

Health Care, Capabilities, and AI Assistive Technologies

Mark Coeckelbergh

Accepted: 4 June 2009 / Published online: 17 July 2009
© Springer Science + Business Media B.V. 2009

Abstract Scenarios involving the introduction of artificially intelligent (AI) assistive technologies in health care practices raise several ethical issues. In this paper, I discuss four objections to introducing AI assistive technologies in health care practices as replacements of human care. I analyse them as demands for felt care, good care, private care, and real care. I argue that although these objections cannot stand as good reasons for a general and a priori rejection of AI assistive technologies as such or as replacements of human care, they demand us to clarify what is at stake, to develop more comprehensive criteria for good care, and to rethink existing practices of care. In response to these challenges, I propose a (modified) capabilities approach to care and emphasize the inherent social dimension of care. I also discuss the demand for real care by introducing the ‘Care Experience Machine’ thought experiment. I conclude that if we set the standards of care too high when evaluating the introduction of AI assistive technologies in health care, we have to reject many of our existing, low-tech health care practices.

Keywords Health care · Capabilities · Good life ethics · AI · Assistive technology · Robots

1 Introduction

Scenarios involving the introduction of artificially intelligent (AI) assistive technologies in health care practices raise several ethical issues. Consider plans to use carebots that would replace nurses, intelligent assistive technologies such as artificial limbs, and intelligent ambient technology built into houses. With regard to ethical reflection on these proposals, a distinction can be drawn between, on the one hand, concerns about the *replacement* of human care by AI assistive technologies and, on the other hand, concerns about care *assisted* by AI technologies (as opposed to care not assisted by such technologies) but without replacing human care. The latter is far less controversial than the former. It appears that AI technologies could contribute to health care in useful and important ways. For

M. Coeckelbergh (✉)
Department of Philosophy, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands
e-mail: m.coeckelbergh@utwente.nl

instance, a robot could help to lift people, which is a heavy job for nurses. An artificial limb could help people to walk (again). By monitoring people in their own homes, it would allow people to stay and feel at home while receiving care. However, once the issue of replacement comes in, there is more controversy. A recurrent objection—for instance in the discussion about care robots¹—is that care provided by AI systems is not as good as human care. In other words, the objection is put in terms of a question of *replaceability*. For instance, Michael Decker has discussed the question of care-giving robots in these terms, using the Kantian argument about means and ends to cast doubt on whether such robots could replace humans (Decker 2008). And Robert and Linda Sparrow have argued against care robots for the reason that some needs can only be met by humans (Sparrow and Sparrow 2006). Moreover, there are worries about privacy and about ‘fooling people’ when such technologies would be introduced in care practices as replacements of human care. But are these arguments sufficient grounds for rejecting the use of AI assistive technologies in care as replacements of human care as such?

In this paper, I focus on the replaceability issue but my arguments are also relevant to concerns about AI technologies that are ‘merely’ assistive. I discuss four potential objections to introducing AI assistive technologies in health care practices as replacements of human care. First, an AI system such as an AI robot or an AI monitoring system is able to deliver care, but it will never really *care* about the human. Second, if AI technologies are able to provide care at all, they cannot provide *good* care since good care necessarily requires contact with humans since we have social and emotional needs. Third, AI assistive technologies may be able to provide good care, but in doing so violate the fundamental principle of privacy, which is why they should be banned. Fourth, AI assistive technologies such as robots are provide ‘fake’ care: they are likely to ‘fool’ people by making people think they receive real care.

I will argue that although these objections cannot stand as reasons for a general rejection of AI assistive technologies (as replacements of human care or as assisting human care), they demand us to develop more comprehensive criteria for good care and rethink existing practices of care. In particular, I will propose a capabilities approach to care and emphasize the inherent social dimension of care. I will also argue that if we set the standards of care too high when evaluating the introduction of AI assistive technologies in health care, we have to reject many of our existing, low-tech practices of health care.

2 Deep Care

The first objection could be put in terms of a deep/shallow contrast. AI assistive technologies that replace human care are supposed to provide ‘shallow’ care only, since they do not really care about the patient. What is lacking is the kind of ‘deep’ feelings that accompany human care. Moreover, one may object that what is required is not only feeling, but reciprocity of feeling: it involves an emotional exchange. Now a defender of AI care as replacement of human care has at least two ways of answering this objection. First, she might argue that although such AI systems do not have feelings now, they may be designed to have them in the future and this will also make possible reciprocity of feeling. Although this is what many technology gurus would say, I believe this scenario is highly unlikely and in any case does not provide much guidance for thinking about care in the *near* future.

¹ Since this is the discussion I know best, I will mainly refer to the literature on care and robots. However, the arguments presented in this paper are relevant to other AI assistive technologies as well.

Second, a better response, therefore, is to note that the ‘deep care’ asked for (care as feeling and as reciprocity of feeling) is not always and not necessarily part of ‘low-tech’ *human* care as it is organized today. In the context of mass care and bureaucratic organisations, human care practices usually lack ‘deep’ care. Much care work is routine work. Furthermore, there is little time for emotional, intimate, and personal engagement with the patient; the ratio patients/carers is too high. And if there is time, such engagement is not always deemed to be appropriate to the extent that the professionalisation of care requires keeping some emotional distance. Moreover, a too strong emotional engagement may hurt when the person cared for is no longer in your care (e.g. because the person is healthy again, has been moved to another care institution, or has died).² Although these are not *necessary* features of care, they are unavoidable within the current way we mainly organise care as a society: as professionalized, bureaucratically controlled mass care. Thus, if ‘deep’ care is what we want, we need not blame AI assistive technologies, but rethink and reorganise our present practices of care.

Now one might argue that although human care is presently not always deep, at least it *could* be improved, whereas with AI technology this is not the case. But this objection misses the possibility that replacement with regard to one care task may actually make room for more and deeper human care with regard to other care activities. Indeed, there is an argument to be made *for* using AI technologies as replacements with regard to some human tasks: they could take over routine tasks, so that more time is available for personal and emotional engagement with the patient. When such routine tasks are done by humans, and without any personal or emotional engagement, this is perhaps particularly undignifying and unethical. Arguably, to be treated like a thing by a machine is less morally degrading than to be treated like a thing by a human being.

3 Good Care

Let us suppose that health care aims at providing good care. To what extent does good care depend on humans? Many of us would say that it very much depends on humans. Robert and Linda Sparrow, for instance, have argued against the replacement of human nurses by robots in elderly care for the reason that robots are incapable of meeting the social and emotional needs of elderly persons, which can only be done by means of contact with humans (Sparrow and Sparrow 2006). Their main concern is, rightly so, the quality of care. But what is good care, and what is the place of social and emotional needs in it?

Both in health care ethics and in ethics of technology, little work has been done in providing systematic and comprehensive criteria of good care. Usually the discussions in medical ethics and health care ethics are based on one or a few principles such as autonomy, consent, and privacy.³ For instance, a classical text in medical ethics defines four basic principles: autonomy, beneficence, non-maleficence, and justice (Beauchamp and Childress 1994). Such principles are useful to exclude some morally unacceptable practices, but they provide only limited guidance when it comes to giving a positive definition of

² Note that it also hurts to see a person’s health gradually diminish. And it is also not always easy, when dealing with a terminally ill patient, for instance, to take and show the ‘right’ emotional attitude towards that person. What does a particular person at a given moment in time need most? Compassion? Encouragement? Which feelings should I show? Should I talk or should I listen?

³ The latter principles may also be understood as requirements that flow from the principle of (respect for) autonomy.

good care.⁴ Compare them with moral principles such as the biblical 10 Commandments: they may forbid certain actions, but they do not tell us what the good life is. But a broader, positive conception of ethics is needed if we wish to decide about whether or not to introduce AI assistive technologies. Sparrow and Sparrow's point is not that care robots violate the elderly person's autonomy and consent, or that they do not benefit the person. They claim that robots cannot provide the same kind of care as humans do. But this raises the question what kind of care that is. Let me, therefore, take a broad, 'good life' ethical approach to good care, rather than limiting myself to the more restricted scope of negative morality. Moreover, it would be great if we had principles that are related in some way, rather than a list of what might appear as an arbitrary selection.

I believe we can achieve this by taking a capabilities approach to health care. There has already been some attention for the relation between capabilities and health care, for instance in the work of Anand (2005) and Ruger (2006), but these authors are more concerned with macro-issues of justice and social choice mainly relevant to economical and health policy decisions; in this paper I reflect on how to evaluate concrete health care practices in a way that could guide health care professionals.

The capabilities approach has been developed by Sen and Nussbaum to evaluate well-being in terms of what people are actually able to do rather than the resources they have (Nussbaum and Sen 1993; Nussbaum 2000). Nussbaum's recent version of the capabilities approach is based on the principle of human dignity (Nussbaum 2006). Let me try to apply that approach to health care.

Arguably, part of what health care should aim at, is respecting, promoting and preserving the dignity of the patients. Both professional and non-professional carers, then, have a corresponding ethical obligation to respect, promote, and preserve the dignity of those who are in their care. Moreover, we can rightly expect that health care institutions and health care systems are set up to achieve the same aim. However, put in terms of dignity alone, this demand provides as little guidance as the principle of autonomy or justice. We want to know what dignity means. I propose to add that respecting the dignity of humans means to treat and respect them (1) as humans, (2) as humans belonging to a particular community and social-cultural context, and (3) as the unique persons they are. However, this formulation is still too vague. We need more specific, positive criteria that allow us to evaluate the quality of care.

Here Nussbaum's capabilities list can help out. Claiming that they are founded on the principle of human dignity, Nussbaum has drafted a list of ten capabilities 'as central requirements of a life with dignity'; they are 'general goals that can be further specified by the society in question' but an 'appropriate threshold level' needs to be reached (Nussbaum 2006, p. 75). Important for Nussbaum is that all capabilities matter; we need to enjoy all of them to live a life with dignity. This is my summary of Nussbaum's list, which I drafted on the basis of the version articulated in *Frontiers of Justice* (Nussbaum 2006). The list includes the following 'central human capabilities':

1. *life*: 'Being able to live to the end of a human life of normal length; not dying prematurely, or before one's life is so reduced as to be not worth living.'
2. *bodily health* (includes nourishment and shelter)
3. *bodily integrity*: free movement, freedom from sexual assault and violence, having opportunities for sexual satisfaction

⁴ Although the principles of beneficence and justice seem to be positive principles, they are commonly used in a negative way.

4. being able to use your *senses, imagination, and thought*; experiencing and producing culture, freedom of expression and freedom of religion
5. *emotions*: being able to have attachments to things and people
6. *practical reason*: being able to form a conception of the good and engage in critical reflection about the planning of one's life
7. *affiliation*: being able to live with and toward others, imagine the other, and respect the other
8. *other species*: being able to live with concern to animals, plants and nature
9. *play*: being able to laugh, to play, to enjoy recreational activities
10. *control over one's environment*: political choice and participation, being able to hold property, being able to work as a human being in mutual recognition

Nussbaum (2006, pp. 76–78)

This list of capabilities can be used as a list of criteria to evaluate health care and the use of AI technology in health care. If good care is care that respects human dignity, and if the principle of human dignity requires that the listed human capabilities be restored, maintained, and perhaps enhanced,⁵ then we have a list of criteria health care practices should meet – with or without technology. For instance, the list includes the social and emotional aspects Sparrow and Sparrow emphasized, but many other dimensions of care and the good life as well. And indeed, it may well turn out that for certain care tasks, a particular AI assistive technology is not able to restore, maintain or enhance some capabilities as well as humans can do. But this has to be decided on a case by case basis; we cannot reject them a priori and in general, as Sparrow and Sparrow do. With regard to some criteria, some care tasks, and in specific situations some AI technologies may be able to replace human care; with regard to other criteria, other technologies, and other situations they can assist human care without being able to replace it.

Let me make some further observations about these criteria for good health care.

First, note that this is a rather expansive view of what health care should do. Someone might object to this and argue (1) that health care should be concerned with health alone and (2) that health has to do with the well-functioning or performance of an organism – that is, the human body. For instance, she might adopt Leon Kass's definition of health (Kass 1985). Such definitions of health are indifferent to conceptions of what a good life consists in.

Second, one might object that such an expansive view of health care puts too heavy obligations on health care professionals. But this remark neglects my claim made above that health care is not just the responsibility of professionals. If health care and the good life are related in the way suggested above, then what makes me healthy or unhealthy has to do with others, with the social environment, with the society I live in. Responsibilities for my health should be distributed accordingly. Having said this, as with the demand for 'deep' care, the demand for good care can still ask too much of those who are immediately concerned with the health of the patient. Therefore, I propose to set the threshold level not so high that no-one could ever achieve it. But the level that will be achievable in practice, depends not only on individual capacities of care professionals or relatives (or, for that matter, on what assistive technologies can do), but also on how we organise care and our society.

⁵ I am aware that to add 'enhancement' to this list is very controversial and there are serious difficulties with defining what enhancement means. Therefore, I leave 'enhancement' out of my discussion in this paper.

Third, it is important to be clear on what such principles can do and cannot do for reasoning about ethics of health care. The issue concerning the role of principles is a topic on its own and I have discussed it to some extent elsewhere (Coeckelbergh 2007). Let it suffice for me to remark here that the principles for good care I derived from Nussbaum, as any ethical or moral principles, do not necessarily *settle* difficult cases or *solve* hard problems in health care practice. One lesson we can learn from pragmatism and reflective equilibrium ethical theory, is that the role of principles only comes in when there is a problem, and that their role must be located within a larger picture of moral practice and moral deliberation. As only *one* of the elements, principles help us to decide what to do but do not determine the outcome of our ethical reflection. There are also the particular case, context, and other principles that need to be taken into consideration. There are the views of various stakeholders involved in the practice, and there is a role for moral imagination when we try to apply the principles to particular cases and attempt to find a new action option when faced with a dilemma. With regard to my proposed account of dignity-promoting health care, this implies that each criterion cannot settle but rather inform and guide moral deliberation and evaluation in particular cases and with particular practices. The proposed criteria can play this role with regard to three questions that I believe moral deliberation and evaluation must answer in health care ethics cases. First, what does criterion X mean in this particular case? Second, is the criterion satisfied? And third, has the criterion been *sufficiently* met? (With regard to the latter question, the threshold level comes in.) Although moral reasoning requires us to take some distance from our practices and our opinions, these questions cannot be answered from a ‘point of nowhere’. It is up to care givers, care receivers, and other stakeholders to decide in practice what it takes to live in dignity. A capabilities approach can help them to reflect on this question.

4 Private Care

A typical objection to employing AI assistive technologies in health care as replacements of human care is that they violate the privacy of the patient. If I have an AI monitoring system in my home, then I might feel that ‘Big Brother’ is watching me, to use an Orwellian metaphor. But how appropriate is this worry, given that current care practices involve the continuous ‘violation of privacy’? If a human nurse washes me in the hospital, how ‘private’ is that? If my medical data are stored in a database that is not under my control, how ‘private’ are they? This is not to say that the concern for privacy is unwarranted, but rather that the privacy issue is not new or unique to the introduction of AI assistive technologies. Furthermore, the principle of privacy needs to be balanced against other principles. Let me give an example using the capabilities list. If the technology restores my communication with others (partner, family, friends, but also medical professionals), if it allows me to participate in the community, and if I am in such a condition that without constant monitoring of my bodily functions, my life expectancy decreases dramatically, then it is not clear why privacy should be the sole or overriding principle. Thus, even as a replacement AI technology may promote the social life of the patient and constitutes one element in a health care system that itself has a social character (depends on various kinds of relations between people) and is not guided by privacy alone. I grant that privacy is *one* of the principles that should guide the design and use of AI assistive technologies; but it is not the only one or necessarily the most important one. Finally, although it is hard to predict the exact outcomes once a technology is introduced, as designers and as users we have some control over what we want a technology to do or not to do. If we think privacy is

important, we can take this into account in the design, use, and regulation of AI technologies.

5 Real Care

A particular concern voiced by Sparrow is that assistive technologies such as robot carers and robot pets are *simulacra*: they are substitutes for the real – the human or the biological pet (Sparrow 2002; Sparrow and Sparrow 2006). The worry is that people are ‘fooled’ when they are given AI assistive systems that resemble the biological ones. I think this argument is not a very strong one, since in practice most of the time people are very much aware that a certain AI autonomous system such as a robot is not really human, even if the robot has a human appearance and even if they respond to the robot as if it were human.⁶ But what if we really would mistake the AI technology for what it is not? What if we are really ‘fooled’? To try to understand this objection, let me develop the following thought experiment: The Care Experience Machine. The idea is inspired by Nozick’s thought experiment ‘The Experience Machine’. Here is Nozick’s idea:

Suppose there were an experience machine that would give you any experience you desired. Superduper neuropsychologists could stimulate your brain so that you would think and feel you were writing a great novel, or making a friend, or reading an interesting book. All the time you would be floating in a tank, with electrodes attached to your brain. Should you plug into this machine for life, preprogramming your life’s desires? [...] Would you plug in? What else can matter to us, other than how our lives feel from the inside? (Nozick 1974, p. 43)

Now suppose there were a “Care Experience Machine” that would give you the experience of receiving all the care you need, without the interventions of humans. Would it be morally wrong to plug patients into this machine? Most of us would object to this proposal, and I sympathize with such a response; I share the intuition. But it is not clear why exactly it would be wrong. Why is the fact that it would be *virtual* care, not real care, a morally relevant distinction?

Let me discuss the thought experiment I introduced. First, the Care Experience Machine thought experiment constitutes a valid objection against *subjectivist* accounts of care, who only consider the subjective experiences of people. Indeed, care should not only make people *feel* that they are cared for; care should actually provide care. But, partly for that reason, the thought experiment is not an objection against an *objectivist* account of care such as the capabilities approach to health care I proposed above. What matters, in that account, is the restoration and maintenance of capabilities. Of course these capabilities mean nothing if they are not experienced by the person in question. But the emphasis is not on what people feel (e.g. if they feel happy) but on the capabilities they actually have as humans. Feelings of happiness, in this objectivist view, is not within the domain of ethics. The ancient Greek term *eudemonia* is not translated as happiness here, but as the good life or human flourishing. And leading the good life does not necessarily make you

⁶ See for instance experiments by Ishiguro and others with the ‘android father’: as far as eye movements of the child go, Ishiguro and others found that the child responds to the android father as if it were her real father - *knowing*, however, that it is not the real father. For what the presence of a robot does to children see for example the experiments by Nishio et al. (2007).

(subjectively) happy all the time. To require this from any account of the good life would render the account irrelevant to human life.

But let me press on with the reality objection. How real must the care be? What if the Care Experience Machine gives you the experience of restored or enhanced *capabilities*? Does it matter, morally speaking, that you are living the *virtual* good life? If AI assistive technologies constitute a ‘good demon’ (*eudemonia* in another sense) that gives you that experience of the good life, what is wrong with it?

I see two options here. Either we say that virtual care is wrong, since it needs to restore *real* capabilities. This is the easy way out taken above, but it adds another requirement to the capabilities approach. Or we admit that virtual good care is good (and that, by extension, the virtual good life is good).

I prefer the first option, but let me make the following qualification. Even if the goal is the restoration of real capabilities, this does not exclude virtual experiences as a means to achieve these real capabilities. For instance, if such experiences involved someone playing virtual reality games that restored her capabilities (for instance, brought her into contact with others or enhanced some brain functions), then leaving aside other possible objections there is nothing wrong with using that technology as an aid within a care practice that aims at capability restoration. Thus, although there might be something wrong with full replacement, it seems that care assisted by virtual reality generating technology is acceptable provided it helps people to achieve real capabilities.

Note that if we really believe that (1) care has to do with the restoration of capabilities and (2) that *humans* are needed not only to restore some capabilities, but also to *experience* exercising these capabilities, then of course the Care Experience Machine is theoretically impossible. But even if we endorse these beliefs, the Care Experience Machine thought experiment helps us to fine-tune and further develop the intuitions we already had: the capabilities approach to health care is an objectivist account, but experiencing having and exercising the capabilities must be part of that account. Thus, the capabilities approach could be improved by putting more emphasis on the ‘insider’ experience of people and less on the need policy makers, managers and other ‘outsiders’ who want an instrument to measure, from the outside, the quality of people’s lives and the quality of care. The question why the real/virtual distinction must be relevant for ethics, however, remains a hard one.

I suspect that one of the issues that make us worry here is not only the real/virtual or real/fake problem, but also the well-known problem of paternalism. Is it right that another should decide for the patient about her treatment? The concern may be that an objectivist account of health care, as any objectivist account, could be used to justify doing things to people ‘in their best interest’ but without their consent. The worry has to do with the principle of autonomy: we want patients to be able to make their own choices about their care. My brief response to this objection is that the principles provided by my capabilities approach to health care should be balanced with other ethical principles, and that in some cases paternalism can be justified. For instance, if a patient lacked the capability to plan her own life, then it can be justified for others to decide about her life if this contributes to, and does not conflict with, the aims of health care: her dignity and the restoration of her capabilities. I will not provide a more elaborate argument for this view here, but it is clear that the issue of paternalism is one of the moral stakes in the discussion.

Note that the demands for deep care as felt care and for real care and experienced capability restoration show another possible weakness of a Nussbaumian account of good care. I find Nussbaum’s capabilities list a particularly useful instrument to articulate what is at stake in health care ethics. However, as it stands my view might be interpreted as being too much directed towards the *outcome* of the care process rather than the care *process* itself

(to the extent that these are different). It might give the impression that health care is only about preparing people for life afterwards, outside the context of health care. This impression is mistaken. The aims of care, if understood in terms of improving capabilities, are partly realised in the care process itself. They should not be disconnected from the enjoyment of capabilities *in* the care process and from anything else that renders the care process itself good. In this respect, it is also recommendable to take into account that when we demand felt care we often demand reciprocity of feeling. Nussbaum's list includes emotions and attachments to others but does not explicitly acknowledge this concern. Although reciprocity of feeling may not be required for all forms of care (for instance current AI assisted care), it is part of our emotional and social capabilities as humans and it is therefore relevant to evaluating care when and in so far as care involves and aims at the exercise and enjoyment of these capabilities. Finally, although objectivist in its definition of the aim of health care and in defining central human capabilities, the account – as it applies to both care givers and care receivers – must be sensitive to personal and cultural differences in experience of care. Different people from different cultural contexts may respond differently to AI technology that is intended to assist care.

6 Conclusion

I have discussed several ethical issues raised by scenarios concerning the introduction of AI assistive technologies in health care practices as replacements of human care. I have shown that many of these objections are not very strong objections against employing these technologies as replacements of human care. Rather, these objections force us to develop more systematic and comprehensive criteria for good care and rethink our existing practices of care. Responding to worries about deep and good care, I proposed a capabilities approach to care. I noted that the role of such principles must be understood as one element within a broader moral deliberation and evaluation. I also argued that we should not set the standards of care too high when evaluating the introduction of AI assistive technologies in health care, since otherwise we would have to reject many of our existing, low-tech health care practices. And if we think our high standards are right, then perhaps we should change these practices. Furthermore, in response to worries about privacy, I emphasized the inherent social dimension of care. Moreover, I discussed problems concerning reality and noted the problem of paternalism. In particular, with the help of the Care Experience Machine thought experiment I clarified the intuitions that the restoration of capabilities must be real and must be experienced and enjoyed by the person in question. I left open whether or not the reality demand can be justified; perhaps we can do no more than simply adding a reality requirement to the capabilities account of good care. I also suggested that in certain cases, it can be justified that the decision to restore a person's capabilities is made by others. Finally, I concluded that a Nussbaumian account of good health care should not be interpreted as being exclusively concerned with the outcomes of care. We must also pay attention to the care process if and in so far as that differs from the outcomes: to feeling and the possibility of reciprocity of feeling, to the (real) experience and enjoyment of capability restoration in the care process, and to personal and cultural differences in care experiences.

More work needs to be done in clarifying the relation between feelings and care, between (objective) capabilities and (subjective) experience, between the demand for privacy and the social character of care, between the real and the virtual good life, and between outcome (aim) and process (not a mere means to reach the aim). We also need more detailed discussions about which particular AI technologies could replace which care

tasks in which health care context. But here is an approach to health care that improves the definition of what is at stake in the ethical discussion concerning AI assistive technologies in health care. The problem is not the technologies themselves, not replaceability as such, and not the (potential violation of) the principles of privacy or autonomy alone, but the question what good care and the good life is: for us as humans, for us in this context, and for us as the unique persons that we are. I proposed a (modified) capabilities approach to health care as one way to articulate what this ethical requirement could mean at the level of moral theory and moral principles.

Acknowledgments Thanks to Nicole Vincent, Nicholas Munn, Aimee van Wynsberghe, and other participants of the *International Applied Ethics Conference 2008* (Hokkaido University, Sapporo, Japan), the January 2009 research seminar of the Philosophy section at Delft University of Technology, and the Good Life meetings at the Philosophy Department of Twente University for the discussions we had about robots and care. I also wish to thank the anonymous reviewers for their helpful comments, which improved the quality of my arguments.

References

- Anand P (2005) Capabilities and health. *J Med Ethics* 31:299–303
- Beauchamp T, Childress IF (1994) *Principles in biomedical ethics*, 4th edn. Oxford University Press, New York
- Coeckelbergh M (2007) *Imagination and principles*. Palgrave Macmillan, Basingstoke/New York
- Decker M (2008) Caregiving robots and ethical reflection: the perspective of interdisciplinary technology assessment. *AI & Soc* 22(3):315–330
- Kass L (1985) The end of medicine and the pursuit of health. In: Kass L (ed) *Toward a more natural science*. The Free Press, New York, pp 157–186
- Nishio S, Ighiguro H, Hagita N (2007) Can a teleoperated android represent personal experience? a case study with children. *Psychologia* 50(4):330–342
- Nozick R (1974) *Anarchy, state, and utopia*. Basic Books, New York
- Nussbaum MC (2000) *Women and human development: the capabilities approach*. Cambridge University Press, Cambridge
- Nussbaum MC (2006) *Frontiers of justice: disability, nationality, species membership*. The Belknap Press of Harvard University Press, Cambridge M.A. and London
- Nussbaum MC, Sen A (eds) (1993) *The quality of life*. Clarendon Press, Oxford
- Ruger JP (2006) Health, capability, and justice: toward a new paradigm of health ethics, policy and law. *Cornell J Law Public Policy* 15(2):403–482
- Sparrow R (2002) The march of the robot dogs. *Ethics Inf Technol* 4:305–318
- Sparrow R, Sparrow L (2006) In the hands of machines? The future of aged care. *Mind Mach* 16(2):141–161