Anthropocentrism, Artificial Intelligence, and Moral Network Theory: An Ecofeminist Perspective

VICTORIA DAVION

Philosophy Department
College of Arts and Sciences
University of Georgia
107 Peabody Hall, Athens, Georgia 30602-1627, USA
Email: vdavion@arches.uga.edu

ABSTRACT

This paper critiques a conception of intelligence central in AI, and a related concept of reason central in moral philosophy, from an ecological feminist perspective. I argue that ecofeminist critique of human/nature dualisms offers insight into the durability of both problematic conceptions, and into the direction of research programmes. I conclude by arguing for the importance of keeping political analysis in the forefront of science and environmental ethics.

KEY WORDS

Ecofeminism, artificial intelligence, intelligence, moral theory

Philosophers and others have noted that problematic conceptions of 'intelligence' and 'reason' remain durable in the fields of artificial intelligence (AI) and moral philosophy. Some wonder why these conceptions remain so highly durable given evidence that challenges them. I argue that an understanding of this durability is enhanced by political analysis, and that ecological feminism can provide a viable theoretical framework from which to conduct such an analysis. In what follows I first present an ecological feminist analysis of dualistic thinking. I then apply this analysis to a conception of intelligence central in AI, and a related concept of reason central in moral philosophy. I show that ecological feminist critique of human/nature dualisms, including anthropocentrism, offers insight into both the durability of concepts and the direction of research programmes. I conclude with some general remarks about the importance of keeping political analysis in the forefront in discussions of science and environmental ethics.

Environmental Values **11** (2002): 163–76

VALUE DUALISMS AND OPPRESSIVE CONCEPTUAL FRAMEWORKS

Analysis of value dualisms plays a prominent role in ecological feminist critiques of western patriarchal cultures. A value dualism is a disjunctive pair in which the disjuncts are seen as oppositional and exclusive, and which place higher value on one disjunct than the other (Warren 1990; Plumwood 1993). Many ecological feminists argue that a reason/nature dualism underlies the conceptual framework of western patriarchal cultures. This dualism is thought to form the basis for a series of related dualisms in which whatever is associated with reason is viewed as fundamentally different from and superior to whatever is associated with nature.

Examples of such dualised pairs involve not only human/nature reason/nature, masculine/feminine, but also mental/manual, civilised/primitive, and human/nature. These pairs function to legitimate a number of oppressions, including sex, race, and class oppression, which can all be seen in terms of the central dualism underlying the system, that of reason/nature.

It is crucial to an accurate understanding of the ecofeminist critique of value dualisms to realise that not all differences are dualisms, and that deconstructing value dualisms does not mean denying all differences between dualised pairs. The problem with value dualisms lies in the construction of dualised pairs as absolutely different in morally relevant ways, which leads to the justification of moral hierarchies.

THE CONSTRUCTION OF VALUE DUALISMS

The construction of dualised identities involves five features, according to Plumwood. These are (1) backgrounding, (2) radical exclusion, (3) incorporation, (4) instrumentalism, and (5) homogenisation (Plumwood 1993). Each of these features functions to validate the relationship of denied dependency that is at the core of dualistic relationships.

Backgrounding involves an oppressor's attempt to use the oppressed, which creates a dependency on the oppressed, and to deny simultaneously that dependency. This is often achieved by denying the importance of the contribution of the oppressed by devaluing the forms of life associated with the oppressed group. A good example is the devaluation of the so-called 'private' realm of the family or household that has been traditionally associated with women. The contribution of the oppressed is backgrounded, devalued, not the focus of attention.

The relation of radical exclusion involves not merely recognising some differences between dualised pairs, but seeing them as radically different. The

number and importance of differences in maximised, shared characteristics that cannot be denied, are viewed as inessential. Thus, any continuity between the pair is either denied or seen as unimportant. This helps in the construction of the idea of radically different and separate 'natures' of oppressed and oppressor, thus justifying oppression and domination, and making it appear natural or inevitable.

Incorporation involves constructing the identity of the devalued side of the dualised pair in terms of the underside's lack of morally relevant features associated with the other side. Since qualities that do not fit into the scheme are ignored, the other is assimilated or incorporated into the oppressor's sense of self, and does not have to be dealt with as an independent entity, important in its own right.

This leads to the next feature, instrumentalism. Those groups seen as morally inferior are constructed as having no morally important independent interests; thus, they are seen as valuable only instrumentally, in so far as they can be of use in promoting the interests of the morally relevant groups. Finally, homogenisation involves denying the differences between those on the underside of dualised pairs; thus, seeing all women, or all slaves, as basically the same, as merely other.

ANTHROPOCENTRISM IN AI

Theorists have noticed that a highly anthropocentric notion of intelligence has dominated AI. The following are some definitions of intelligence from prominent sources in the field (all quotes from Preston 1991):

> [Artificial Intelligence is] the science of making machines do things that would require intelligence if done by men [sic]. (Minsky 1968: v)

> Artificial intelligence is the study of ideas that enables computers to be intelligent. (Winston 1984: 1)

> Artificial intelligence is the study of mental faculties through the use of computational models. (Charniak and McDermott 1985: 6)

> Artificial Intelligence (AI) is the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the characteristics which we associate with intelligence in human behaviour – understanding language, learning, reasoning, solving problems, and so on. (Barr and Feigenbaum 1981: 3)

> Artificial intelligence is a field of study concerned with designing and programming machines to accomplish tasks that people accomplish using their intelligence. (Schutzer 1987: 1)

> Artificial intelligence is a field of study that encompasses computational techniques for performing tasks that apparently require intelligence when performed by humans. Such problems include diagnosing problems in automobiles, computers, and people, designing new computers, writing stories and symphonies, finding mathematical theorems, assembling and inspecting products in factories, and negotiating treaties. (Tanimoto 1990: 6)

As Preston makes clear, the use of the word 'intelligence' in many of these definitions indicates a reference to some pretheoretical notion of intelligence that the authors assume their audiences will share, and that this notion is a strongly anthropocentric one.

This strongly anthropocentric notion of intelligence led to research programmes focusing on recreating human intelligence understood strictly in terms of abstract formal reasoning, which was taken to be the essential component of human intelligence. Interestingly, the assumption that the essential component of human intelligence is abstract formal reasoning is challenged by problems in modelling 'common-sense knowledge' and perception. What has come to be known as 'the problem of common- sense knowledge' in AI refers to the fact that programs which could solve abstract formal problems were unable to solve even simple, real-life problems such as getting a cat out from under the bed (Preston 1991). The reason for this failure is that real-life problems require more than consideration of entities with formal properties and relations to each other. They require knowledge about objects in the world. A program that could solve fairly high-level algebra and word problems failed to understand a story accessible to a three-year-old child.

The other problem concerns attempts to model skills involved in perception and movement. AI researchers assumed that it would be fairly easy to model perceptual capacities such as vision. However, this turned out to be very difficult to model. And, other capacities involved in managing the world such as motor control, navigation, and manipulation have also been underestimated in terms of complexity. Hence, capacities clearly shared with other animals (vision, navigation) were assumed to be less complex than the capacity for abstract reasoning, often assumed to be exclusively human, or to be possessed by humans to the greatest degree.

According to Preston:

> The original anthropocentric premise is that human-level intelligence is the most significant and interesting form of intelligence. But this view is challenged by the realization that the great bulk of intelligent behavior depends on aspects of behavior which we share with infra-human intelligences, and which are on the whole more complex and difficult to understand than the uniquely human aspects of intelligence. (Preston 1991: 270)

The anthropocentrism obvious both in the assumption that the most important and interesting kind of intelligence is human intelligence, and that the most interesting aspects of human intelligence are those which are uniquely human is undermined.

Preston notes that the failure of early AI research programmes to recognise what she calls 'peripheral intelligence' as intelligence at all can be attributed to the lack of recognition of the complexity involved in peripheral intelligence to begin with. However, '...an equally important factor was the sheer weight of the Western intellectual tradition, which routinely elevates reason and language and denigrates the senses' (Preston 1991: 269).

The conception of intelligence as abstract reasoning is an excellent example of dualistic thinking at work. We can see this by applying features involved in the construction of dualised identities offered by Plumwood. Backgrounding is evident in the assumption that human intelligence is just abstract reasoning, something unrelated to infra-human abilities, or that humans supposedly do best. The relation of radical exclusion, seeing differences between dualised pairs not simply as differences but as essential differences, is demonstrated in assumptions that only human-level intelligence is interesting, and that it is radically different from infra-human abilities (or intelligences). This is also evident in the intuitive conception of intelligence originally adopted by AI researchers, that the only important and interesting aspects of intelligence are the ones that were taken to be prototypically, if not exclusively, human.

In addition, ecological feminism provides a political framework for understanding assumptions evident in dualistic thinking. Ecological feminism can help in understanding the two factors that Preston speculates led to such strong anthropocentrism in AI. A central project of ecological feminist philosophy is to show how elevating reason over nature, in combination with the association of privileged men with language and reason, and the association of nature, women, and other men with the senses and the emotional, contributes to sexism, naturism, and racism in western patriarchies (Griffin 1978; Merchant 1983; Warren 1990). According to Plumwood:

> The key exclusions and denials of dependency for dominant conceptions of reason in western culture include not only the feminine and nature, but all those human orders treated as nature and subject to denied dependency. (Plumwood 1993: 42)

Thus:

> The set of interrelated and mutually reinforcing dualisms which permeate western culture forms a fault line which runs through the entire conceptual system… Each of them has crucial connections to other elements and has a common structure with other members of the set. The can be seen as forming a system, an interlocking structure. (Ibid.: 42)

Hence, ecological feminist analysis argues that the conception of intelligence evident in AI is not only anthropocentric, but it is androcentric and racist as well, as it is part of a structure that justifies white male supremacy. Understanding its political function is crucial in understanding its durability.

At this point it is important to deal with a possible misunderstanding of ecofeminist analysis. Ecofeminists do not claim that those who use problematic dualistic concepts are necessarily attempting to promote white supremacist patriarchy in doing so. Quite the opposite. The key point for ecofeminists is that dominant portions of the western tradition have laid the groundwork so that dualistic anthropocentric concepts promoting racism, sexism, and naturism *seem intuitive*, and no argument for them is needed. The assumption that perceptual abilities shared with other animals are not as complex as abstract formal reasoning is certainly a factor in why they were not the focus of AI work on intelligence. However, the 'sheer weight of the western tradition,' which Preston names as the other factor in why perceptual capacities were not central, is not another factor at all. Rather, ecological feminist analysis explains how failing to see abilities shared with other animals as interesting and complex is actually part of the dominant western tradition itself. Hence, the reason AI theorists were able to count on people's intuitions in defining intelligence in a highly anthropocentric manner is because the fault line referred to by Plumwood is so solidly in place that no further explanation is needed. Our dualistic intuitions simply kick in. Therefore, it would be no surprise to find the very theorists who rely on such intuitions denying that they had *any* intention of maintaining or endorsing white male supremacy in deploying particular conceptions of intelligence. I have no doubt that this is true. In fact, this is the point. And, it is exactly such *obvious* intuitions that ecological feminists argue require a political analysis.

In the next section of this paper I argue that the conception of 'moral reasoning' in the dominant tradition in western philosophy shares many problematic assumptions with the concept of 'intelligence' discussed above. And, I argue that ecological feminist analysis can once again provide a political analysis of why this problematic assumption remains durable.

NEURAL NET RESEARCH AND MORALITY

Recent developments in the areas or neural net research (NNR) area suggest that moral learning, deliberation, and action may not be matters of applying abstract rules to particular situations in the way that the dominant tradition in western philosophy has assumed. While failures in AI challenge the idea that rule-based reasoning is the most complex form of human behaviour, neural net research suggests that our moral capacities may not involve this kind of reasoning, challenging the notion that human's moral capacities result from some ability

that is supposedly radically different from anything that other animals do. Because ecological feminist critique explains the threat involved in displacing the rule-based picture of moral reasoning for western patriarchal order, it can be useful to theorists who have noticed a resistance to exploration of implications of neural net research for moral thinking. My discussion of neural net research (NNR) will draw heavily from the works of Paul M. Churchland (1996) and Owen Flanagan (1996) .

Churchland argues that NNR provides strong reasons for believing that moral learning, deliberation, and action are not a matter of applying abstract rules to concrete situations. Rather, humans' moral capacities are more like perceptual capacities, which seem to be the result of well-tuned neural networks. Learning to respond to social reality is not radically different from learning to respond to other aspects of the environment. Dealing with all the necessities for survival require the same sorts of neuronal resources and coding strategies. In the case of moral reasoning, 'The job may be special, but the tools available are the same' (Churchland 1996: 92). How do we know the tools are the same? Simply because these appear to be the only tools that are present in the brain. Neural net theory uses a prototype activation model to explain how creatures learn about and respond to their physical environments. According to the prototype activation model, we learn about the physical world by learning how to recognise and respond to prototypical situations. We then learn how to redeploy prototypes in response to the continuous unfolding of new situations.

Neural network research indicates that one's ability to recognise and discriminate among perceptual properties goes far beyond one's ability to articulate the basis of these discriminations in words. Thus, one's ability to recognise a particular taste usually goes beyond one's ability to describe the taste. The ability to recognise a face is another good example. 'In fact, the cognitive priority of the preverbal over the verbal shows itself upon examination, to be a feature of almost all of our cognitive categories' (Ibid.: 101). Thus:

> One's ability to recognize instances of cruelty, patience, meanness, and courage, for instance, far outstrips one's capacity for verbal definitions of those notions. One's diffuse expectations of their likely consequences similarly exceed any verbal formulas that one could offer or construct, and those expectations are much more penetrating because of it. All told, moral cognition would seem to display the same profile or signature that in other domains indicates the activity of a well-tuned neural network underlying the whole process. (Ibid.: 101)

Churchland maintains that in all cases, the ability to respond to our environments is not a matter of applying rules to a particular situation, but an ability to activate a correct prototype for the situation and respond to it appropriately. With respect to moral learning and behaviour, I shall refer to this as 'moral network theory', following Owen Flanagan (1996). According to this approach, moral understanding and learning is a process of learning how to

recognise a wide variety of complex situations and how to respond to them appropriately. '[T]here is a straightforward analogy between the way a submarine sonar device that needs to learn to distinguish rocks from mines might acquire the competence to do so and the way a human might acquire moral sensitivities and sensibilities' (Flanagan 1996: 25). Hence, according to Flanagan:

> One way to teach a mine-rock device would be simply to state the rule specifying the necessary and sufficient characteristics of rocks and mines. The trouble is that these are not known (indeed part of the mines producer's job is to make them as physically indistinct as possible). Despite these efforts at disguise, there are bound to be (or so we hope) subtle features that distinguish mines from rocks, so it would be good if the device could be trained in a situation where it starts by guessing mine or rock, and then, by being clued into the accuracy of its guesses, develops a profile for recognizing rocks from mines. . . Eventually the mine-rock detector (which of course is never perfect at its job) comes to be able to make judgments of kind very quickly, based on a small number of features, and it responds accordingly. (Ibid.: 25)

Flanagan maintains that children learn about morality in a similar way.

> According to moral network theory, the fundamental process is the same for moral learning. Children learn to recognize certain prototypical kinds of situations, and they learn to produce or avoid the behaviors prototypically required or prohibited in each. Children come to see certain distributions of goodies as fair or unfair distribution. They learn to recognize that a found object may be someone's property, and that access is limited as a result. They learn to discriminate unprovoked cruelty, and to demand or expect punishment for the transgressor and comfort for the victim. (Ibid.: 28)

In all cases, our abilities far exceed any rules we might articulate.

Described in this way, moral response is harder to construct dualistically, as something for which only humans have the capacity. There are interesting similarities in the way that Flanagan describes moral response, and the way that some researchers describe animal responses to what look like unfair situations. The book *When Elephants Weep* (Mason and McCarthy 1995), documents many instances in which it appears that animals respond with what seems like a sense of justice. Researchers explain how Nim Chimpsky, a chimpanzee taught to sign, became indignant when he felt he was treated unfairly.

> Nim Chimpsky learned when to expect praise and when to expect blame and accepted these artificial standards. If he broke a toy punishment did not surprise him and he accepted it. But if one of his teachers punished him for something that others ignored or if one failed to praise him for something that others rewarded, Nim became sulky. (Mason and McCarthy 1995: 214)

Researchers suggest that Nim became upset due to lack of predictability or violation of settled expectations. They comment, '...but this is a big part of legal justice among humans' (Ibid.: 214). In commenting on the chimpanzees of the Arnhem colony, Frans de Waal states:

> [They] ... seem to react to a sense of unjust treatment towards others. In one instance the chimpanzee Puist 'kidnapped' a one-year old infant from his mother, and carried him up a tree, where he screamed with fright. After the mother recovered her child she attacked Puist, although Puist was larger and more dominant. The male Yeroen rushed up to them and stopped the fight by seizing Puist and flinging her away. This was unusual, because on all other occasions Yeroen had intervened on Puist's side. (Ibid.: 214)

Frans de Waal concludes that Yeroen agreed with the mother chimp that she had cause of complaint. In another instance:

> Puist appeared aggrieved on her own behalf, backing up Luit in a dispute with a large male. The male made a threatening display at Puist, who stretched out her hand in appeal to Luit. Luit did not respond, and Puist rushed him, apparently because he had violated the tradition of supporting one's supporters. This kind of solidarity is part of many human notions of fairness. (Ibid.: 215)

Finally:

> Observers of wild coatimundis in Arizona suggest that they have a system of entitlement through a variety of squeals. When cuffed by an older male for lagging behind the troop, a cub would crouch submissively and utter the 'don't beat me' squeal, which seemed to indicate resistance. On several occasions, when a subadult animal committed the usual act of trying to take food from a cub and cuffed it, the cub would utter a different squeal and an adult female would come and drive the subadult away, apparently enforcing a tradition of tolerance toward cubs. This may be merely different cub feelings being expressed in different situations of threat, but it is telling that there is a difference. Enforcing and cushioning hierarchy also plays a part in human justice systems. (Ibid.: 215)

IMPLICATIONS FOR ANTHROPOCENTRISM

If it is true that moral reasoning and moral behaviour are not matters of the rational person applying rules to particular situations, and if moral deliberation has its origins in the same kinds of processes that allow other animals to function successfully in their own environments, then this severely challenges the idea that our moral natures spring from something other animals, even very simple ones, lack. Radical exclusion is once again challenged. Churchland argues:

> Social and moral cognition, social and moral behavior, are no less activities of the brain than is any other kind of cognition or behavior. We need to confront this fact squarely and forthrightly, if we are ever to understand our moral natures. We need to confront it if we are ever to deal both effectively and humanely with our too-frequent social pathologies. And we need to confront it if we are ever to realize our full social and moral potential. (Churchland 1996: 92)

Yet, he goes on to say:

> Inevitably these sentiments will evoke discomfort in some readers, as if by being located in the purely physical brain, social and moral knowledge were about to be devalued in some way. Let me say, most emphatically, that devaluation is not my purpose. (Ibid.: 92)

Churchland notices that researchers have not expended the same energy concentrating on how beings navigate in the social world as they have spent on learning how beings navigate in what he calls 'the purely physical world'. He believes that if they did, we would find that navigation of the social world involves well-tuned neural nets. Experimental neuroscience in the twentieth century has focused on finding neuroanatomical (structural) and neurophysiological (activational) correlates of perceptual properties that are purely natural or physical in nature. The central programmatic question has been: Where in the brain, and by what processes, do we recognise such properties as colour, shape, motion, sound, taste, aroma, temperature, texture, bodily damage, relative distance, and so on? According to Churchland, this research has resulted in an ability to map areas of the brain that seem centrally involved in each of the functions mentioned. Researchers have learned that the rear half of a typical primate cerebral cortex is involved in the perception of 'purely physical' properties, and more specifically which areas of the brain are involved in the perception of the various different ones.

Scientists have yet to map the front lobe of the brain. Churchland believes there are 'social areas' of the brain, involving the as-yet unmapped areas. He asks:

> Might some of these areas be principally involved in social perception and action? Might they be teaming with vast vectorial sequences representing social realities of one sort or another? Indeed, once this question is raised, why stop in these areas? Might the so-called primary sensory cortical areas – for touch, vision, and hearing especially – be as much in the business of grasping and processing social facts as they are in the business of grasping and processing purely physical facts? These two functions are certainly not mutually exclusive. (Ibid.: 100)

Churchland contends that the answers to at least some of the above questions are almost certainly yes, and that the reason that we do not have intricate maps for social features which are comparable to existing brain maps for physical features

is not because they do not exist, but because 'we have not looked for them with a determination comparable to the physical case' (Ibid.: 100).

Churchland does not speculate as to why identifying something in the realm of the purely physical would devalue it. And, it is a bit unclear exactly what is meant by 'purely physical' here. However, ecological feminists might argue that a deeply embedded mind/body dualism is kicking in, associating mind with reason and body with nature. Using the fault line idea suggested by Plumwood, we can then see why challenging such a dualism would also challenge many other key dualisms at the core of white male supremacy. These include, most obviously, human/nature and mental/physical. Therefore, challenging the core of dualistic thinking which has not only dominated western philosophy, but which has been used politically to justify a complex hierarchy including anthropocentrism, sexism, racism, homophobia, and so forth.

The point here is not to argue that moral network theory is the correct model for moral response either in humans or in other animals. Rather, the point is that ecological feminist analysis provides an interesting lens through which to view the resistance to moral network theory that Churchland notices. As far as I know, there is no dominant western tradition that claims animals do not perceive, and hence, seeing moral response as perceptual and emotional, rather than formal and objective, challenges key dualisms forming the fault line mentioned by Plumwood above. Most obviously, this challenges a number of arguments against the moral considerability of other nonhuman animals.

A CLOSING NOTE ON ANTHROPOCENTRISM

In closing, I shall examine implications of this discussion for the ongoing anthropocentric/nonanthropocentric debate in environmental ethics. This debate concerns whether a paradigm shift from anthropocentrism to non-anthropocentrism is needed to achieve an adequate level of environmental protection and concern. Ecological feminists and other radical environmentalists argue such a shift is not only useful for environmental protection but is morally required (Naess 1973; Griffin 1978). Others, such as Bryan Norton, argue such a shift is unnecessary and politically divisive (Norton 1987; Weston 1982; Dobson 1990). An adequate ecological understanding of human dependence on natural systems along with moral concern for the well-being of future generations can furnish the basis for an environmental ethic with the same practical consequences as a shift to nonanthropocentrism. Norton argues that because we can generate the same practical consequences without a paradigm shift, the anthropocentric/nonanthropocentric debate is at best a useless detour, and at worst politically divides people who would be better off working together. Hence, according to Norton, 'the theory that environmentalists should be sorted into two camps according to commitment or lack of to the principle that nature

has independent value ...[leads] us to no important differences between environmentalists and their critics among deep ecologists' (Norton 1987: 236). In defence of this claim Norton offers a convergence principle which states that 'policies serving the interests of the human species as a whole, and in the long run, will serve also the "interests" of nature, and vice versa' (Ibid.: 240). Understanding this requires adopting the 'ecological world view'. On this view:

> Two great ideas come together here: Darwinian biology has taught us that humans are, basically, evolved animals; ecology has taught that evolution works within almost unbelievably complex and interrelated organic systems on the interlocking levels ranging through molecules, cells, organs, organisms, habitats, ecological systems, the biota as a whole, and ultimately the abiotic system. (Norton 1987: 204)

Therefore:

> The natural history of *Homo sapiens*, viewed as a highly evolved and highly intelligent but physically dependent being that has survived in a hostile world, can stand as a guide to human behavior. (Ibid.: 206)

The anthropocentric/nonanthropocentric debate raises many interesting questions I will not delve into here. I agree with those who argue that an ethic affording only instrumental value to nonhuman beings cannot offer the same level of protection and concern for them as an ethic affording them moral value in their own right (Naess; Plumwood). I want to focus on a different issue. Norton's claim that the science of ecology can provide the necessary understanding of human dependence on the rest the natural world relies on the idea of objective value-free science, untainted by the moral and political contexts in which it is generated. Feminists, among others, have pointed out that science is deeply influenced by background assumptions determined by the political and moral contexts from which it emerges (Harding 1986; Code 1991). This includes everything from which research agendas are formulated, to what gets funding, to which results are taken as credible. Science that challenges anthropocentrism, including notions of the hyperseparation of humans from nature, men from women, and so forth, may not be taken seriously, if it gets generated at all. As my discussions of AI and neural net research clearly demonstrate, background assumptions are crucial to the determination of what kinds of research are important. Science beginning with anthropocentric, dualistic assumptions will generate very different research programmes from science which does not. The anthropocentric/nonanthropocentric debate is not only a moral debate about the moral importance of nonhuman beings, it also has implications for how we can expect research programmes in science to be formulated and carried out. Hence, whether we adopt an anthropocentric or a nonanthropocentric paradigm can be expected to impact on what science will teach us about nature and ecology . Even if it is true that there is an objective, value-free, scientific perspective which can

reveal the ecological world view, I hope my analysis shows concretely that it matters very much whether one begins with anthropocentric or nonanthropocentric assumptions. This impacts not only on which questions are asked, but also on which answers are taken seriously.

ACKNOWLEDGEMENT

I want to thank reviewers from *Environmental Values* for extremely helpful comments on an earlier draft of this essay.

REFERENCES

Barr, Avron and Edward Feigenbaum 1981. *The Handbook of Artificial Intelligence,* Vol. 1. Los Altos, Calif.: William Kaufmann.

Charniak, Eugene and Drew McDermott 1985. *Introduction to Artificial Intelligence*. Reading, Mass.: Addison-Wesley.

Churchland, Paul M. 1996. 'The Neural Representation of the Social World', in Larry May, Marilyn Friedman and Andy Clark (eds) *Mind and Morals: Essays on Ethics and Cognitive Science*. Cambridge, Mass. and London: MIT Press.

Code, Lorraine 1991. *What Can She Know? Feminist Theory and the Construction of Knowledge.* Ithaca, N.Y. and London: Cornell University Press.

Dobson, Andrew 1990. *Green Political Thought.* London: Routledge.

Flanagan, Owen 1996. 'Ethics Naturalized: Ethics as Human Ecology', in Larry May, Marilyn Friedman and Andy Clark (eds) *Mind and Morals*: *Essays on Ethics and Cognitive Science.* Cambridge, Mass. and London: MIT Press.

Griffin, Susan 1978. *Woman and Nature: The Roaring Inside Her.* New York: Harper and Row.

Harding, Sandra 1986. *The Science Question in Feminism.* Ithaca, N.Y. and London: Cornell University Press.

Mason, Jeffrey Moussiaeff and Susan McCarthy 1995. *When Elephants Weep*. New York: Delacorte Press.

Merchant, Carolyn 1983. *The Death of Nature: Women, Ecology, and the Scientific Revolution*. New York: Harper and Row.

Minsky, Marvin (ed.) 1968. *Semantic Information Processing.* Cambridge, Mass.: MIT Press.

Naess, Arne 1973. 'The Shallow and the Deep, Long-range Ecology Movements: A Summary', *Inquiry* **16**: 95–100.

Norton, Bryan G. 1987. *Why Preserve Natural Variety?* Princeton: Princeton University Press.

Plumwood, Val 1993. *Feminism and the Mastery of Nature*. London: Routledge.

Preston, Beth 1991. 'AI, Anthropocentrism, and the Evolution of 'Intelligence', *Minds and Machines* I: 259–77.

Schutzer, Daniel 1987. *Artificial Intelligence: An Applications-oriented Approach.* New York: Van Nostrand Reinhold.

Tanimoto, Steven L. 1990. *The Elements of Artificial Intelligence.* New York: Computer Science Press (W. H. Freeman).
Warren, Karen J. 1990. 'The Power and Promise of Ecological Feminism', *Environmental Ethics* **12**(2): 125–46.
Weston, Anthony 1992. *Towards Better Problems.* Philadelphia: Temple University Press.
Winston, Patrick Henry 1984. *Artificial Intelligence* (2nd edition). Reading, Mass.: Addison-Wesley.