strong that it creates a boomerang effect, whereby exposure to an argument from a mistrusted outgroup (e.g. one's political opponents) actually reinforces the subject's contrasting opinion [69, 70]. This effect is mediated by a sense of being threatened by proponents of the opposing opinion [71]. In this context, it is worth recalling that half of US citizens fear artificial intelligence [30].

Humans may be prone to a particularly strong sense of injustice when their actions are judged by machines. One study demonstrated the curious result that when humans judged mistakes made by artificial agents, perceived injustice was a stronger predictor of mistrust of the system than was perceived harm [72]. Compare also the recent UK public protests in response to A-level school grades' being generated by an algorithm (dubbed 'mutant' by the Prime Minister in response [73]).

Because machines are not human and are unlikely to have existences that humans can relate to in relevant contexts, they will be unable to engage in many of the methods of moral persuasion that, as we argued above, are a crucial mechanism of moral progress and may serve to mitigate against these kinds of backfiring reactions. Machines will be restricted to the least effective methods (such as offering sensible reasons for actions [74]). There is, therefore, little to suggest that morally progressive machines represent the solution to the concern we have raised about calcification of prevailing norms.

5 | ONE REASON WHY THE PROBLEM IS ACUTE: MORAL PSYCHOLOGY, THE ENVIRONMENT, AND THE URGENT NEED FOR MORAL PROGRESS

To this point, we have raised our concern in relatively general terms. Moral machines may hinder the mechanisms of moral progress, and that generates the danger of progress-inhibiting value calcification. We have also argued that the politics and psychology of human reactions to artificial intelligence precludes any simple solution to this problem.

This way of raising the concern may make the dangers we have identified seem abstract and remote. This impression is mistaken, however. We conclude by providing a concrete illustration of the problem that shows it to be an acute and urgent concern.

An overwhelming body of evidence indicates that humanity is currently doing enormous damage to the Earth's climate and biodiversity [3, 4]. If we do not alter our course, human activity is likely to bring about catastrophic increases in global mean surface temperatures and devastating losses of ecosystem services and species within centuries or even decades, implying, for example, serious impacts on global food supplies [75]. Moreover, we are running out of time to make the necessary changes. In 2015, in recognition of some of these risks, governments around the world agreed to try to limit global mean surface temperature increases to 1.5°C above pre-industrial levels. A 2018 report by the Intergovernmental Panel on Climate Change estimated that for even a 50% chance of achieving that target, global carbon emissions must be cut by around half within the next 10 years and to 'net zero' by 2050. Moreover, that figure assumes the availability soon of carbon capture and storage technologies that have yet to be shown to work at the necessary scale. As the IPCC claimed, this will require radical transformations in energy, industry, transport, building, agriculture, forestry, and land use. Meanwhile, as a result of our activities, the sixth mass extinction is already under way [76–78].

This parlous state indicates a clear need for the kind of moral progress captured by our minimal principle. Further, there are reasons to think the environmental crisis might be especially subject to the risks we have associated with machine morality. Some properties of human psychology make it particularly challenging to grasp the full implications of the crisis and thus also the extent of change necessary [79]. For example, problems that are seen as distant in time and space [80, 81] and not caused by malign intent [79, 82] tend not to be perceived as moral concerns requiring urgent attention. Collective decision making currently reflects these psychological biases, and it is therefore particularly likely that machine morality will reflect the inadequate status quo.

Many populations do now espouse high levels of environmental concern—for example most Europeans explicitly claim to place a very high value on caring for the environment, caring no less than they claim to care for other people [83]. However, the lack of action (in proportion to the crisis) indicates a disconnect between explicitly espoused values and more strongly internalised implicit values that drive behaviour [84–86]. In part because of this potentially disconcerting value-action gap, individuals are particularly prone to token environmental behaviours, whereby they feel that because they have done 'something', they feel they have done 'enough' [54]. This frequently causes actions that feel environmentally good to have an overall negative effect because they licence compensatory negative behaviour [55, 87]. This implies that consultation with machines said to be environmentally moral, but which in fact implement a status quo or insufficiently progressive morality, could be particularly counterproductive in reinforcing a false sense of moral adequacy. The prevalence of corporate greenwashing [88] already demonstrates how companies target consumers' implicit-explicit value gap, with marketing targeting explicit environmental values but products themselves targeting

better-internalised desires for low-cost products with little consideration for the environment [89]. It seems reasonable to assume that some commercial moral machines might do the same.

Outside the mainstream cultures of developed Western societies, environmentalist values are sometimes much more deeply embedded within cultural norms [90]. Studies in the field of cultural evolution illustrate that cultural change to reduce environmental impact is psychologically realistic: evolved cultural practices can maintain environmental impact at sustainable levels [91], and rapid moral evolution is, in principle, possible [92–94]. Environmentally progressive values might also be encoded in machines, of course. We note, however, that because the environment is an area where progressive values are already often espoused, it is well known that there are strong psychological defence mechanisms that operate to ignore or derogate such values when they require sacrificial action, even in cases where the individuals themselves espouse the values [46, 47, 95]. Only a limited set of progressive message types can bypass these defence mechanisms to induce moral progress through properly internalised values—for example, role-modelling by prominent ingroup members at a grass-roots level [96]. As outlined above, these are the types of message which machines would be ill-equipped to deliver.

In summary, evidence indicates that although the environmental values which most people display are generally insufficient to motivate necessary environmentalist behaviour [97], this could change. A shift away from egoistic individual rights-based norms, towards more strongly internalised environmentalist duty-based norms, might well result in increased environmentalist behaviour and therefore represent much-needed moral progress. However, the risks of stagnation or backlash associated with machine morality indicate that it may be particularly ill-suited to this task.

6 | CONCLUSION

We have argued that powerful incentives structure the choices of engineers concerning which moral values the machines they build should reflect, as a result of which the machines they build will tend to reflect prevailing moral values. We also argued that mechanisms of moral progress are likely to be inhibited by moral machines. The upshot of these two arguments is that moral machines present a real risk of calcifying prevailing moral values. Moreover, the risk we have identified is not a mere abstract, theoretical possibility. On the contrary, as humans face the catastrophic dangers of climate change and ecological destruction as a result of their ongoing engagement in activities that they take to be licenced by prevailing moral norms, the risk of value calcification associated with moral machines is clear, acute, and urgent.

However, we also acknowledge that it is difficult to assess how great this risk is, because there are currently few machines making moral decisions, and thus little directly relevant data. Although there is no clear reason to believe that encoding the

status quo could assist with moral progress, and reasons to believe moral machines will indeed encode the status quo, we grant that there are plausible arguments why moral progress might sometimes be assisted by moral machines that are progressive [14–18]. Further, we agree that artificial intelligence in general might prove crucial in supporting decisions of moral relevance, including decisions regarding the environment [98, 99]. However, in line with previous critics of machine ethics [100–103], we suggest that automatic systems can be constructed for human decision support without computing moral values.

Our aim has been to identify and articulate a risk that has hitherto gone largely unrecognised. For reasons of space and the limits of available evidence, we have not offered an assessment as to the gravity of the risk—either in general or in the present moment—or proposed solutions to it. However, we believe we have said enough to warrant the inclusion of this risk alongside other, more familiar risks associated with AI and related technologies in the thinking of those who consider the creation of moral machines.

ORCID

Ben Kenward  <https://orcid.org/0000-0002-8472-7172>

Thomas Sinclair  <https://orcid.org/0000-0002-9194-8555>

REFERENCES

- Danaher, J.: The rise of the robots and the crisis of moral patiency. *AI Soc.* 34(1), 129–136 (2019)
- Cave, S., et al.: Motivations and risks of machine ethics. *Proc. IEEE.* 107(3), 562–574 (2019)
- IPCC: Summary for policymakers. In: Masson-Delmotte, V., et al. (eds.) *Global warming of 15°C An IPCC Special Report on the impacts of global warming of 15°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty*, pp. 32. World Meteorological Organization, Geneva (2018)
- IPBES: Summary: For policymakers. In: Díaz, S., et al. (eds.) *The Global Assessment Report on Biodiversity and Ecosystem Services*, pp. 56. IPBES Secretariat, Bonn (2019)
- Gert, B., Gert, J.: The definition of morality. In: Zalta, E.N. (ed.) *The Stanford Encyclopaedia of Philosophy*. Fall 2020 Edition. Stanford University (2020)
- Awad, E., et al.: The moral machine experiment. *Nature.* 563(7729), 59–64 (2018)
- Simpson, T.W., Müller, V.C.: Just war and robots killings. *Phil. Q.* 66(263), 302–322 (2016)
- Malle, B.F.: Integrating robot ethics and machine morality: the study and design of moral competence in robots. *Ethics Inf. Technol.* 18, 243–256 (2016)
- IEEE: The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. Available from <https://standards.ieee.org/content/ieee-standards/en/industry-connections/ec/autonomous-systems.html> (2019). Accessed 04 Nov 2020
- Bryson, J.J.: Robots should be slaves. In: Wilks, Y. (ed.) *Close Engagements with Artificial Companions*, pp. 63–74. John Benjamins, Amsterdam (2010)
- Cervantes, J.-A., et al.: Artificial moral agents: a survey of the current status. *Sci. Eng. Ethics.* 26(2), 501–532 (2020)
- Lin, Z.J., et al.: The limits of human predictions of recidivism. *Sci. Adv.* 6(7), eaaz0652 (2020)
- Kehl, D.L., Kessler, S.A.: Algorithms in the criminal justice system: Assessing the use of risk assessments in sentencing. *Responsive Communities Initiative, Berkman Klein Center for Internet & Society, Harvard Law School*. Available from <http://nrs.harvard.edu/urn-3:HUL.InstRepos:33746041> (2017)
- Savulescu, J., Maslen, H.: Moral enhancement and artificial intelligence: moral AI? In: *Beyond Artificial Intelligence*, pp. 79–95. Springer, New York (2015)
- Klincewicz, M.: Artificial intelligence as a means to moral enhancement. *Stud. Log. Gramm. Rhetor.* 48(1), 171–187 (2016)
- Giubilini, A., Savulescu, J.: The artificial moral advisor. The ideal observer meets artificial intelligence. *Philos. Technol.* 31(2), 169–188 (2018)
- Anderson, M., Anderson, S.L.: Introduction to Part V. In: Anderson, M., Anderson, S.L. (eds.) *Machine Ethics*, pp. 495–498. Cambridge University Press, Cambridge (2011)
- Seville, H., Field, D.G.: What can AI do for ethics?. *The Society for the Study of Artificial Intelligence and the Simulation of Behaviour*, Birmingham (2000)
- Gordon, J.-S.: Building moral robots: ethical pitfalls and challenges. *Sci. Eng. Ethics.* 26(1), 141–157 (2020)
- Santos-Lang, C.: Our responsibility to manage evaluative diversity. *SIGCAS Comput. Soc.* 44(2), 16–19 (2014)
- Wong, P.-H.: Cultural differences as excuses? Human rights and cultural values in global ethics and governance of AI. *Philos. Technol.* 33(4), 705–715 (2020)
- Ecoffet, A., Lehman, J.: Reinforcement Learning Under Moral Uncertainty. arXiv:200604734 preprint (2020)
- Deacon, R.: Human rights as imperialism. *Theoria.* 50(102), 126–138 (2003)
- Hagendorff, T.: The ethics of AI ethics: an evaluation of guidelines. *Minds Mach.* 30(1), 99–120 (2020)
- IBM: Building trust in AI. Available from <https://www.ibm.com/watson/advantage-reports/future-of-artificial-intelligence/building-trust-in-ai.html>, (2019)
- Glikson, E., Woolley, A.W.: Human trust in artificial intelligence: review of empirical research. *Annals.* 14(2), 627–660 (2020)
- Yokoi, R., Nakayachi, K.: Trust in autonomous cars: exploring the role of shared moral values, reasoning, and emotion in safety-critical decisions. *Hum. Factors*, 0018720820933041 (2020)
- Malle, B.F., et al.: (eds.): Which robot am I thinking about? The impact of action and appearance on people's evaluations of a moral robot. In: *Proceedings of the 2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 7–10 March IEEE Press, Christchurch, New Zealand (2016)
- Banks, J.: Good robots, bad robots: morally valenced behaviour effects on perceived mind, morality, and trust. *Int. J. Soc. Robot.* (2020)
- Blumberg Capital. Artificial Intelligence in 2019: Getting past the adoption tipping point 2019 Available from <https://www.blumbergcapital.com/ai-in-2019/>, (2019). Accessed 09 Nov 20
- Criado, N., Such, J.M.: Digital discrimination. In: Yeung, K., Lodge, M. (eds.) *Algorithmic Regulation*. Oxford University Press, Oxford (2019)
- Caliskan, A., Bryson, J.J., Narayanan, A.: Semantics derived automatically from language corpora contain human-like biases. *Science.* 356(6334), 183–186 (2017)
- Saka, E.: Big Data and Gender-biased Algorithms. *The International Encyclopedia of Gender, Media, and Communication*, pp. 1–4. Wiley, Hoboken (2020)
- Smart, J.J.C.: An outline of a utilitarian system of ethics. In: Smart, J.J.C., Williams, B. (eds.) *Utilitarianism: For and Against*, pp. 1–74. Cambridge University Press, Cambridge (1973)
- Sumner, L.W.: Mill's theory of rights. In: West, H.R. (ed.) *The Blackwell Guide to Mill's Utilitarianism*. Blackwell, Hoboken (2006)
- Lyons, D.: Mill's theory of justice. In: Goldman, A.I., Kim, J. (eds.) *Values and Morals: Essays in Honour of William Frankena, Charles Stevenson, and Richard Brandt*. Springer, New York (1978)
- Ivry, T., Teman, E.: Shouldering moral responsibility: the division of moral labour among pregnant women, rabbis, and doctors. *Am. Anthropol.* 121(4), 857–869 (2019)

38. Keshet, Y., Liberman, I.: Coping with illness and threat: why non-religious Jews choose to consult rabbis on healthcare issues. *J. Relig. Health.* 53(4), 1146–1160 (2014)
39. Reason, J.: *Human Error*. Cambridge University Press, Cambridge (1990)
40. Harford, T.: *Crash: How Computers are Setting Us up for Disaster*. The Guardian (2016)
41. Beduschi, A.: International migration management in the age of artificial intelligence. *Migration Studies* (2020). <https://doi.org/10.1093/migration/mnaa003>
42. Molnar, P., Gill, L.: *Bots at the Gate: A Human Rights Analysis of Automated Decision-making in Canada's Immigration and Refugee System*. University of Toronto, Toronto (2018)
43. Wacquant, L.: *Habitus*. In: Becker, J., Zafirovski, M. (eds.) *International Encyclopaedia of Economic Sociology*, pp. 315–319. Routledge, London (2004)
44. Deakin, S.: Evolution for our time: a theory of legal memetics. *Curr. Leg. Probl.* 55(1), 1–42 (2002)
45. Laurin, K.: Inaugurating rationalisation: three field studies find increased rationalisation when anticipated realities become current. *Psychol. Sci.* 29(4), 483–495 (2018)
46. Norgaard, K.M.: People want to protect themselves a little bit: emotions, denial, and social movement nonparticipation. *Socio. Inq.* 76(3), 372–396 (2006)
47. Adams, M.: *Ecological Crisis, Sustainability and the Psychosocial Subject*. Palgrave Macmillan UK, London (2016)
48. FeldmanHall, O., et al.: What we say and what we do: the relationship between real and hypothetical moral choices. *Cognition.* 123(3), 434–441 (2012)
49. Haidt, J.: Moral psychology for the twenty-first century. *J. Moral Educ.* 42(3), 281–297 (2013)
50. Cameron, C.D., et al.: Implicit moral evaluations: a multinomial modelling approach. *Cognition.* 158, 224–241 (2017)
51. Hofmann, W., Baumert, A.: Immediate affect as a basis for intuitive moral judgement: an adaptation of the affect misattribution procedure. *Cognit. Emot.* 24(3), 522–535 (2010)
52. Jost, J.T.: The IAT is dead, long live the IAT: context-sensitive measures of implicit attitudes are indispensable to social and political psychology. *Curr. Dir. Psychol. Sci.* 28(1), 10–19 (2018)
53. Kurdi, B., et al.: Relationship between the implicit association test and intergroup behaviour: a meta-analysis. *Am. Psychol.* 74(5), 569–586 (2019)
54. Sachdeva, S., Jordan, J., Mazar, N.: Green consumerism: moral motivations to a sustainable future. *Curr. Opin. Psychol.* 6, 60–65 (2015)
55. Sörqvist, P., Langeborg, L.: Why people harm the environment although they try to treat it well: an evolutionary-cognitive perspective on climate compensation. *Front Psychol.* 10, 348 (2019)
56. Moody-Adams, M.M.: Moral progress and human agency. *Ethical Theory & Moral Pract.* 20(1), 153–168 (2017)
57. Anderson, E.: *Social movements, experiments in living, and moral Progress: case studies from Britain's abolition of slavery*. The Lindley Lecture, University of Kansas, Department of Philosophy Lawrence, Kansas (2014)
58. Jamieson, D.: Slavery, carbon, and moral progress. *Ethical Theory & Moral Pract.* 20(1), 169–183 (2017)
59. Perloff, R.M.: *The Dynamics of Persuasion: Communication and Attitudes in the Twenty-first Century*. Routledge (2020)
60. Smithey, L.A.: Social movement strategy, tactics, and collective identity. *Sociol. Compass.* 3(4), 658–671 (2009)
61. Minson, J.A., Monin, B.: Do-gooder derogation. *Soc. Psychol. Personal. Sci.* 3(2), 200–207 (2011)
62. Bolderdijk, J.W., Brouwer, C., Cornelissen, G.: When do morally motivated innovators elicit inspiration instead of irritation? *Front. Psychol.* 8, 2362 (2018)
63. Zane, D.M., Irwin, J.R., Reczek, R.W.: Do less ethical consumers denigrate more ethical consumers? The effect of willful ignorance on judgements of others. *J. Consum. Psychol.* 26(3), 337–349 (2016)
64. Kurz, T., et al.: Could vegans and lycra cyclists be bad for the planet? Theorising the role of moralised minority practice identities in processes of societal-level change. *J. Soc. Issues.* 76(1), 86–100 (2020)
65. Sparkman, G., Attari, S.Z.: Credibility, communication, and climate change: how lifestyle inconsistency and do-gooder derogation impact decarbonisation advocacy. *Energy Res. Soc. Sci.* 59, 101290 (2020)
66. Hogg, M.A., Smith, J.R.: Attitudes in social context: a social identity perspective. *Eur. Rev. Soc. Psychol.* 18(1), 89–131 (2007)
67. Mackie, D.M., Gastardo-Conaco, M.C., Skelly, J.J.: Knowledge of the advocated position and the processing of in-group and out-group persuasive messages. *Pers. Soc. Psychol. Bull.* 18(2), 145–151 (1992)
68. Syrdal, D.S., et al.: The negative attitudes towards robots scale and reactions to robot behaviour in a live human-robot interaction study. *Adaptive and emergent behaviour and complex systems: Proc. of the 23rd Convention of the Society for the Study of Artificial Intelligence and Simulation of Behaviour, AISB 2009*, pp. 109–115. SSAISB (2009)
69. Hart, P.S., Nisbet, E.C.: Boomerang effects in science communication. *Commun. Res.* 39(6), 701–723 (2011)
70. Ma, Y., Dixon, G., Hmielowski, J.D.: Psychological reactance from reading basic facts on climate change: the role of prior views and political identification. *Environ. Commun.* 13(1), 71–86 (2019)
71. Hoffarth, M.R., Hodson, G.: Green on the outside, red on the inside: perceived environmentalist threat as a factor explaining political polarization of climate change. *J. Environ. Psychol.* 45, 40–49 (2016)
72. Sullivan, Y., de Bourmont, M., Dunaway, M.: Appraisals of harms and injustice trigger an eerie feeling that decreases trust in artificial intelligence systems. *Ann. Oper. Res.* (2020)
73. Paulden, T.: A cutting re-mark. *Significance.* 17(5), 4–5 (2020)
74. Blair, J.A.: The persuasive ineffectiveness of arguing and arguments. In: *OSSA Conference Archive*, vol. 10, Windsor (2020)
75. FAO: The state of the world's biodiversity for food and agriculture. In: Bélanger, J., Pilling, D. (eds.) *FAO Commission on Genetic Resources for Food and Agriculture Assessments*, Rome. <http://www.fao.org/3/CA3129EN/CA3129EN.pdf> (2019)
76. Ceballos, G., Ehrlich, P.R., Dirzo, R.: Biological annihilation via the ongoing sixth mass extinction signalled by vertebrate population losses and declines. *Proc. Natl. Acad. Sci. USA.* 114(30), E6089–E6096 (2017)
77. Ceballos, G., et al.: Accelerated modern human-induced species losses: entering the sixth mass extinction. *Science advances.* 1(5), e1400253 (2015)
78. Barnosky, A.D., et al.: Has the Earth's sixth mass extinction already arrived? *Nature.* 471(7336), 51–57 (2011)
79. Markowitz, E.M., Shariff, A.F.: Climate change and moral judgement. *Nat. Clim. Change* 2(4), 243–247 (2012)
80. Singh, A.S., et al.: The perceived psychological distance of climate change impacts and its influence on support for adaptation policy. *Environ. Sci. Pol.* 73, 93–99 (2017)
81. Spence, A., Poortinga, W., Pidgeon, N.: The psychological distance of climate change. *Risk Anal.* 32(6), 957–972 (2012)
82. Greene, J.D., et al.: Pushing moral buttons: the interaction between personal force and intention in moral judgement. *Cognition.* 111(3), 364–371 (2009)
83. Bouman, T., Steg, L.: Motivating society-wide pro-environmental change. *One Earth.* 1(1), 27–30 (2019)
84. Farjam, M., Nikolaychuk, O., Bravo, G.: Experimental evidence of an environmental attitude-behaviour gap in high-cost situations. *Ecol. Econ.* 166, 106434 (2019)
85. Beattie, G., McGuire, L.: Consumption and climate change: why we say one thing but do another in the face of our greatest threat. *Semiotica.* 2016(213), 493–538 (2016)
86. Thomas, G.O., Walker, I.: The development and validation of an implicit measure based on biospheric values. *Environ. Behav.* 48(5), 659–685 (2014)
87. Maki, A., et al.: Meta-analysis of pro-environmental behaviour spillover. *Nat. Sustain.* 2(4), 307–315 (2019)
88. Lyon, T.P., Montgomery, A.W.: The means and end of greenwash. *Organ. Environ.* 28(2), 223–249 (2015)

89. Van't Veld, K.: Eco-labels: modelling the consumer side. *Annu. Rev. Resour. Econ.* 12(1), 187–207 (2020)
90. Gratani, M., et al.: Indigenous environmental values as human values. *Cogent Soc. Sci.* 2(1), 1185811 (2016)
91. Brooks, J.S., et al.: Applying cultural evolution to sustainability challenges: an introduction to the special issue. *Sustain. Sci.* 13(1), 1–8 (2018)
92. Lindström, B., et al.: The role of a common is moral heuristic in the stability and change of moral norms. *J. Exp. Psychol. Gen.* 147(2), 228–242 (2018)
93. Fogarty, L., Kandler, A.: The fundamentals of cultural adaptation: implications for human adaptation. *Sci. Rep.* 10(1), 14318 (2020)
94. Wheeler, M.A., McGrath, M.J., Haslam, N.: Twentieth century morality: the rise and fall of moral concepts from 1900 to 2007. *PLoS One.* 14(2), e0212267 (2019)
95. Lamb, W.F., et al.: Discourses of climate delay. *Glob. Sustain.* 3, e17 (2020)
96. Grabs, J., et al.: Understanding role models for change: a multilevel analysis of success factors of grassroots initiatives for sustainable consumption. *J. Clean Prod.* 134, 98–111 (2016)
97. Persson, I., Savulescu, J.: *Unfit for the Future: The Need for Moral Enhancement*. OUP, Oxford (2012)
98. Vinuesa, R., et al.: The role of artificial intelligence in achieving the Sustainable Development Goals. *Nat. Commun.* 11(1), 233 (2020)
99. Rolnick, D., et al.: Tackling climate change with machine learning. *arXiv:190605433 preprint* (2019)
100. Van Wynsberghe, A., Robbins, S.: Critiquing the reasons for making artificial moral agents. *Sci. Eng. Ethics.* 25(3), 719–735 (2019)
101. Bryson, J.J.: Patience is not a virtue: the design of intelligent systems and systems of ethics. *Ethics Inf. Technol.* 20(1), 15–26 (2018)
102. Ryan, M.: In AI we trust: ethics, artificial intelligence, and reliability. *Sci. Eng. Ethics.* 26(5), 2749–2767 (2020)
103. Nyrup, R., Whittlestone, J.: Why Value Judgements Should Not Be Automated (2019). <https://doi.org/10.17863/CAM.41552>

How to cite this article: Kenward, B., Sinclair, T.: Machine morality, moral progress, and the looming environmental disaster. *Cogn. Comput. Syst.* 1–8 (2021). <https://doi.org/10.1049/ccs2.12027>