ROUTLEDGE STUDIES IN SCIENCE, TECHNOLOGY AND SOCIETY

The Good Life in a Technological Age

Edited by
Philip Brey, Adam Briggle,
and Edward Spence

Contents

PART V
Medical and Agricultural Technology and Well-Being

PART VI
Technology Design and Policy

20 Care Robots, Virtual Virtue, and the Best Possible Life

Mark Coeckelbergh

Growth of the elderly population and shortage of caregivers place increased pressure on health care systems in North America, Europe, and Japan. Aging rapidly, these societies are on their way to what the U.S. President's Council on Bioethics called 'mass geriatric' societies with 'more long-lived individuals than ever before in human history' (President's Council on Bioethics 2005, p. xvii). The problem is not so much the absolute numbers of elderly people but the ratio between older and younger people. This poses a problem for elderly care. Of course many of us welcome a longer life span. But although we age, we do not always age healthy: the 'oldest' old (80+, 85+) often suffer from chronic diseases, such as dementia, before they die. Hence, there is a need for more caregivers for long-term care. But it is unlikely that children take up these tasks:

> precisely because many individuals have taken advantage of modern freedom's opportunities for education, careers, and geographic mobility, many elderly persons will live in greater isolation from loved ones, separated from children and grandchildren who have settled elsewhere or whose lives are defined primarily by work and school. (President's Council on Bioethics, p. 2)

At the same time, state funding—which pays for professional caregivers—is hardly rising. Thus, we need more caregivers, while their availability decreases for various reasons (see also Lynn 2004). Hence, we find ourselves 'on the threshold of a crisis in long-term care' (President's Council on Bioethics 2005, p. ix–x).

If this analysis is right, then one possible response is the following future scenario: let *robots* take over care tasks from humans. First, we delegated care from family, neighbors, and friends to professionals and institutions. This professionalization of care created all kinds of problems (Illich 1976) as well as benefits. But if this solution is inadequate in the light of the demographical and political situation, then why not delegate care to robots?

This scenario is not as exotic as it may appear at first sight. It is seriously considered or even promoted by some governments (e.g., Japan)

and seems to be in line with developments in other fields. We increasingly delegate tasks to automated systems in all kinds of domains of life, from coffee machines to autopilot systems. Medicine and health care is no exception; although machines usually assist rather than *replace* human care. Human lives depend on heart-lung machines, electronic monitoring devices, pace makers, robotic surgery, and other systems. Moreover, in response to social-economic problems with the current health care systems, some suggest that we should use electronic monitoring and care equipment to enable people to live longer in their own homes and to save hospital costs.

Is this care robot scenario desirable? First, we need to specify what kind of robots and care tasks we are talking about. Care robots could appear human in some way (humanoid robots) or look and act like a pet (pet robots). As *personal* robots they would "share physical and emotional spaces with the user" (Arras and Cerqui 2005, p. 13). Care tasks would range from therapy to entertainment and companionship: robots could assist ill and elderly people by monitoring them, delivering drugs, moving them around, helping them with domestic tasks, entertaining them, talking to them, or keeping them company.

In robotics, concrete steps are being taken toward the development and use of such personal or 'social' robots, as well as less interactive assistive robots in health care. Consider Paro (a baby seal therapeutic robot), My Spoon (a feeding robot), or HAL (assistive limb suit), and ongoing research projects such as the Companionable project,[1] which aims to enhance the independence and quality of life of elderly and disabled people. Therapeutic and companion robots could be regarded as more social and personal than an assistive limb suit or feeding robot, which, in turn, are more social and personal than, for example, a cleaning robot or a car manufacturing robot.

This chapter will focus on intelligent, autonomous, and interactive robots that take on 'personal' or 'social' dimension and perform the previously mentioned therapy, monitoring, entertainment, and companionship functions. Let me also assume for the sake of argument that in these scenarios, care robots would *replace* real humans or pets.

The replacement assumption invites one of the primary ethical worries: is it right that robots would replace humans in care? In the robot ethics literature, this is called the *replaceability* question (Decker 2006, 2008). There are various approaches to the question. For instance, Decker has analyzed the question from a Kantian perspective; and Robert Sparrow and Linda Sparrow have understood the problem as being about "genuine social interaction" and respecting people as "ends in themselves" (Sparrow and Sparrow 2006, p. 149).

In this chapter, I reformulate the research question concerning evaluation of the 'care robot' in terms of the good life: Can interaction with care robots contribute to the good life, and what is the difference with human care?

THE GOOD LIFE AND THE BEST POSSIBLE LIFE

What is 'the good life'? There are many theories of well-being (Crisp 2008). Usually, the term 'good life' or the 'flourishing life' refers to Aristotle's view that we are to realize our potential and our aim (*telos*) that lies in our nature as humans. His theory, which gave rise to the virtue ethics tradition in moral philosophy, can be understood as implying that one of the things we need for a good life is friendship. I will explore the applicability of the Aristotelian notion of complete friendship, which he thought was constitutive of the good life, and consider some alternative models of interaction in which human good can appear.

But first we should ask what the good life means given the specific situation of ill and elderly people. Human flourishing, *the* good life, seems easier when one is young, healthy, and in touch with one's nearest and dearest. But these conditions are not always fulfilled. Ill and elderly people cannot aspire to the same level of human flourishing as others and have specific needs and vulnerabilities. It is more appropriate, then, to speak about *the best possible life* (BPL) in their situation.

On reflection, however, this term seems appropriate with regard to all of us. How absolute can the criteria for human flourishing be? Are they relative to one's personal situation and one's unique individual character, or are there needs and aspirations shared by all humans? Let me respond to this familiar problem in ethical theory by distinguishing between three equally important 'good life' questions that must be raised to evaluate care robots and other technology, related to three kinds of norms: one concerns general criteria that apply to all humans (criteria for *the* good life, that is, the best possible life conceivable without taking into account contextual or other limitations: the good life for a healthy human being able to use and develop all her abilities), one concerns universal criteria that apply to the particular situation, context, and technology (criteria for the *BPL* given that situation), and one concerns criteria that are unique to the person and that person's interaction with the technology (criteria for the life of person X). This approach amounts to what I coin a *vertical moral pluralism* of norms: there are different kinds of moral norms that range from general to personal levels of application, but none of them takes priority by definition.

Of course, such an approach raises the question how to balance the different criteria in case of tension between them. I believe this is a question that is difficult (if not impossible) to solve *a priori*, disregarding the particular situation. Moreover, this scheme is not by itself able to guide an ethics of (elderly) care; we need to specify the criteria and what they mean for elderly people in specific situations. In the remainder of the chapter, I shall start with Aristotle's account of friendship (formulated at the most general level) as a way to flesh out the BPL criterion at a less general (but not personal) level: the BPL for elderly people that find themselves in a situation of needing care (that is, more and perhaps to some extent different care than

a young and healthy person needs) and of receiving professional care (in an institution or in their own home).

FRIENDSHIP

What kind of interaction between humans is constitutive of the good life, that is, makes us flourish as humans? In contrast to human rights theory or the capabilities approach, the good life approach does not ask about *minimum* criteria for human life, but asks about how to *maximize* the quality of that life. Part of Aristotle's answer to that question is his account of friendship. In the *Nicomachean Ethics,* he argues that the most worthy relationship between people is what he calls 'perfect friendship,' which requires that one desires the other for the sake of the other and one desires the good for the other (book 8, p. 1156b and p. 1157b25; see also Badhwar 1987). It is a relationship between moral equals (p. 1156b); we acknowledge and enjoy virtue in the other. In imperfect friendship, by contrast, the other is a source of advantage or pleasure (p. 1156a15–20). It is an *instrumental* relationship.

Perhaps this notion of friendship could guide nonprofessional, noninstitutional care by family and friends, provided that there is a relation between moral equals (this is far from obvious: our contemporary understanding of friendship seems closer to pleasure-based friendship than to moral equality). However, this Aristotelian criterion is certainly far too demanding to guide relations between care professionals and care receivers. Professional caregivers and care receivers (and maybe nonprofessional caregivers also) need not be friends in this Aristotelian sense to achieve a good care relation which contributes to the BPL. Yet we do not want to resort to 'minimum' morality embodied in ethical codes, contracts, and care protocols; we need a concept that refers to 'warmer' and less instrumental relationships that concerns the good rather than (only) the (legal or deontological) right.

A further problem is that Aristotle's theory of friendship has little to say on human-technology relationship. Using the example of a bottle of wine, he argues that the notion of friendship is not applicable to lifeless objects because we do not wish good to them for their own sake but only for our own use (we want the bottle of wine to remain good for us to drink it), whereas to friends we wish good because of themselves (book 8, p. 1155b30). Therefore, *insofar as* robots count as 'lifeless objects,' Aristotelian theory is not very helpful. Within its framework it is inappropriate to talk of moral symmetry between humans and robots, or to discuss a robot's contribution to the good life. (Unless, of course, robots are *not* viewed as lifeless objects. This would trigger further discussion of 'friendship' with robots. But even if we create this opening, it seems unlikely that the relation could be defined in terms of moral equality as opposed to use or pleasure. Generally, we tend to regard robots as a kind of slaves. This excludes the

possibility of perfect Aristotelian friendship. Or is there another possibility? I will discuss this further in my section of 'virtual virtue'.) Moreover, given its historical roots in the culture of the ancient Greek city-state, Aristotle's theory appears to be addressed to active, rational, independent, and healthy male citizens of high social rank, whereas care receivers embody dependency, vulnerability, inactivity, and nonparticipation in the public life. Perhaps they even lost their rational capacity, for example, when they suffer from dementia.

COMPANIONSHIP

Some of these problems could be avoided by appealing to *companionship* rather than Aristotelian friendship. This notion has something of the 'warm' connotation of friendship, but is less demanding. It does not require moral symmetry, without supposing a totally instrumental relationship. The aim is social: being together, doing things together. The (virtue of the) other is not the end, although a somewhat Aristotelian argument can be made that there is a noninstrumental dimension to the extent that we are social beings and that therefore being together realizes that *social* dimension of what we already are, the social *telos* of humans. We are social by nature. Furthermore, the term lacks the erotic element: desire for the other is not needed; desire for company suffices. Nevertheless, the end of companionship cannot be reduced to pleasure or utility: companionship is 'more' to the extent that it realizes the social end of humans.

Given these features, we can apply the term to relations between humans and nonhumans, such as 'social' animals (pets) and 'social' robots. People sharing their physical and emotional space with pets or personal robots do not typically experience their relation to the pet robot or care robot in instrumental terms. Rather, empirical research (e.g., Turkle 2005, see following) suggests that these nonhumans *appear* as their companions and are treated and interacted with accordingly. At the same time, humans—if mentally healthy—are aware of the 'real' asymmetry. They usually know that the robot is not really human or not a biological pet. Therefore, the contribution of such interactions to human flourishing as the BPL must be judged on the basis of what people experience. Appearance turns out to be ethically relevant. If people *experience* companionship, then this contributes to the BPL in the given situation and personal circumstances. I will expand this argument in the next sections.

VIRTUAL VIRTUE

However, if appearance matters so much to the good life, can we push this argument further and consider the possibility of moral symmetry in

human-robot relations? Are robots necessarily like Aristotle's bottle of wine? They sometimes appear to us as being 'more' than lifeless objects. Perhaps an advanced humanoid robot could appear to us as a morally *equal* companion, despite us knowing that it is a robot. In that case, we would ascribe 'virtue' to the robot on the basis of its appearance. A robot with *virtual* virtue—whatever its designed or intended functions may be—we would call a 'good' robot; and human good can emerge in relations with such a robot. A robot with virtual virtue, therefore, could contribute to the BPL in a way that differs qualitatively from 'mere' companionship. It would provide virtual friendship and thereby potentially contribute more to the BPL of people.

The ascription of virtual virtue is less odd than it may seem. After all, we are used to virtual characters in novels, films, and computer games. And very little is necessary for us to ascribe a personality to an object such as a computer (Reeves and Nass 1996). Moreover, we judge human virtue on the basis of habits, that is, on the basis of consistently good behavior. If robots acquired 'good habits,' then by the same moral-social 'mechanism,' we would probably ascribe virtue to them.

Of course, Aristotle was not a behaviorist and he would require that virtuous actions are done for the right reasons. This means that in addition to observation of virtuous behavior of the robot, the ascription of virtual virtue to the robot would require that there is also the appearance of reasoning and intention. It requires that we 'forget' that the robot is following an algorithm, that we 'forget' that it is a machine. On the part of the designer, it demands the creation of the *illusion* of virtue and its requirements in terms of behavior and virtual motivation. Is this deception? And if so, is it morally wrong?

Using the term virtual virtue enables us to redeem the Aristotelian notion of friendship in the context of human-care robot relations, at least if we are content with *virtual* friendship with care robots and disregard potential moral concerns this concept might raise. (Next I will discuss the deception objection.)

CARESHIP

But why take great pains to apply notions such as friendship and companionship to health care when we already have the concept of *care* itself? Using that concept, we could try to develop criteria specific to the kind of life ill or elderly people lead: criteria specific to those who do not meet the Aristotelian assumptions regarding autonomy and independence and can only attain the *best possible* life in their vulnerable situation. Such a BPL ethics would require that ill and elderly people enjoy interaction with those who care for them or who appeal to their capacity to care. They participate in, and contribute to, a care relation. Here 'care' does not have the narrow, more instrumental meaning of 'medical care' or 'health care,' and it is

not required that one desire the other (as in Aristotle's perfect friendship); rather, one desires *the best* for the other in the given situation of illness or old age and given that person's unique character and identity. I believe this criterion of 'careship' can be met by others than friends and relatives only, but is more demanding than what protocols and codes require from health professionals. (Hence, the criterion is in tune with the intuition that professional care is only a 'second best' option—the best option being that relatives, friends, and others who care *about* the person are the caregivers.)

However, while this criterion for achieving the BPL or model of care could in principle be met by health professionals if certain contextual conditions are fulfilled (e.g., sufficient staff and time available, good intentions, the possibility to develop a care relationship, etc.), it cannot be met by robots. If robots cannot be conscious or cannot have (real) emotions, they can 'deliver' care but not 'really' *care* about their patients. Only if we remove the reality demand, can we create an opening: care robots could *appear* to care. But this invites the deception objection: It would amount to making the person believe that the robot really cares, while in reality it does not.

This objection relies on making an ontological distinction between appearance and reality, an assumption that may be contested (for example—but not exclusively—by behaviorists). But even if one accepts the distinction, one might argue in response to the objection that apparent care does not count *as deception* if care receivers are aware that the robot is a robot. They would *know* this but relate to the robot on the basis of its *apparent* care and this could contribute to their BPL. At most, there is temporary *self-deception*, but one that enhances the BPL. Personal robots are never 'mere' robots in this experiential, perceptual sense. They are what Sherry Turkle, Will Taggart, and others call 'relational artefacts' (Taggart, Turkle, and Kidd 2005). In their study of human-robot interaction in nursing homes, a pet robot looking like a baby seal was treated by elderly people in a nursing home *as if* it needed nurturing. This suggests that if people *feel* that the robot needs their care or that the robot cares about them, this helps them achieve their BPL, depending on their situation and character. The emphasis on the moral significance of feeling and appearance lends support to the claim that both human care and robot care should be adapted to the specific situation ('case') of the person and to the human-robot relation that emerges between the unique, individual person and the robot. For example, Turkle's research shows that people respond very differently to the same robot according to their character. The stress on appearance renders this care ethics of BPL sensitive to individual and contextual differences.

THE DECEPTION OBJECTION

Nevertheless, it seems that there is something morally problematic here. Does virtual friendship and virtual care respect people? Sparrow and

Sparrow admit that people can feel happy when interacting and forming 'relationships' with robots, the authors take issue with the cause, which they call *delusion*:

> In most cases, when people feel happy, it will be because they (mistakenly) believe that the robot has properties which it does not (. . .) It is these delusions that cause people to feel loved or cared for by robots and thus to experience the benefits of being cared for. (Sparrow and Sparrow 2006, p. 155)

Why is this problematic? According to the authors, failure to apprehend the world accurately is a moral failure. Moreover, they claim that our well-being is not served by 'illusions:'

> What most of us want out of life is to be loved and cared for, and to have friends and companions, not merely to *believe* that we are loved and cared for, and to *believe* that we have friends and companions, when in fact these beliefs are false. (Sparrow and Sparrow 2006, p. 155)

The authors conclude that using robots in this way would express a gross lack of respect for older people.

I agree that *if* virtual friendship and virtual care were to constitute a delusion, the use of robots for this purpose would be morally wrong. However, provided that the elderly people or patients in question are mentally healthy, the authors mis-describe what happens in virtual care. Just as in interactions with virtual agents in a computer-generated virtual world, most mentally healthy people have a twofold awareness: They are aware of the virtual other, but they are also aware that this other is not real. If we really wish to take seriously ill and elderly people as autonomous beings, which the authors want, then we should not assume that they live in delusion. Rather, we should assume that they can enjoy interaction with their robot nurse and robot companion without believing that they are interacting with a human being.

Of course, some people can no longer distinguish between the real and the virtual. But this implies they *already* live in 'delusion' or illusion (for example, they may mistake the *human* nurse for someone else or mistake their spouse for someone else), and it is unclear why care robots would worsen their condition or not respect them.

Moreover, in *human* health care such people are *already* treated as non-autonomous persons and for a good reason: they *lack* the capacity for autonomy. But that does not mean that we necessarily disrespect them. Paternalism can be justified under certain conditions. And respect can also be granted on the basis of recognizing a person as human or on the basis of the particular relation we have to that person. The questions whether or

not care robots can help mentally confused elderly people and how we can best respect such people and enhance their well-being are different questions that must be separated from the deception issue.

Sparrow and Sparrow are right to point to subconscious processes that go on when humans interact with robots. But surely these processes do not stop when we interact with humans. In human-human interaction too, we can be deceived about the other—and worse: *by* the other. 'Worse' because I assume that intended deception is morally worse than nonintended deception, which may not always be the case.[2]

Finally, human-human interaction often involves nonintended deception and partly *depends* on appearance in the sense that as social beings we always play roles to *some* extent (roles of a nurse, a doctor, a partner, a friend, etc.) and it is difficult to make a strict distinction between appearance and reality in human social life. At least it is not *obvious* what reality is as opposed to appearance (e.g., what is the real self or real person?). Hence, given the role of appearance in human-human relations, including care relations, putting all emphasis on the deception objection when evaluating care robots is misplaced. As it happens, thinking about human-robot relations and human-robot care is a lens through which we can look afresh at what we mean by good human care and relations. Perhaps there is something dubious about the role appearance plays in human-robot care relations, but it is not obvious why it is morally wrong *per se*.

However, there may be other, better objections against delegating care to robots not discussed in this chapter.[3] Moreover, is the problem well-defined? Are robots the problem?

NON-TECHNOLOGICAL SOLUTIONS? REFORMULATING THE PROBLEM

Why consider robots at all in response to problems in health care and elderly care? Discussing a technological solution may bypass reflection on the causes of the problem as outlined in the beginning of this chapter. Must we not ask why there is a shortage of care workers? Must we not ask why we should have professional care at all instead of care by relatives and friends? How could our society end up with this problem in the first place?

Once we pursue this line of inquiry, it is tempting to move immediately to a normative view. For instance, with Ivan Illich, one may argue against the professionalization and institutionalization of modern health care by referring to the ethic of the 'Good Samaritan' (responding to people, feeling dis-ease at the sight of suffering, caring reciprocity as opposed to following rules or conventions) and to more communal forms of living together (see, for example, Illich and Cayley 2005; Kahn 2009).

I hesitate to embrace this normative view without further qualification. A weak reply to the 'Samaritan' objection is that we cannot turn back

time—that we have to start from (late) modern society as it is. This is a weak reply because however wide the gap between norm and practice is, philosophers have a duty to imagine and discuss normative ideals. A stronger reply is that both the objection and the weak reply assume (1) that we already have, and agree about, a conception of good care and good community, and (2) that there was a time when that conception of good care and community was realized. But we must doubt if the past was really better, be critical of the current health care system, and scrutinize conceptions of care that allow us to evaluate that system.

The 'Good Samaritan' model of care is yet another model that, like the Aristotelian, is not simply applicable to robot care or to current professionalized care. It is based on the model of reciprocal love in *human* relations, an ideal that is close to the Aristotelian model of perfect friendship. Robots can be programmed to help any human person in need, regardless of how they otherwise appear, and thus act 'impartially' as good rule followers,[4] but they cannot love. They cannot really feel the suffering of the care receiver. However, this, by itself, is not an argument against using robots in health care. Moreover, professionalization and the preservation of systematic knowledge do not only have disadvantages. The intuitive 'Samaritan' response to the care problem invites us to further reflect on what good care is, what a good community is, what love is, and what the differences are between human-robot relations and human-human relations. Therefore, normative evaluation of robots in health care and elderly care should involve more conceptual and empirical analysis of the care problem and of the moral ideals that might help us to create better care futures. This includes being more critical of our current norms and practices that do not involve robots. Philosophy of technology is not (only) about technology, but about us.

CONCLUSION

I started this chapter with a methodological proposal: evaluate care robots in terms of the good life. In particular, we should shift our attention to criteria for BPL for unique and vulnerable people in particular situations. I have discussed friendship, companionship, and careship as models for evaluating BPL in relation to human-robot interaction. More empirical research is needed to evaluate how people in health care contexts could attain the BPL. This notion of BPL could also be used to evaluate other human practices.

But what is my normative conclusion? There is little doubt that robots could contribute to better lives for ill and elderly people. Therefore, there is nothing intrinsically wrong with giving them *a* place within health care. However, there is disagreement about what place that should be if our aim is the BPL of ill and elderly people—that is, the BPL of all elderly people, not just those who happen to like robots. At first sight, it seems unlikely and undesirable that robots would replace human caregivers. They cannot

be our 'Aristotelian' friends and they cannot *really* care about us. A way to get around these problems is to lower the Aristotelian standard and take seriously what I have called 'virtual virtue' and 'virtual care.' In virtual friendship and virtual care, human good could emerge that is experientially similar in form to human friendship and human care. If this happens, it contributes to people's BPL, does not necessarily violate the moral requirement that people are not deceived, and if it did, this would be less problematic than appeared at first sight given the mental condition of the people in question and given the importance of appearance in human social life.

In the current technological-social situation, the conditions for this kind of good to arise are seldom fulfilled. Robots cannot produce much 'virtual virtue' yet, and our current health care systems do not stimulate the development of virtual virtue—or real virtue for that matter. To the extent that these conditions are not fulfilled, therefore, we should make sure that the use of care robots does not lead to a reduction of human contact, as Sparrow and Sparrow rightly fear (Sparrow and Sparrow 2006, p. 152), and use robots only as a supplement to human caregiving, not as its replacement.

However, if robots cannot always provide what ill and elderly people need to reach BPL, it must be noted that professional caregivers cannot necessarily do this either. This is not just a matter of individual ethics; we should assume that health care professionals usually have good intentions. But good intention is not enough. As previously suggested, current health care systems may not stimulate the kind of relations we hold most worthy or most conducive to BPL. Therefore, a comprehensive discussion about using robots in health care should involve a reassessment of the way we deal with the ill and elderly in our society. For instance, in contrast to most robots, humans *can* intentionally deceive. If they are human, we want caregivers to *really* care. If our current social institutions do not support that kind of care, we should not blame the robots but imagine better, *best possible* ways of organizing care and society.

NOTES

1. This is an acronym for 'Integrated Cognitive Assistive and Domotic Companion Robotic Systems for Ability and Security.' The project is funded by the EU 7th Framework Programme. The robotic system would be able to recognize emotional states of people, distinguish between normal and exceptional behavior, remind people to take their medication, and so on.
2. Sometimes deception by robotic and other caregivers may be required (Wagner and Arkin 2010).
3. For example, Adam Briggle thinks delegating care to robots is like the ostrich sticking its head in the sand, because it gives us the illusion that we live in a world without ageing. In other words, we deceive ourselves if we let technology take away the 'burden' of care. (Compare this situation also to the initial life of the Buddha before his departure: He lived in the illusion of a world without suffering.)

4. Note that even robotic rule-following is not straightforward: as Asimov's stories informed by the 'Laws of Robotics' show, rules may conflict.

REFERENCES

Aristotle. (1908). *Nicomachean Ethics* (William David Ross, trans.). Oxford: Clarendon Press.

Arras Kai O. and Daniela Cerqui. (2005). "Do we want to share our lives and bodies with robots? A 2000 people survey," Technical Report Nr. 0605–001, Autonomous Systems Lab, Swiss Federal Institute of Technology Lausanne (EPFL).

Badhwar, Neera Kapur. (1987). "Friends as Ends in Themselves," *Philosophy and Phenomenological Research*, vol. 48, pp. 1–23.

Crisp, Roger. (2008, December). "Well-Being." In *Stanford Encyclopedia of Philosophy*. Stanford: Stanford University Press. Retrieved September 21, 2010, from http://plato.stanford.edu/entries/well-being/Decker, Michael. (2006). "Replaceability of Humans by Autonomous Robots? Ethical Reflection within the Framework of an Interdisciplinary Technology Assessment." Paper presented at the ETHICBOTS workshop, Naples, Italy, October 18, 2006.

Decker, Michael. (2008). "Caregiving Robots and Ethical Reflection: The Perspective of Interdisciplinary Technology Assessment," *AI & Society*, vol. 22, pp. 315–330.

Illich, Ivan. (1976). *Medical Nemesis: The Expropriation of Health*. New York: Random House.

Illich, Ivan, and David Cayley. (2005). *The Rivers North of the Future: The Testament of Ivan Ilich*. Toronto: Anansi.

Kahn, Richard. (2009). "Critical Pedagogy: Taking the Illich Turn," *The International Journal of Illich Studies*, vol. 1, pp. 37–49.

Lynn, Joanne. (2004). *Sick to Death and Not Going to Take it Anymore!* Berkeley: University of California Press.

President's Council on Bioethics. (2005). *Taking Care: Ethical Caregiving in our Aging Society*. Washington, DC: U.S. Printing Press.

Reeves, Byron, and Clifford Nass. (1996). *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge: Cambridge University Press.

Sparrow, Robert, and Linda Sparrow. (2006). "In the Hands of Machines? The Future of Aged Care," *Minds and Machines*, vol. 16, pp. 141–161.Taggart, Will, Sherry, Turkle, and Cory D. Kidd. (2005). "An Interactive Robot in a Nursing Home: Preliminary Remarks." Paper presented at the 27th Annual Conference of the Cognitive Science Society Android Science workshop, Stresa, Italy, July 2005.

Turkle, Sherry. (2005). "Relational Artefacts/Children/Elders: The Complexities of CyberCompanions." Paper presented at the 27th Annual Conference of the Cognitive Science Society Android Science workshop, Stresa, Italy, July 2005.

United Nations Programme on Ageing. (2010). "The Ageing of the World's Population," Retrieved November 25, 2010, from http://www.un.org/esa/population/publications/worldageing19502050/

Wagner, Alan R., and Ronald C. Arkin. (2010). "Acting Deceptively: Providing Robots with the Capacity for Deception," *International Journal of Social Robotics*, vol. 3, pp. 5–26.