

Potential Impacts of HRI upon Different Understandings of Morality

Theodore A. Metzler

Wimberly School of Religion, Oklahoma City University
2501 N. Blackwelder
Oklahoma City, OK 73106-1493
tmetzler@okcu.edu

Abstract

This paper examines compatibility of some existing understandings of morality with the likely phenomenon of increasing human-robot interaction (HRI) in the United States. Particularly for cases involving autonomous humanoid robots, it is argued that impacts of HRI may differ dramatically as a function of the received understandings of morality encountered.

Introduction

In the following, we are addressing human interaction with artificially intelligent artifacts that increasingly resemble – both in appearance and in behavior – whole biological creatures, such as dogs, seals, and *people*. Researchers already engaged in studying effects of this newer class of robots upon humans during HRI have suggested a useful term that will be borrowed here to identify them collectively: *robotic others* (Kahn 545-46).

At this time of writing – early 2006 – the reality of HRI with robotic others appears to be most advanced in Japan. Anthony Faiola has supplied *The Washington Post* with the following rather striking report of conditions last year in that nation:

Analysts say Japan is leading the world in rolling out a new generation of consumer robots. [...] Though perhaps years away in the United States, this long-awaited as-seen-on-TV world [...] is already unfolding in Japan, with robots now used as receptionists, night watchmen, hospital workers, guides, pets and more. [...] “We have reached the point in Japan of a major breakthrough in the use of robot technology and our society is changing as a result,” said Kazuya Abe, a top official at NEDO, the national institute in charge of coordinating science research and development. “People are and will be living alongside robots, which are seen here as more

than just machines. This is all about AI” – artificial intelligence, Abe said – “about the creation of something that is not human, but can be a complement or companion to humans in society. That future is happening here now.” (2)

According to Faiola, part of this future that is happening “now” in Japan includes also a robotic baby seal (worth \$10 million in development grants from the Japanese government) that “is meant to provide therapy for the elderly who are filling Japanese nursing homes at an alarming rate while often falling prey to depression and loneliness” (3).

Obviously, it is no recent discovery that humans interact with – and are affected by – their tools and other technological artifacts. The fairly new class of robotic others just described, however, appears to warrant special attention. Writing in a 2004 issue of *Human-Computer Interaction*, for example, researchers Sara Kiesler and Pamela Hinds observe that “people seem to perceive autonomous robots differently than they do most other computer technologies” (3). In particular, they point out “When we build autonomous robots to look human, we may encourage anthropomorphic mental models of these systems” (3). Please notice that the suggestion, here, is focused simply upon *how humans perceive* robotic others – whether, or to what extent, such perceptions are *technically justified* is a separate issue. The prior topic, concerning *human responses*, has not gone unnoticed, even in the technical community of artificial intelligence and robotics research. Rodney Brooks illustrates such awareness in his 2002 book, *Flesh and Machines: How Robots Will Change Us*, voicing important questions of the following kind:

Is there, or will there ever be, enough similarity to us in our humanoid robots that we will decide to treat them in the same moral ways we treat other people and, in varying degrees, animals? (154)

Nor has the significance of our human responses to robotic others been neglected in the religious community;

theologian Philip Hefner illustrates this point with the following representative comment in *Technology and Human Becoming*:

[Alan] Turing and his colleagues and their descendants have created a significant mirror of ourselves. What we want and who we are coalesce in this mirror. Whether we think this mirror is adequate or not is another question. (31)

Indeed, there appears to be recognition, in domains that span both science and religion, that human responses to *robotic others* may present some important implications for *people* – implications that even might include altered understandings of ourselves.

The ways in which people respond to robotic others, however, can be conditioned by cultural factors – and, yes, this broad rubric is intended here to include differing religious beliefs. Anthony Faiola, within the report mentioned previously, acknowledges influence of this type:

“In Western countries, humanoid robots are still not very accepted, but they are in Japan,” said Norihiro Hagita, director of the ATR Intelligent Robotics and Communication Laboratories [...]. “One reason is religion. In Japanese (Shinto) religion, we believe that all things have gods within them. But in Western countries, most people believe in only one God. For us, however, a robot can have an energy all its own.” (3)

Influences from Buddhist thinking might also contribute to the apparent cultural disposition of the Japanese people to perceive robotic others differently from ways expected in Western countries, such as the United States. Machiko Kasuhara, in an essay titled “The Century of Household Robots,” suggests “The relationship between human beings and other creatures is different [in Japan] from that in Europe, mainly because of historically different religious backgrounds” (2). In particular, he cites a lesser sense of difference “between the lives of human beings and other animals in Buddhist theory” and claims that this Buddhist perspective has strongly influenced the “Japanese way of thinking” (2). In any event, it seems reasonable to bear in mind that any judgments regarding “human” responses to robotic others may need to take into account possible cultural and religious effects.

The conspicuous emergence specifically *in Japan* of an HRI phenomenon involving robotic others most likely reflects a number of other factors characterizing that particular country; e.g., demographic, economic, and political variables. MIT’s Rodney Brooks notes generous support by the Japanese government for development of

“robots that can interact with people in close quarters,” which he associates with several conditions peculiar to that nation; for example, Japan simultaneously has an aging and ethnically homogeneous population, a low birth rate, and traditional cultural resistance to use of imported labor (135-7). Although this profile is in some respects different from that of the United States, it also displays elements that arguably could promote spread of the HRI phenomenon to the U.S. Martha Pollack, in a recent issue of *AI Magazine*, observes that the cohort of people aged 65 and older in the U.S. also is increasing steadily (10). Moreover, she reports that the U.S. federal government pays nearly 60 percent of a “\$132 billion annual nursing home bill,” adding that “technology that can help seniors live at home longer provides a ‘win-win’ effect, both improving quality of life and potentially saving enormous amounts of money” (9). In the presence of such factors, it should be rash to assume that the U.S. could not experience significant future growth in the HRI phenomenon we have been considering.

Although similar observations might equally apply to a number of other Western nations, the author of the present essay – as a citizen and lifetime resident of the United States – believes it would be appropriate to focus the following discussion principally upon features of the setting for HRI in the U.S.

The Cultural Setting for HRI Growth in the United States

Diversity historically has been a prominent feature of the U.S. Accordingly, the country harbors a broad range of distinct philosophical and religious positions that are embedded in quite different worldviews. In particular, the following remarks briefly will reconnoiter (a) views expressed by individuals closely associated with U.S. robotics research and (b) ideas that one might typically expect to encounter in religious communities of the U.S.

Profile (a)

Carnegie Mellon University roboticist Hans Moravec gives our reconnoitering a brisk start by offering an admirably straightforward criterion for determining whether a robot possesses a conscious soul:

... we might grant a conscious soul to a robot by interpreting its behavior as expressing the action of such a soul: the more humanlike its interaction with us, the easier the attribution.” (76)

If this seems a bit quick, rest assured that Moravec is willing to take a vote on the matter:

... So, it may be appropriate to say “God” has granted a soul to a machine when the machine is accepted as a real person by a wide human community. (77)

These are strong claims, and they clearly are pertinent to dispositions one might have to ascribe moral status to a robotic other. Indeed, I have argued elsewhere that traditional religious and legal conditions for moral status normally *do* require conscious awareness – and that at least a representational model of *self*-awareness plausibly could be realized in artificially intelligent machines (Metzler 41). Nevertheless, software recursively representing itself – nay, even itself in relations with representations of an external environment – is strictly not yet a license to claim solution for the mystery of what philosophers have called “raw feels” or *qualia*. Having actually worked extensively with patients suffering various impairments of consciousness, neurologist Antonio Damasio is quite careful in his articulation of this distinction:

The “looks” of emotion can be simulated, but what feelings feel like cannot be duplicated in silicon. Feelings cannot be duplicated unless flesh is duplicated, unless the brain’s actions on flesh are duplicated, unless the brain’s sensing of flesh after it has been acted upon by the brain is duplicated. (314-15)

In contrast, the position expressed in Hans Moravec’s comments bears the marks of a much bolder Turing Machine functionalism. Theologian Philip Hefner has furnished an especially concise summary of the latter philosophical position:

... it is not the biology or the physics of the brain that is critical for what it does, but rather the *logical structure of its activities*. Therefore, those activities can be represented in any medium that replicates that structure of logic, including machines. (30)

A significantly different voice, though, can be heard from another quarter of the U.S. artificial intelligence community that is somewhat less directly involved in robotics research. Although he clearly is also no stranger to the technical methods that enable the growing phenomenon of robotic others, Ray Kurzweil (*The Age of Spiritual Machines*) adds an important nuance to his account of robotic potential:

Just being – experiencing, being conscious – is spiritual, and reflects the essence of spirituality. Machines, derived from human thinking and

surpassing humans in their capacity for experience, will claim to be conscious, and thus to be spiritual. They will believe that they are conscious. They will believe that they have spiritual experiences. They will be convinced that these experiences are meaningful. And given the historical inclination of the human race to anthropomorphize the phenomena we encounter, and the persuasiveness of the machines, we’re likely to believe them when they tell us this. (153)

Immediately, one senses in Kurzweil’s remarks a reluctance to issue strong ontological claims about consciousness that appear to conflate third-person observations of behavior with first-person experiences of awareness. Elsewhere, in fact, he explicitly acknowledges that “there’s just no way to truly penetrate another entity’s subjective experience,” adding “that’s where philosophy and religion are supposed to take over” (154). Kurzweil still draws attention, however, to a point of special concern in this essay – viz., *regardless* of what the ontological “ground truth” might be regarding their conscious awareness, sufficiently humanlike robotic others are likely to be *accepted and treated* by humans as peers.

If one’s understanding of consciousness already resembles Turing Machine functionalism, then any such transition to regarding robotic others as peers will most likely *feel* like it involves more continuity than change. After all, if conscious awareness can be identified as a special type of computation then the robotic other plausibly could be well on its way to meeting even one of the necessary conditions for moral status. Nevertheless, there still may be some difficult changes buried in this process – particularly, what about the requirement for *freedom*? Besides consciousness, aren’t “free” moral choices commonly counted as another *necessary* condition for moral responsibility?

Well ... if one happens to be a roboticist who really wishes to build a robotic moral agent, it should be sufficient to recall some definitions from that Introductory Philosophy course. According to a so-called *compatibilist* definition for “freedom,” your robot need only be capable of conducting some appropriate internal “deliberation” regarding each choice, and able to reach a decision that depends upon *no external coercion*. The internal deliberation might be realized with multiple AI methods – the compatibilist definition for freedom offers one an enormous design space.

In sum, accepting robotic others – even as peers with *moral* status – appears to represent a change for which U.S.

technical spokespersons such as Hans Moravec and Rodney Brooks (perhaps even Ray Kurzweil) might *readily* find resources to enable a graceful transition.

Profile (b)

When we turn to reconnoitering ideas one might encounter in religious communities of the United States, the nation's hallmark cultural diversity must be acknowledged at once. No attempt will be made in the following, therefore, to canvass all possible religious perspectives; rather, attention will be focused upon some specific ways in which religious understandings of the human person and of moral agency that are represented in this country might bring resources to the HRI phenomenon much *different from* those most accessible to technical communities associated with robotics research.

We might properly begin, though, by acknowledging that technical and religious communities are by no means incapable of various kinds of overlap. For example, robotics theologian Anne Foerst – explaining, in *God in the Machine*, “... I am influenced in my language by my stay at MIT” – consequently prefers to substitute “self-awareness” for the word “consciousness” (96). Nonetheless, she is able to conclude her book with a call for a world “in which a peaceful coexistence of all different forms of culture and creed, and of all different humans – *and our robotic children* – becomes possible” [emphasis added] (190). It is not the aim in this essay, therefore, to argue that perspectives of technical and religious communities always are – or must always remain – mutually exclusive. Rather, the following observations are intended only to highlight specific points at which some careful dialogue and honest intellectual work might be required to negotiate changes needed to integrate certain received religious views of morality with any notable expansion of HRI phenomena in the United States.

First, not all religious understandings of moral agency are satisfied by the *compatibilist* definition for “freedom.” U.S. theologian Nancy Murphy and South African cosmologist George F.R. Ellis (*On the Moral Nature of the Universe*) call repeatedly for something more subtle than the compatibilist account of freedom as a necessary condition for authentic moral choice; for example:

A moral universe requires much more than lawlike regularity, however; it is necessary that the laws and regularities within the physical universe allow for the existence of *intelligent* beings who can sense and react in a conscious way and who have *free will*. We

here touch on issues that science has not seriously begun to comprehend: we do not understand the nature of consciousness nor the “free will” we experience, [...] . We assume freedom of action, albeit constrained by many biological, psychological, and social factors, for without this the concept of morality is meaningless. (207)

Of course, so-called *libertarian* (or *categorical*) accounts of freedom – which would insist that not all moral deliberation is comprised of totally deterministic processes – can come packaged with philosophical difficulties of their own. Murphy and Ellis, for example, also caution explicitly against a glib enlistment of quantum mechanics that would claim “some random event at the quantum level actually makes up my mind for me” (36). Indeed, it is not a promising approach to rest moral responsibility on a coin toss. Perhaps a robotic other in the form of a stochastic automaton *could* eventually manage to yield convincingly humanlike moral choices (and even be acceptable as a genuine moral agent, in principle, among compatibilists) – but there currently remains a philosophical and theological culture that would yet challenge its moral status on grounds of an inadequate kind of “freedom.” Called upon to produce a full account of their more adequate alternative, the challengers might be likely to exercise a Kantian move, and to postulate freedom as an inscrutable primitive. This is not the place, however, to continue tracing these well-worn philosophical paths of dispute; it is sufficient for present purposes to point out that there are communities of opinion regarding freedom in the U.S. that can be expected to resist any serious encouragement from HRI experiences to ascribe moral status to robotic others.

In addition, some religious perspectives resist ready ascription of *consciousness* to robotic others. The present author repeatedly has encountered students voicing observations of the general form “Robots can't have consciousness; they don't have souls.” Recalling Jewish philosopher Martin Buber's celebrated work, *I and Thou*, it is tempting to interpret such comments as indicating dispositions to reserve “I-Thou” relations for humans, and to engage robotic others in terms of the “I-It” relation. To the extent that this interpretation is sound, a number of Martin Buber's pronouncements appear unlikely to be compatible with expanded HRI fostering acceptance of robotic others as peers; for example:

If a man lets it have the mastery, the continually growing world of It overruns him and robs him of the reality of his own I, till the incubus over him and

the ghost within him whisper to one another the confession of their non-salvation. (46)

Another spiritual leader who enjoys considerable respect in the United States is His Holiness the Dalai Lama, who – besides being an individual we can safely assume has devoted some serious attention to the subject of consciousness – is known especially for his openness to dialogue with the sciences. However, he recently has recorded some observations *about* consciousness (*The Universe in a Single Atom*) that some scientists might regard as calls for some fairly uncomfortable alterations of their own methods:

[...] it is, I think, clear that the third-person method – which has served science so well in so many areas – is inadequate to the explanation of consciousness. What is required, if science is successfully to probe the nature of consciousness, is nothing short of a paradigm shift. That is, the third-person perspective, which can measure phenomena from the point of view of an independent observer, must be integrated with a first-person perspective, which will allow the incorporation of subjectivity and the qualities that characterize the experience of consciousness. I am suggesting the need for the method of our investigation to be appropriate to the object of inquiry. (133-34)

Indeed, what is called for here might be characterized more properly as The Mother of All Paradigm Shifts! Consciousness, after all, is that unique private locus through which each of us directly experiences his or her intentionality and purposeful action. And modern science, ever since it began to take form in the seventeenth century, has progressively eliminated recognition of *precisely* that subjective experience of purposeful action (aka final causality) from its scope of concerns. The Dalai Lama's observation is both profound and potentially relevant to the concerns we have just heard expressed by Martin Buber. The "I" of the *I-Thou* primary word bears meaning for many people that simply is not captured adequately in third-person descriptions of "patterns," neural or otherwise.

Some Results from Scientific Research Concerning Human Responses to HRI Experience

It has been estimated that empirical HRI research is a fairly recent development, not much more than a decade old (Keisler 4). In the following we shall focus upon reports of some studies concerning human interactions with

Sony's robotic dog AIBO that have been published within the past two years.

The first of these reports was submitted to the 2004 IEEE International Workshop on Robot and Human Interactive Communication, bearing the title "Social and Moral Relationships with Robotic Others?" (Kahn, et al.). Immediately, at least one of the included quotations of human subjects' accounts of their interactions with AIBO (posted to Internet discussion) is sufficiently revealing, for the purposes of the present essay, to warrant unedited presentation:

The other day I proved to myself that I do indeed treat him as if he were alive, because I was getting changed to go out, and tba [AIBO] was in the room, but before I got changed I stuck him in a corner so he didn't see me! Now, I'm not some socially introvert guy-in-a-shell, but it just felt funny having him there! (547)

If one finds in these remarks a somewhat surprising depth of apparent social bonding with (fairly unsophisticated) robotic artifacts, the authors furnish numerous illustrations of similar human responses. For example, 42% of the human subjects reportedly "spoke of AIBO having *intentions*," 38% "spoke of AIBO having *feelings*," and 39% "spoke of AIBO as being *capable of being raised, developing, and maturing*" (547). In this study, however, indications of human attribution of (what the authors term) "moral standing" to the robotic dogs did not seem equally impressive; again, their own account deserves to be heard as written:

One of the most striking results in our AIBO Discussion Forum Study was that while AIBO evoked conceptions of life-like essences, mental states, and social rapport, it seldom evoked conceptions of moral standing. Members seldom wrote that AIBO had rights (e.g., the right not to be harmed or abused), or that AIBO merited respect, deserved attention, or could be held accountable for its actions (e.g., knocking over a glass of water). In this way, the relationship members had with their AIBOs was remarkable one-sided. They could lavish affection on AIBO, feel companionship, and potentially garner some of the other psychological benefits of being in the company of a pet. But since the owners also knew that AIBO was a technological artifact, they could ignore it whenever it was convenient or desirable. (549)

Thus far, could one characterize the reported human interactions with AIBO as suggesting I-Thou or I-It relations? The empirical evidence, it seems, displays some ambiguity. We must acknowledge, further, that the authors of the study also explicitly note a relationship that “seems of a strange hybrid unidirectional form, where the human is able ultimately to control or at least ignore the robotic other with social and moral impunity,” suggesting some questions about possible transfer effects upon human-human relations (549).

In 2005, additional research involving human interaction with the robotic dog AIBO yielded reports of results similar to, as well as different from, the foregoing. A study involving interaction of 80 preschool children with AIBO, as well as a stuffed dog, was followed by an investigation of interactions of 72 children (ages 7 to 15), AIBO, and a *live* dog. In the latter study, the majority of the children reportedly “affirmed that AIBO had mental states (56%), sociality (70%), and *moral standing* (76%)” [emphasis added] (Melson 1651). Although evidence of dispositions to grant AIBO moral standing appear to be something of a mixed bag, authors of the 2005 report believe we are warranted in expecting that “as robots become increasingly life-like in their behavior, so will people treat them as if they were mental, social, and moral beings – thus raising robotic others toward the level of biological others” (Melson 1652). On the other hand, the authors’ interpretations also include further recognition of some ambiguity in human attitudes toward the robotic others. In particular, they suggest that HRI might be challenging “traditional ontological categories,” observing that a 7-year-old boy who was asked whether AIBO was alive replied “Um, he’s kind of alive, he’s kind of both” (Melson 1652).

From the vantage point of the present essay the ambiguities revealed in these studies are potentially important, for they may also be read as evidence of plasticity in the ways in which humans may come to regard increasingly lifelike (and humanlike) robotic others. Even more important within the present context, perhaps, are the possibilities for change in the ways in which people come to regard and understand *themselves* as a result HRI experiences. In any case, the results of these studies – together with the foregoing profiles of moral notions – furnish a useful background for identifying specific effects upon moral thinking and behavior that an expanded HRI phenomenon in this country could be expected to foster.

Suggested Moral Implications of HRI Growth in the United States

Acknowledging that most of the variables in the discussion that follows cannot be assigned sufficiently precise values to support crisp logical arguments, we might structure an account of possibilities by imagining – at the top of the hierarchy – people choosing either to reject or to accept robotic others as peers. For each disjunct in this choice, ensuing moral implications can be expected to differ as a function of whether the chooser approaches HRI experience with cultural-religious profiles described under (a) or under (b) in the foregoing discussion. Profile (a), for example, will afford philosophical and/or theological assumptions that encourage ready attribution of consciousness to robotic others – and most likely an understanding of “freedom” that supplies the artifacts with enough compatibilist resources for moral status. In contrast, profile (b) choosers will come to the decision with a skeptical view of such philosophical positions as Turing Machine functionalism, and a firm insistence upon categorical freedom as a necessary condition for any authentic moral responsibility. In terms of this choice structure, we may begin with the cases in which – despite all the evidence of an apparent human predisposition to accept humanlike entities socially – the choosers are unable to embrace the robotic other as a moral peer.

One might reasonably expect the class of individuals fitting this description to consist predominantly of profile (b) choosers. Profile (a) individuals, after all, should be fairly easily satisfied – particularly, if the robotic other could at least show some evidence of being capable of familiar “moral deliberation.” Even exceptional “outliers,” such as Anne Foerst, should be pleased – indeed, it is she who (referring to the “social” MIT robot, Kismet) has suggested the robotic other as a possible vehicle for overcoming our countless human forms of exclusivity:

[...] humanoid robots such as Kismet will become a definite part of the community of persons. If they don’t, it means that we will also exclude human beings. Discussing the question of Kismet’s personhood can therefore be used as a corrective that helps us to test our own myths of personhood for their inclusivity. In this sense, Kismet is our destiny, because it can help us to turn this world into a better and more accepting place. (189)

For the more numerous profile (b) choosers in this class, however, some painful changes could be in store. One of the more deleterious possibilities, in fact, has already been

suggested in the 2004 report of HRI research by Peter Kahn and his colleagues. If, as the research seems to indicate, we can readily develop various social bonds with artifacts that display sufficiently lifelike (or, ideally, humanlike) behavior, then that level of social acceptance might plausibly encourage the human chooser progressively to categorize other humans *with the robotic others*. To the extent that such identification takes place, might the next step – allowing one morally to ignore selected *humans* “whenever it is convenient or desirable” – become a rather small step? Fortunately, the conjecture being offered here seems, in principle, to be one that could be investigated scientifically.

The class of humans capable of *accepting* robotic others as peers should predominantly be comprised – as suggested previously – of profile (a) individuals. For the (presumably rare) profile (b) humans who *could* join this class, there might – again – be problematic changes ahead. In this case, however, the “danger” would arise from the possibility that acceptance was ambiguous, and that subsequent reflection on apparent inability of robotic artifacts to *satisfy* the preferred categorical notion of freedom might reflexively erode the human individual’s perception of her or his *own* moral status – “Don’t blame me, I’m just a robot,”

Conclusion

Broadly, the dynamic of possibilities appears to resemble the form of a dilemma for profile (b) individuals: for them, either rejection or acceptance of robotic others may generate difficult and unwelcome moral changes. If this analysis is correct then these individuals in the United States, under conditions of significantly increased HRI experiences with robotic others, may be expected to respond in either of at least two ways. First, they may choose to change substantively their understandings of moral agency, and to adopt the perspectives of profile (a). Second, they may choose, instead, to defend continuity of their positions and to call for changes in the thinking of their neighbors representing profile (b). In either case we appear to be facing an increasing need for dialogue characterized by fairness, humility, and intellectual rigor among the communities of science, technology, and religion. To the extent that such dialogue is successful, we might seriously hope to begin filling the prescription written years ago by Alfred North Whitehead for “a deeper religion and a more subtle science” (185).

References

- Brooks, R. 2002. *Flesh and Machines: How Robots Will Change Us*. New York: Vintage Books.
- Buber, M. 1958. *I and Thou*. Trans. Ronald Gregor Smith. New York: Collier Books.
- Damasio, A. 1999. *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. New York: Harcourt Brace & Company.
- Faiola, A. “Japan embraces new generation of robots.” *The Washington Post* 12 Mar. 2005. 31 Oct. 2005. <<http://www.detnews.com/2005/technology/0503/12/tech-114940.htm>>.
- Foerst, A. 2004. *God in the Machine: What Robots Teach Us About Humanity and God*. New York: Dutton.
- Hefner, P. 2003. *Technology and Human Becoming*. Minneapolis: Fortress Press.
- His Holiness the Dalai Lama. 2005. *The Universe in a Single Atom*. New York: Morgan Road Books.
- Kahn, P. Jr., N. Freier, B. Friedman, R. Severson, and E. Feldman. 2004. “Social and Moral Relationships with Robotic Others?” *Proceedings of the 2004 IEEE International Workshop on Robot and Human Interactive Communication*. 545-50.
- Kiesler, S., and P. Hinds. 2004. “Introduction to this Special Issue on Human-Robot Interaction.” *Human-Computer Interaction* 19.1: 1-8.
- Kusahara, M. 2005. “The Century of Household Robots.” 31 Oct. 2005. <<http://www.f.waseda.jp/kusahara/robot.html>>.
- Kurzweil, R. 2000. *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. New York: Penguin Books.
- Melson, G., P. Kahn, Jr., A. Beck, B. Friedman, T. Roberts, and E. Garrett. “Robots as Dogs? – Children’s Interactions with the Robotic Dog AIBO and a Live Australian Shepherd.” *CHI 2005*. ACM : 1649-52.

Metzler, T. 2000. "And the Robot Asked 'What do you say I am?': Can Artificial Intelligence Help Theologians and Scientists Understand Free Moral Agency?" *The Journal of Faith and Science Exchange* IV: 37-48.

Moravec, H. 1999. *Robot: Mere Machine to Transcendent Mind*. New York: Oxford University Press.

Murphy, N., and G. Ellis. 1996. *On the Moral Nature of the Universe: Theology, Cosmology, and Ethics*. Minneapolis: Fortress Press.

Pollack, M. 2005. "Intelligent Technology for an Aging Population: The Use of AI to Assist Elders with Cognitive Impairment." *AI Magazine* 26.2 : 9-24.

Whitehead, A. 1967. *Science and the Modern World*. New York: Free Press.