

Chapter 18

Against Human Exceptionalism: Environmental Ethics and the Machine Question



Migle Laukyte

Abstract This paper offers an approach for addressing the question of how to deal with artificially intelligent entities, such as robots, mindclones, androids, or any other entity having human features. I argue that to this end we can draw on the insights offered by environmental ethics, suggesting that artificially intelligent entities ought to be considered not as entities that are extraneous to the human social environment, but as forming an integral part of that environment. In making this argument I take a radical strand of environmental ethics, namely, Deep Ecology, which sees all entities as existing in an inter-relational environment: I thus reject any “firm ontological divide in the field of existence” (Fox W, Deep ecology: A new philosophy of our time? In: Light A, Rolston III H (eds) *Environmental ethics: An anthology* Blackwell, Oxford, 252–261, 2003) and on that basis I introduce principles of biospherical egalitarianism, diversity, and symbiosis (Naess A, *Inquiry* 16(1):95–100, 1973). Environmental ethics makes the case that humans ought to “include within the realms of recognition and respect the previously marginalized and oppressed” ((Gottlieb RS, Introduction. In: Merchant C (ed) *Ecology. Humanity Books, Amherst*, pp ix–xi, 1999)). I thus consider (a) whether artificially intelligent entities can be described along these lines, as somehow “marginalized” or “oppressed,” (b) whether there are grounds for extending to them the kind of recognition that such a description would seem to call for, and (c) whether Deep Ecology could reasonably be interpreted in such a way that it apply to artificially intelligent entities.

Keywords Moral responsibility · Environmental ethics · Deep ecology · Artificial intelligence · Artificial agency

M. Laukyte (✉)

Department of Private Law, Universidad Carlos III de Madrid, Madrid, Spain
e-mail: migle.laukyte@uc3m.es

18.1 Introduction

Time and again we have raised the question of how in the future we ought to treat artificially intelligent entities, such as robots, mindclones, androids, bemans, or any other entity having intelligence, autonomy, or other features that would make it similar to a human being.¹ Furthermore, with human enhancement, and with the prospect of technologies like those that try to build robots with a biological brain grown in an incubator or to upload a human brain onto a computer (Kurzweil 2006; Rothblatt 2014), it is no longer ontologically clear what it is to be human or how we should draw the line between human and nonhuman, and where we should place transhumans, namely, individuals who “transcend human biological inheritance, modifying their DNA, their bodies, or the substrate for their minds” (Rothblatt 2014, 307).

These technological scenarios confront us with the ethical problem of inclusion and exclusion: Are the new entities worthy of consideration as moral beings? And, if so, on what basis? Depending on the way we answer these questions, we will come out with different ways of treating these new entities, thus fundamentally shaping the social environment in which we are going to live in the future and which we are going to pass on to the future generations.²

This paper offers an approach to the question of how to deal with artificially intelligent entities: I propose that we draw on the insights offered by environmental ethics, suggesting that artificially intelligent entities ought to be considered not as entities extraneous to our social environment, but as forming an integral part of that environment. The argument I will be unpacking builds on the radical strand of environmental ethics known as Deep Ecology,³ whose underlying premise is that all entities exist in an inter-relational environment: Deep Ecology thus rejects any “firm ontological divide in the field of existence” (Fox 2003, 255), and on that basis it introduces principles of biospherical egalitarianism, diversity, and symbiosis (Naess 1973). Environmental ethics makes the case that we humans ought to “include within the realms of recognition and respect the previously marginalized and oppressed” (Gottlieb 1999, ix), so in this paper I consider whether (a) artificially intelligent entities can be described along these lines, as somehow “marginalized” or “oppressed”; (b) whether there are grounds for extending to them the kind of

¹A mindclone is a cyberversion of a human being, with a human mind uploaded on a digital support, whereas a beman is not a replication of human mind but an entity that is cyberconscious on its own account. On mindclones, bemans, and other possibilities offered by artificial intelligence, see Rothblatt (2014). In the interest of clarity, I will use the term *human* to refer to human beings, *nonhuman* to refer to all other living and nonliving entities (animals, mountains, rocks, machines), and *artificial* to refer to artificially intelligent entities and other artificial forms of life.

²On the moral treatment of artificial agents and its justification on different grounds, such as rationality, interactivity, and autonomy, see Floridi and Sanders (2004), Tavani (2011) and Coeckelbergh (2009 and 2010).

³Other radical theories are Social Ecology, Political Ecology, and Ecofeminism (Keulartz 1995).

recognition that such a description would seem to call for; and (c) whether the ideas of Deep Ecology can be applied to artificially intelligent entities.

The discussion is organized as follows: In Sect. 17.1, some of the key notions, related to artificially intelligent entities, environmental ethics, and Deep Ecology, are explained. In Sect. 17.2, the focus is on why and how the ideas of Deep Ecology could apply to artificial intelligence, focusing in particular on some of the eight principles of Deep Ecology: the argument is that these principles are applicable not only to biological entities and the biosphere in general, but also to artificially intelligent entities. In Sect. 17.3, the focus shifts to the main difficulty with the idea of bringing Deep Ecology to bear on artificially intelligent entities. This is the idea that nature—or the environment at large—is a breathing and evolving organism made of living sentient entities (such as animals, fish, and plants) and as such is thus worthy of moral consideration. The difficulty is that this description—namely, being alive or sentient—is usually not attributed to artificially intelligent entities. This critical point is addressed by offering a way out of the impasse, arguing that being alive and sentient are not essential requirements for moral consideration, while pointing out alternative approaches that have been developed in that regard, so much so that even Deep Ecologists themselves as well as many environmental ethicists agree that landscapes and mountains, for example, are also worthy of moral consideration. Having addressed those issues, the paper finishes with a few closing remarks.

18.2 Some Notes on Terminology: Who's Who?

Before taking up the arguments for and against extending Deep Ecology from the natural environment to artificially intelligent entities—from the natural world to the artificial world, thus providing the concept of the environment with a new and more inclusive meaning we need to make some clarifications about the terminology used in this paper.

I begin with the idea of an artificially intelligent entity. As suggested earlier, artificially intelligent entities are any kind of entity having an artificially built intelligence and other features associated with intelligence, such as autonomy and the ability to make reasoned decisions. This artificial intelligence I regard as similar to human intelligence: It may outsmart human beings in some respects (Bostrom 2014), while falling short in others. The point is not to rank different forms of intelligence on any scale of excellence: It is rather to determine whether they have the kinds of features that would trigger the question of moral consideration. This means that it does not matter how this intelligence is achieved: It can be via human whole-brain emulation or by uploading a human brain onto a digital device (Rothblatt 2014)—or indeed any embodiment of artificial intelligence, be it digital, virtual, or physical—so long as it resembles human intelligence. This importantly means that I regard it as essential to artificially intelligent entities that they have social or interactive capacities, and that in light of that behaviour we can ascribe some kind of emotion or other to them. The artificially intelligent entities taken

into account here can thus be described as embodying a general-purpose artificial intelligence, an intelligence that is not related to any particular task but applies across different environments and contexts and to a range of different problems (just like human intelligence).

A problem comes up in regard to the sentience of such entities: are they sentient or nonsentient? This is the criterion by which we usually determine whether we have a duty of ethical consideration: If the entity is sentient, we owe some consideration; if it is not, we can exploit it in any way that we think is beneficial to us (this I will call the anthropocentric view). There are many reasons why this way of thinking is wrong, but I will focus on two of them. For one thing, Deep Ecologists already acknowledge that nonliving entities, such as mountains, are inherently worthy of moral consideration, so it is beside the point whether or not artificially intelligent entities are sentient. And, for another, by developing artificially intelligent entities, we might also develop a different and new kind of sentience that will challenge our idea of what sentience and nonsentience are.

Let us turn now to environmental ethics: This is the branch of ethics that focuses on the interaction between humans and “nonhuman nature within the context of ecological systems” (Keller 2010, 3). It branches into several subareas, but what links them all together is the juxtaposition of, and competition between, two values that are attributed to nature, namely, its instrumental value (nature as a means to an end) and its intrinsic value (nature as an end in itself).⁴

We can now consider Deep Ecology.⁵ This is a field of environmental ethics that departs from mainstream environmentalism by moving away from the previously mentioned anthropocentric view on which nature is worthy of protection only insofar as that is instrumental to human welfare. Deep Ecology, by contrast, envisions a deeper way of dealing with environmental issues, not only from a philosophical perspective but also from a political one. It does so by looking at nature as valuable in itself, regardless of whether it is useful to human beings: It thus assigns intrinsic value to ecosystems (Baard 2015). This view has been termed biospheric egalitarianism. And the reason why it describes itself as deep is that, in reframing our understanding of nature, it calls on us to fundamentally change our way of relating to it, not as a means to an end but as an end in itself. This can be achieved by “engaging in a process of ever-widening identification with others” (Keulartz 1995, 118), an identification which is not limited to other human beings but extends to the entire biosphere, and which would therefore be impossible without “a more sensitive openness to ourselves and nonhuman life around us” (Devall and Sessions 1985, 65).⁶

⁴Environmental ethics quite often deals with the juxtaposition between anthropomorphism and nonanthropomorphism, holism, individualism, and other ideas (Keulartz 1995).

⁵The seminal study on Deep Ecology is Naess (1973), and the view has since been developed in numerous works. For an overview, see Naess (1986, [1989] 2001, 2005, 2008) and Keller (2008, 206–11).

⁶In making this identification, however, Deep Ecology does not argue that all beings enjoy the same moral standing. Naess himself concedes that a ranking of beings is inevitable, pointing out that this

Deep Ecology is based on eight principles premised on the idea that humans no longer form the centre of discourse. This does not amount to removing the human being from the spectrum of moral consideration, but it does mean that since the human being is deeply intertwined with nature, the two components of this relation—namely, humans and nature—are to be regarded as forming a whole rather than as separate entities.

Let us see, then, what these eight principles of Deep Ecology are:

1. The well-being and flourishing of human and nonhuman life on Earth are valuable in themselves, from which it follows that the value of nonhuman life-forms is not a function of its usefulness to humans.
2. The flourishing of human and nonhuman life is dependent on the richness and diversity of life forms, and this diversity is itself inherently valuable.
3. The inherent value of the richness and diversity human and nonhuman life means that humans do not have a right to reduce such richness and diversity, except to satisfy vital needs.
4. Human life and cultures can flourish even with a substantially smaller human population.
5. Human interference with the nonhuman world is excessive.
6. We must therefore make structural economic, technological, and policy changes.
7. The underlying change will have to be ideological: a change in attitude that consists in appreciating the quality of life itself rather than aiming for an increasingly higher standard of living as measured by economic growth.
8. The foregoing principles entail a duty to join together in an effort to implement the necessary changes (Naess 1986, 2; 2008, 111–12).

But before explaining how these principles could be applied to artificially intelligent entities (in Sect. 17.2), we still have a more fundamental question to address, namely, why choose environmental ethics, and Deep Ecology in particular, as a basis for reasoning about artificially intelligent entities?

Let us first consider three main reasons for framing the discussion on the basis of environmental ethics. The first reason is that we do not yet have any sufficiently broad ethics for artificial intelligence: Azimov's three laws of robotics cannot help us solve this problem, so we need a more solid ground on which to build an ethical approach to artificial intelligence. And the second reason is that we want to avoid the errors we made in the past in framing an ethical approach to nonhuman entities. This suggests looking for moral guidance outside the realm of specific disciplines, such as the philosophy of technology or the philosophy of artificial intelligence. Environmental ethics makes it possible to give broad scope to the question of the moral consideration of nature and the environment, and to do so in such a way as to address the paradigmatic shift we confront in the human approach to the ecosystem, in that the "biosphere [. . .] has become a human trust and has something of a moral

also entails a raking of duties: The duties we owe to fellow human beings are higher than those we owe to other beings, such as mice. This is an easy choice—humans versus mice—but there are choices that are neither clear nor easy, especially when it comes to ranking different species. This is why Naess (2005) describes ranking a complex process, and not straightforwardly moral.

claim on us” (Jonas 1984, 8). There is, finally, a third reason why environmental ethics seem promising as an approach to the ethics of artificial intelligence: This relatively young and dynamic field of moral inquiry takes in different ideas from the other moral theories,⁷ and in so doing it offers some insights that can be helpful in dealing with nonhuman otherness.

The idea of drawing on environmental ethics to address the moral problem of artificial intelligence is not new. Gunkel, for example, points out that environmental philosophy, animal rights philosophy, and the machine question all seek “to think outside the restrictions of anthropocentric privilege and human exceptionalism” and consequently “to dissolve the kind of human centric view of the universe that is being broken open by what we can say is a Copernican Revolution” (quoted in Kellogg 2014). He turns to environmental philosophy because in it he finds “a thinking of otherness that is no longer tied to either human centrism or biocentrism” (ibid.).⁸

But why Deep Ecology in particular? I will point out four reasons. First, seeing the human-centric (or anthropocentric) approach to environmental ethics as problematic, Deep Ecology takes an ecocentric approach: Instead of placing the human being at the centre of the discussion, it places humans next to other (biological) entities. Accordingly, Deep Ecology rejects the position that regards “humans as isolated and fundamentally separate from the rest of Nature, as superior to, and in charge of, the rest of creation” (Devall and Sessions 1985, 65). This is a good starting point, because it enables us to address the issue of artificially intelligent entities without narrowly preselecting humans as the main lens through which to understand what is worthy of moral consideration.

Second, Deep Ecology does not confine itself to strictly philosophical inquiry but advocates a wider and more profound social change: It takes us from a purely theoretical discussion to a more practical level, asking us to consider the need for institutional and political change, and that is exactly what may be needed in dealing with artificially intelligent entities (Bostrom 2014; Rothblatt 2014; Kurzweil 2006).

Third, Deep Ecology proceeds not from normative prescriptions but from principles. Unlike prescriptions, principles are broad and flexible, making it possible to interpret and shape them in ways that will meet the demands of a discussion on artificially intelligent entities. And, as we will see in Sect. 17.2, a useful link can be established between Deep Ecology and artificial intelligence (see Coeckelbergh 2010).⁹

⁷Thus, for example, environmental ethics introduced the question of justice in the debate on environmental problems. On this development, see Armstrong (2012). The problem with traditional theories is that they are all anthropocentric and no longer adequate to deal with current problems (Troster 2008, 392), but that need not be the case, considering that Deep Ecology draws inspiration from well-established theories like those of Spinoza, Heidegger, and Whitehead (Keulartz 1995).

⁸In this connection, see Coeckelbergh (2010) and Rothblatt (2014), drawing parallels between robot ethics and animal ethics in a way that brings environmental ethics into the picture.

⁹The bearing that Deep Ecology has on artificial intelligence and computer ethics is also briefly mentioned in Floridi and Sanders (2001).

Finally, the fourth reason is that, unlike anthropocentrism, with its short-term vision of environmental problems, Deep Ecology takes the long view (Baard 2015): this is precisely what we need if we want to have any kind of discussion about the future of artificial intelligence and its place in our human and natural environment. If we are to work toward any kind of fruitful coexistence of human, natural, and artificial life-forms, we need to extend our view over the long stretch. Indeed, short-term thinking would make the discussion irrelevant from the start, considering that artificially intelligent entities of the kinds that would make these problems real have yet to be created.

In the following Sects. 17.2 and 17.3, I will introduce arguments for and against applying Deep Ecology to artificially intelligent entities, exploring the reasons why the principles of Deep Ecology could extend to artificially intelligent entities, and considering how the arguments against such an extension could be defeated.

18.3 Deep Ecology as an Approach to the Problem of Artificially Intelligent Entities

Let us consider the arguments in favour of applying the principles of Deep Ecology to some of the challenges that artificially intelligent entities may give rise to in the future. I will argue that the insights Deep Ecology offers in dealing with the environment and the moral status of nature can also shed light on the question of the moral, social, and political implications of artificial intelligence. I begin by pointing out that, while the object of discussion may differ (the environment and nature as against artificial intelligence), the problem is the same (exploitation) and so is the decision-making entity (the human being). I elaborate on this point by taking the eight principles of Deep Ecology and applying them to artificial intelligence so as to see whether these principles are applicable to something more inclusive than nature, the environment, and the ecosystem.

As we saw, the first principle invites us to consider human and nonhuman life—or, as Jonas (1984, 8) puts it, “extrahuman” life—as intrinsically valuable, regardless of its contribution to human welfare. I submit that this principle can be extended to artificially intelligent entities because they, too, can be seen as a form of nonhuman life. Note that *life*, in the term *human and nonhuman life*, is understood by Deep Ecology to include rivers and landscapes: these are “nonliving” (Devall and Sessions 1985), and so also nonhuman, forms of life. And if rivers and landscapes are considered in this way—as nonhuman forms of life—so can artificially intelligent entities, as artificial forms of life.¹⁰ So the question here is not What is valuable to

¹⁰Here is what one commentator has written on the prospect of artificial life: “Many agree it is only a matter of time before artificial life creates machines that are alive, are intelligent, reproduce their own kind, have their own purposes, set their own goals, and evolve autonomously. These machines will be as much a part of the natural world as features in the landscape or existing forms of life,

human life? but How can human and nonhuman life, including artificial or synthetic life, be made compatible, and indeed coherent, weaving into a single fabric the value that can be recognized as intrinsic to all?

The same applies to the second principle. This principle recognizes the richness and diversity of life-forms, and artificially intelligent entities can be counted as a life-form, however much artificial or synthetic. These life-forms are not yet known to us, but on a Deep Ecology approach they can be regarded as valuable in themselves, just like other life-forms.

On the third principle, human beings do not have a right to reduce the richness and diversity of life-forms except to satisfy their vital (existential) needs, and for no other reason, and on the fourth principle human society can flourish consistently with a smaller human population: if we extend these two principles to the moral question of our treatment of artificially intelligent entities, we can see that there seem to be no vital needs in virtue of which to justify reducing the diversity of an environment inclusive of these entities, nor do these entities seem to have much influence on the growth of the human population: If artificially intelligent entities contribute to the richness and diversity of life-forms, and if they bear little relation to population growth, they are protected under these two principles.

The fifth principle asks us to reduce human interference in the nonhuman world. This principle raises something of a paradox because, if on the one hand such human interference is in large part responsible for the environmental problem we face today, on the other hand humans need to keep interfering in the nonhuman world so as to deal with and solve that very problem. The paradox is solved, then, by looking at the “nature and extent of such inference” (Devall and Sessions 1985, 72): Not all interference is of the same kind or equally extensive. This principle is more difficult to apply to artificially intelligent entities than the other principles of Deep Ecology because, on the one hand, artificially intelligent entities are artificially created, and so anything they do is ipso facto artificial, but on the other hand, they are so inextricably bound up with their human makers, and the interaction is so close, that it is difficult to draw a neat line of separation.¹¹ I would therefore count this as the most problematic principle of Deep Ecology. And the problem is compounded by the fact that the boundary between beneficial and harmful human interference is blurred, such that, when dealing with artificial intelligence, we probably need to make case-by-case judgments.

The sixth principle calls for a structural change in policy, moving away from a *laissez faire* model of self-regulating production, consumption, and growth that does not concern itself with the problem of externalities—i.e., the social and environmental costs of free-market capitalism—toward an environmentally sustainable model

and their evolution will affect the course of existing forms of life. [...] machines might play an unprecedented role in the next major evolutionary transition, and the challenge here is to predict and explain this role. Machines may well be the central players in the transition, as will be the case if autonomously evolving machines get established in the natural world” (Bedau et al. 2000, 373).

¹¹The interaction between humans and artificially intelligent entities will be developed further in the next section.

that does take those costs into account. In this respect, too, artificially intelligent entities can be seen to play a dual role. For on the one hand, as a product of this free-market model, they are part of the problem, but on the other, as a technology, they can also be part of the solution, contributing to the paradigm shift toward sustainability.¹²

Closely bound up with the sixth principle is the seventh, positing an inherent quality of life that cannot be reduced to the standard of living, in which the attempt is to measure the quality of life by the amount of goods and services produced in the economy, that is, by the total output of the economy, or gross domestic product, a measure plagued by the problem that it counts any economic transaction as growth regardless of whether it is sustainable or unsustainable. On the seventh principle, the idea is that while the quality of life may be dependent on material quality (as measured by access to goods and services), it cannot be *equated* with this measure, nor can it be severed from the problem of the whole—the problem of what the acquisition of material quality entails for the good of the planet as an interconnected whole. We can see that this principle clearly applies to artificial intelligence: even if Naess warns us against this neophilia (Baard 2015), the technology can be used to improve the quality of life—by relieving humans of the burden of carrying out tasks that do not seem to have any inherent value—and it can do so in an environmentally sustainable way.

The eighth principle calls on us to implement the first seven. Naess ([1989] 2001, 26, 45) observes that this would require “a substantial reorientation of our whole civilization,” with “new criteria for progress, efficiency, and rational action,” and “new social forms for co-existence.” This rethinking of society and civilization is clearly open-ended and open to interpretation, and one can expect a good deal of disagreement over the practical details, but there is no doubt that in the solution we can fit the idea of our coexistence with artificially intelligent entities.

What we can appreciate from this rundown is that there is no principled reason why we should be prevented from applying the principles of Deep Ecology to the question of artificially intelligent entities and our treatment of them. On the contrary, these principles can be useful from the outset in framing the issue of artificially intelligent entities in a constructive way. The issue is not so much about these entities themselves as it is about us, how we ought to interact with them, and the place we should find for them. If there is an overarching principle that captures the whole of Deep Ecology, it is that we live in a holistic system of interdependent components that are valuable in themselves, and whose interaction is essential to the life of the system itself: although the standing paradigm of social organization based on the idea of the market economy as a self-regulating system seems consistent with that overarching principle, we have learned from experience that a literal application

¹²As a societal model on which to base our interaction with nature, sustainability is also advocated in Kortetmäki 2016. Even if sustainability is consistent with the policy changes called for under the sixth principle, Naess was critical of the idea as such, arguing that it is anthropocentric and therefore out of keeping with the holism by which Deep Ecology is underpinned (Baard 2015).

of this idea comes with costs that are both social and environmental. Deep Ecology offers a way to reconsider that idea in such a way as not to repeat the mistakes of the past, and artificial intelligence can certainly be part of that solution.

That, in a snapshot, is Deep Ecology in connection with the problem of artificially intelligent entities. But there are a couple of important arguments that work against this approach. Let us therefore see what these arguments are and how we might respond.

18.4 A Critique of Deep Ecology

In this section we will consider two arguments against the idea of drawing on Deep Ecology as an approach to the issue of how we ought to relate to artificially intelligent entities. The first is the argument that Deep Ecology is ideally suited to dealing with conscious, sentient, or living organisms as part of our social and legal environment, and that these are not characteristics we can use to describe artificially intelligent entities.

There are many commentators who have responded to this objection. Thus, Rothblatt (2014) draws a parallel between consciousness and the cyberconsciousness ascribable to our mindclones, arguing that different entities have different forms of consciousness, and that these differences are irrelevant to whether artificially intelligent entities can be regarded as worthy of moral consideration, and hence whether there are reasons for bringing them under the protection of the law. Developments in synthetic (artificial) life applications suggest that the same reasoning applies when noting the property of being alive and sentient as obstacles to moral consideration. Nor does any other biological sort of property seem to keep us from recognizing artificially intelligent entities as having a moral status (Floridi and Sanders 2004).

The argument against such moral recognition seems to take an approach that consists in checking off a list of properties acting as necessary conditions to be met, but some alternative approaches have been proposed. One example is Coeckelbergh (2010), suggesting that we attack the problem by focusing on the social relations between humans and robots, without having to look at the properties ascribable to robots or their ontological features. This relational approach suggests that we would treat industrial robots in one way and domestic robot assistants in another.

The same approach is proposed by Gunkel (Kellogg 2014), who points out the increasing social interactivity of robots. Nor is this approach confined to philosophical inquiry: Engineers, too, have recognized how important it is to take the social aspects into account in dealing with human interaction with artificially intelligent entities. Thus, Farshchi (2016) argues that we should switch from building human-machine interfaces to creating machine-human interfaces. Where does the difference lie? A human-machine interface is based on natural-language processing and speech recognition, while in building a machine-human interface we are focused on understanding people and their emotional states. Such emotion-reading interfaces would hugely contribute to building social relations between

machines and human beings. Examples are JIBO, the world's first social home robot, which "communicates and expresses using natural, social and emotive cues"¹³; Pepper, whose "number one quality is his ability to perceive emotions" and adapt accordingly via an Emotion Engine¹⁴; and CoBot robots, based on symbiotic human-machine interaction.¹⁵ All these examples show that research in robotics is appreciating the crucial role of machine social skills, which are indispensable if we are to achieve a deeper interaction with greater empathy between human beings and machines.

The relational approach—as opposed to the argument that moral consideration necessarily requires sentience and biological life—also finds support in the work of ecologists such as Kortetmäki (2016, 92), who takes the example of the lakewater pollution caused by the Talvivaara nickel mine in Finland: In making the case that lakes are worthy of moral consideration, she does not stress that they are valuable as ecosystems on which we depend, or that the environmental damage done to them is detrimental to our health or kills a variety of living organisms, but rather argues that the issue is about "the lakes themselves as places to which the people have special relations." The same special relationship humans establish not only with places but also with the artificially constructed world, and even more so with artificially intelligent entities.

This relational approach to artificial intelligence—offering an alternative to the property-based approach—might seem inconsistent with Deep Ecology on account of the eight principles, which seem to work as a checklist. But that is not what we should take away from Deep Ecology. Indeed, its central insights on the question of whether other entities (natural or artificial) ought to be recognized as having a moral status revolve around the appreciation that we share the same interactive environment with them.

The second of the two previously mentioned arguments against the idea of applying Deep Ecology to artificial intelligence raises a problem of coherence. The argument proceeds from the fundamental distinction between nature and artificially intelligent entities: Nature is not a human creation, while artificially intelligent entities are. Ergo: If it is wrong for humans to exploit nature, why should it be right to do so with something (or someone) they created themselves and which (or who) would not exist without the human beings that developed them in the first place?

The flaw in this argument lies in its factual premise, in that a large chunk of nature *is* in fact created by human beings. Let us consider domestic animals, like dogs. Many dog breeds are created by human beings, and if it weren't for such human breeding, those breeds would not exist. The same applies to different kinds of other animals and plants developed using different techniques (such as genetic engineering and plant breeding). So, on this reasoning, we should draw a distinction among nonhuman life forms that are developed by human beings and nonhuman

¹³More information about Jibo is available at www.jibo.com

¹⁴More information available at <https://www.aldebaran.com/en/cool-robots/pepper>

¹⁵More information available at <http://www.cs.cmu.edu/~coral/projects/cobot/>

life forms whose development is not owed to human intervention, and we should therefore preserve the latter and ignore the former. A moment's thought, however, should suggest that it may not be a good idea to extract moral consequences from such a distinction: The "authorship" of some species is not a license to treat these species however we like. This applies to animals and plants, and it should also apply to artificially intelligent entities.

Furthermore, if we want to discuss the role that humans play in the nonhuman world, we should frame the discussion in terms not of authorship but of stewardship, which implies a duty of care and responsibility to nonhumans: This is what Jonas (1984) argues as concerns nature, and what Floridi and Sanders (2004) argue as concerns artificial agents. And although Deep Ecologists do not like the idea of stewardship (Devall and Sessions 1985), the idea may well serve as another starting point for a discussion of what it is to relate to that which surrounds us.

18.5 Conclusions

In this paper we considered the question of how we ought to relate to artificially intelligent entities as entities forming part of our natural and constructed environment. It was suggested that that one solution may come from Deep Ecology, a theory that in the description of its founder, Arne Naess, "asks deeper questions: we ask why and how, where others do not" (Devall and Sessions 1985, 74). This is why I believe that Deep Ecology can provide an interesting lens through which to discuss the moral standing of artificially intelligent entities: we see these entities as things, and we seldom ask the deeper questions that go beyond the anthropocentric conception, from which comes a spectrum of stances ranging from unfiltered consumerism to the antagonistic "we-against-them" mindset.¹⁶

The underlying idea of this paper is precisely that we should not draw any sharp distinctions between the natural environment made of living organisms (plants and animals) and the artificial environment we shape either by design or as a consequence of what we do with the designs we put out into the world. If we can appreciate the inherent value of that overall environment and the relations it depends on for its own sustenance, we can see that its constituent entities may be worthy of moral consideration independently of their usefulness to human welfare: We can thus include artificially intelligent entities in that group (comprising a growing range of entities), and to that end we need not necessarily rely on a standard list of properties such as sentience, consciousness, intelligence, or the ability to use a language.

¹⁶In addition, even if we stick to the notion of artificially intelligent beings as things, we could still ask deeper questions about them: In this way, as Holy-Luczaj (2015, 59–60) argues, we could "stop regarding them [things] as (easily) replaceable disposables," and "such a transformation [would] likely change the patterns of our consumption and thereby [have] a positive proenvironmental impact."

One reason suggesting that this may not be an appropriate set of metrics, or at least that the property-based approach may not work as a standalone solution, is that artificially intelligent entities may even *outstrip* biological entities in their capacities (such as the use of language), yielding the counterintuitive conclusion that they are worthy of even greater moral consideration than other beings in that respect. But the point here is not to rank different sorts of entities according to their degree of moral worth: It is rather to see whether they can be included as participants in our environment by looking at the role they play within that environment, and to see what moral consequences can be extracted on that basis.

That is why Deep Ecology seems to offer itself as an appropriate vantage point: It enables us to frame the moral and legal problem of artificial intelligence on an inter-relational approach closer to the kind of approach that has already been shown to work in tackling the great moral and political issues of inclusion and exclusion we have faced in the past. And it can do so drawing on philosophical insights from a broad range of inquiries. As Palmer (1998, 164) has noted, “a small and controversial new philosophical school [gains] revealing conceptual closeness to relatively illustrious philosophical ancestors,” and that is precisely one of the aims of this paper: to show that environmental ethics in general, and Deep Ecology in particular, can draw on a broad range of insights from the past in dealing with a problem that is facing us now in the present and is poised to become even more pressing in the future. There are many aspects of environmental ethics and Deep Ecology that I do not discuss here, but I hope to have at least offered some good reasons for looking at artificially intelligent entities as entities forming part of a shared environment, for I submit that from this vantage point we can make some headway in dealing with some of the moral and legal issues their use and development might give rise to.

The idea of applying Deep Ecology to artificial intelligence runs parallel to the contemporary legal and ethical discussions on artificial intelligence and robotics: A debate is underway on whether to recognize electronic personhood for robots and whether the problem of their use can be managed within the current legal framework. It seems to me that before these questions can be given any definite solution, we need to challenge our assumptions. Naess (2008, 311), for example, argues that what we need is not a “shift from humans towards nonhumans, but an extension and deepening of care.” This highlights a contrast between two different paradigms, suggesting that if we are to properly deal with artificial intelligence, we need to effect something along the lines of Naess’s “substantial reorientation of our whole civilization” (*ibid.*).

Acknowledgements This paper is part of the project ALLIES (Artificially Intelligent Entities: Their Legal Status in the Future) that has received funding from the Universidad Carlos III de Madrid, the European Union’s Seventh Framework Programme for research, technological development and demonstration under grant agreement n° 600371, el Ministerio de Economía, Industria y Competitividad (COFUND2014-51509) el Ministerio de Educación, cultura y Deporte (CEI-15-17) and Banco Santander.”

References

- Armstrong, A.C. 2012. *Ethics and justice for the environment*. Abingdon: Routledge.
- Baard, P. 2015. Managing climate change: A view from deep ecology. *Ethics & The Environment* 20 (1): 23–44.
- Bedau, M.A., et al. 2000. Open problems in artificial life. *Artificial Life* 6: 363–376.
- Bostrom, N. 2014. *Superintelligence: Paths, dangers, strategies*. Oxford: Oxford University Press.
- Coeckelbergh, M. 2009. Distributive justice and co-operation in a world of humans and non-humans: A contractarian argument for drawing non-humans into the sphere of justice. *Res Publica* 15 (1): 67–84.
- . 2010. Robot rights? Towards a social-relational justification of moral consideration. *Ethics and Information Technology* 12 (3): 209–221.
- Devall, B., and G. Sessions. 1985. *Deep ecology: Living as if nature mattered*. Salt Lake City: Peregrine Smith Books.
- Farshchi, S. 2016. Let's bring Rosie home: 5 challenges we need to solve for home robots. *IEEE Spectrum* January 13. http://spectrum.ieee.org/automaton/robotics/home-robots/lets-bring-rosie-home-5-challenges-we-need-to-solve-for-home-robots/?utm_source=computerwise&utm_medium=email&utm_campaign=011916.
- Florida, L., and J.W. Sanders. 2001. Artificial evil and the foundation of computer ethics. *Ethics and Information Technology* 3: 55–66.
- . 2004. On the morality of artificial agents. *Minds and Machines* 14 (3): 349–379.
- Fox, W. 2003. Deep ecology: A new philosophy of our time? In *Environmental ethics: An anthology*, ed. A. Light and H. Rolston III, 252–261. Oxford: Blackwell.
- Gottlieb, R.S. 1999. Introduction. In *Ecology*, ed. C. Merchant, ix–xi. Amherst: Humanity Books.
- Holy-Luczaj, M. 2015. Heidegger's support for deep ecology reexamined once again: Ontological egalitarianism, or farewell to the great chain of being. *Ethics & The Environment* 20 (1): 45–66.
- Jonas, H. 1984. *The imperative of responsibility: In search of ethics for the technological age*. Chicago/London: University of Chicago Press.
- Keller, D.R. 2008. Deep ecology. In *Encyclopedia of environmental ethics and philosophy*, ed. J. Baird Callicott and R. Frodeman, 206–211. Detroit/New York: Gale Cengage Learning.
- , ed. 2010. *Environmental ethics: The big questions*. Chichester: Wiley-Blackwell.
- Kellogg, P. 2014. Do machines have rights? ethics in the age of artificial intelligence. *Aurora*, <http://aurora.icaap.org/index.php/aurora/article/view/92>.
- Keulartz, J. 1995. *The struggle for nature: A critique of radical ecology*. London/New York: Routledge.
- Kortetmäki, T. 2016. Is broad the new deep in environmental ethics? A comparison of broad ecological justice and deep ecology. *Ethics & The Environment* 21 (1): 89–108.
- Kurzweil, R. 1999. *The age of spiritual machines: When computers exceed human intelligence*. New York: Viking Press.
- . 2006. *The singularity is near: When humans transcend biology*. New York: Viking Press.
- Naess, A. 1973. The shallow and the deep, long-range ecology movement: A summary. *Inquiry* 16 (1): 95–100.
- . 1986. The deep ecological movement: Some philosophical aspects. *Philosophical Inquiry* 8: 1–2.
- . [1989] 2001. *Ecology, community and lifestyle: Outline of ecosophy*. Cambridge: Cambridge University Press.
- . 2005. Theoretical dimension of deep ecology and ecosophy. In vol. 10 of *The selected works of Arne Naess*, ed. A. Naess, and A. Drengson, 546–550. Dordrecht: Springer.
- . 2008. *The ecology of wisdom: Writings by Arne Naess*. Berkeley: Counterpoint.
- Nash, R.F. 1989. *The rights of nature: A history of environmental ethics*. Madison: The University of Wisconsin Press.
- Palmer, C. 1998. *Environmental ethics and process thinking*. Oxford: Clarendon Press.

- Rothblatt, M. 2014. *Virtually human: The promise—and the peril—of digital immortality*. New York: St. Martin's Press.
- Tavani, H.T. 2011. *Ethics and technology: Controversies, questions and strategies for ethical computing*. 3rd ed. Hoboken: Wiley.
- Troster, L. 2008. Caretaker or citizen: Hans Jonas, Aldo Leopold, and the development of Jewish environmental ethics. In *The legacy of Hans Jonas: Judaism and the phenomenon of life*, ed. H. Tirosh-Samuelson and C. Wiese, 373–396. Leiden/Boston: Brill.
- Yampolskiy, R.V. 2012. Artificial intelligence safety engineering: Why machine ethics is a wrong approach? In *Philosophy and theory of artificial intelligence*, ed. V. Müller, 389–396. Berlin: Springer.