Ethical and social aspects of cognitive technologies in health and social care.

Chair's report on a parallel session at the EUCogIII Members' Meeting, University of Sussex, 23 October 2013

Steve Torrance
University of Sussex.

Note

The EUCogIII parallel session on Ethical and Social Aspects of Cognitive Technologies in Health and Social Care started with a 15 minute context-setting talk by Mark Coeckelbergh, followed by a five minute response by myself, and then around 70 - 80 minutes' discussion, with some brief concluding points from the opening speakers. The following report expands on some of the points made in my Chair's talk, as well as giving a flavor of discussion from the floor and relevant points from other sessions of the conference. This report draws on further reflection and discussion, as well as additional material and developments that have followed since the original meeting.

Report

Robotic and other cognitive technologies are playing increasingly prominent roles in medical treatment and in personal and social care, assisting in care for people in hospitals, care homes and private houses and other domestic contexts. Often it is tempting to anthropomorphize such technologies, and to see them as artificial agents entering into social relationships, and working, with various degrees of operational autonomy, in close proximity to many individuals, often those that make up the most vulnerable members of the population. Artificial agents in these applications will be fulfilling functions that may routinely require interactive and affective sensitivity, practical knowledge of a range of rules of professional conduct, and general ethical insight, autonomy and responsibility. They will be working with people who are in fragile states of health, who have physical or cognitive disabilities of various kinds, who are very young or very old, and so on.

The areas of health and social care cover a wide range of applications from clinical and hospital settings to residential and home-based care for the old and for people with cognitive and physical disabilities. At this EUCogIII meeting the area of smart medical agents was less widely addressed than that of personal care agents. This may be because medical AI and robotics includes a more diverse and technical set of issues, issues which are less easy to grasp or generalize about by an audience who are not specialists in the medical application domains. Also, there is perhaps more scope for the deployment of specialized AI- and/or robotics-based technologies in a sector which is already highly adapted to the employment of sophisticated technical solutions for many focused clinical applications. (As examples, papers discussing medical applications included in an April 2014 symposium on medical and social care at the AISB50 Convention in London feature the following medical systems: the Da Vinci and the Cyberknife robotic surgery systems, and a largescale clinical decision support system being developed by Oxford, UCL and the Royal Free Hospital, London. See [1, 2, 3] respectively.) For whatever reason, the discussion at the EUCogIII meeting in October 2013 mostly covered the role of cognitive technologies in social care rather than in medical care, and that will be the main focus of discussion here.

The scope for possible application of AI agents and related technologies in the social or personal care sphere includes, but is far from exhausted by, care for the elderly. This is an area where large numbers of frail people require daily assistance from human professional carers, and where there are crucial demographic factors which will see the demands for such care increase steadily over coming decades. As Robert and Linda Sparrow pointed out, in

their landmark 2006 article [4], the care of the elderly has been one of the most popular application areas used to market robotic and related AI technologies. This is despite the fact that such elderly and frail people, often living on their own, can live very precarious lives, where the use of technologies to supplement or replace human carers will thus raise deep issues of ethics and social responsibility, but also great possibilities for improving their lives and the quality of human-delivered care. The temptation is great for tech developers, care providers and governments to seek ways to massively deploy assistive technologies to try to address the needs of this population – and the needs of related groups, such as those with cognitive and physical disabilities – people who require care workers in large numbers for physical support and companionship. This application sector is nonetheless fraught with challenges, in terms of technical as well as ethical viability.

Alongside the growing demand for carers there is a shrinkage of available human labour because of the trends towards an increasingly greying population in Europe and elsewhere in the world. These factors are exacerbated by a shrinkage in available funding to support such care. It was ironic that, on the day of this EUCogIII meeting, a front-page news report carried by the London *Guardian* bore the following headline: "Council funding cuts force care firms to pay less than the minimum wage" [5]. The article – one of a series of *Guardian* reports that had come out on this issue, detailing the cash crisis in the system for financing public care in the UK – provided a useful backdrop to the afternoon's discussion. Local authorities are required by the Government to commission care in their localities for elderly and vulnerable people at home and in residential care. In the UK the finances these local authorities receive from Central Government have been reducing, in real terms, for some considerable time. So the care contracts that Councils are able to offer to care provider organizations pressurize such organizations to pay care workers at the UK national minimum wage (currently GBP 6.31 per hour, but set to rise in Spring 2014). As the *Guardian* article made clear, care provider companies can sometimes even go below the minimum wage via workarounds such as not paying care workers for necessary travel times between visits, etc.

Another consequence of this cash squeeze is that care professionals are allotted less and less time to spend in the presence of a client. The Conference had been told, in the opening keynote talk that morning by Prof. Tony Prescott (University of Sheffield), that many elderly recipients of care support now receive just 15 minutes of the care worker's time per visit [6]. This had in fact been the subject of a *Guardian* article on the previous day, and has been covered in many other media reports in the UK. See, for example, [7, 8]. The chances are that things are going in the same direction in many parts of the EU.

This prompts a number of conflicting reflections, on the part of developers of AI and robotics solutions for social and health care. (1) Given the lack of available money for human-delivered care support, the situation is ripe for introducing robotic and other AI-based solutions, provided these could be deployed in ways that economized on human staff costs, and provided, too, that their use respected ethical standards (a big if, as was emphasized in much of the discussion at the meeting). (2) Future funding to support proper R&D on such solutions (with adequate ethical safeguards) is likely to be tight, because of overall economic uncertainties. (3) Commercial companies may well exploit the situation by putting pressure on care providers to accept solutions that are inadequate in either technical or ethical respects, incompletely tested and with poor managerial frameworks. However: . . . (4) There is often an unspoken assumption that existing human delivery of such care functions provides a paragon of acceptability which it would be hard for machines to match up to. So discussions of the drawback of automated care often centres on threats to human dignity, independence, safety that would be engendered by automating delivery of care. In fact the economic strains in the existing care system mean that, notwithstanding the high degrees of dedication and commitment among professional carers, the profession suffers, as we've observed, from chronically poor pay, but also inadequate training, high levels of stress, and high staff turnover. All of these factors mean that the current experience of care delivery by those who

use such services is often far from ideal – so the baseline against which automated care technologies are to be introduced is perhaps lower than what is generally assumed.

As was stressed in a discussion from the floor (and also in Madeline Drake's panel presentation on the following day), the introduction of even relatively primitive forms of assistive technology may have the effect of transforming the lives of service users for the better. For example, introducing customized "smart home" facilities to enable people with cognitive impairments to live independently in their own flats may enable them to attain a quality of life that could not be provided otherwise, given the kinds of human staffing levels that funds would currently permit. Designers of more complex cognitive and assistive technology systems may have something to learn from observing how relatively established smart systems are employed in the practices of care-provider organizations, and the ways in which they do (or don't) enhance the lives of care service users. Smart surveillance technologies, often viewed with ethical suspicion, can be welcomed by certain care providers for some families where child sexual abuse is suspected, where such technologies, installed to monitor unusual behaviour in the home, can help beneficially to keep families together where otherwise they may be broken up - e.g. if the standard practice for that kind of case is to remove the child(ren) at risk into foster care [9].

Clearly different technologies may confer different benefits and risks in this area. EU-funded support for further R&D in assistive care technologies would clearly be desirable, if only to see what developments could be brought into widespread deployment across different member countries. It may be that the sophisticated robotic companions that are often foregrounded in this kind of research, would be too complex and expensive to realize in a way that could achieve genuine net improvements across the European care sector (for example, costs of commissioning, setting up managerial frameworks, technical backup, staff training, fail-safe systems, etc., may outweigh savings in staff costs.) However, less complex technologies, which still incorporate many AI and robotic system features, may have the potential to be very successful and widely used.

Such examples may correct a tendency to assume that the only "cognitive technologies" relevant in this area must be rich, deeply anthropomorphic (or zoomorphic) systems with comprehensive cognitive architectures incorporating features which include sophisticated learning abilities, flexible deployment of learned knowledge and a set of autonomous, system-owned, goals. (See [10], for example.) Such systems are of course central to the research challenges at the apex of the cognitive systems research community. They will no doubt also be important in certain care and medical contexts. Nevertheless simpler, less ambitious systems, that still involve AI and/or robotic technology, may still be highly relevant to the markets in these application areas, and easier to deploy in the kinds of numbers that would make them economically viable as more than experimental tools.

Artificial intelligence and robotic technologies can both be seen in contrasting ways – as leading-edge and exotic at one extreme and as incorporating the relatively primitive and mundane on the other. It has been pointed out many times - for example, by Larry Tesler, as quoted by Doug Hofstadter [11] - that AI can be viewed as "whatever hasn't been done yet" [1]: that is, as soon as research goals have been achieved and become mainstream they cease to be called AI. But there is another way of looking at AI – that now, in the 2010s, it's all around us – in how we wordprocess, how we google, how our phones, cars and (increasingly) our homes work. Similarly, while robots may be seen in humanoid or personoid terms (i.e. the traditional stars of scifi literature and film), it is also possible, as Jordan Pollack urges, to see robots as *any* physical devices that operate by means of information-processing and programmed control systems, including printers and ATMs [12].

[1] Hofstadter called this Tesler's Theorem. See [11: 601].

Those whose conception of AI or robotics tends to gravitate towards the more exotic (Tesler) end of this spectrum, will (validly) see the major challenges of AI, Robotics, and Cognitive Engineering in terms of these sophisticated systems. However, the key issues concerning embedding of cognitive technologies in wide-distribution applications require one (equally validly) to shift one's viewpoint towards the other (Pollack) end.[2] The implications of this duality of vision for considering robots in medical and care contexts are clear. At one end of this continuum we see agents that may replace nurses, doctors, care staff or informal companions, and which thus raise heady ethical and social problems for the future, but which may be seen in mass use only in the very remote future. At the other end we see useful tools that already are incorporated in many settings of medical practice, and, perhaps increasingly, personal care practice. (Is a smart home a robot? An ultrasound machine? In Pollack's terms, yes to both, but surely neither are intelligent agents capable of ethical decision-making, in the sense of an Asimov story or the philosophical visions of the Machine Ethics community.)

These points are worth bearing in mind, when thinking of the recent Eurobarometer Survey on public attitudes towards robots [13] (discussed by Mark Coeckelbergh in his introductory talk to the session, and several times in the floor discussion, and also previously in Tony Prescott's earlier keynote talk). In one way the Eurobarometer findings reproduce the duality of vision just discussed. EU citizens see robots both as "instrument-like machines" and as "human-like machines" (with a preponderance to the former). As against this, however, only 12% of citizens claim to have had experience of using robots – whereas, according to the Pollack conception, anyone who has used an ATM or a modern car with smart electronics should count as having used one.[3]

What implications does this have for robots in medical or care contexts? Significantly for our present discussion (and as was emphasized by several discussants at the EUCogIII meeting), 60% of respondents to the Eurobarometer Survey were in favour of *banning* the use of robots in the care of children, the elderly or the disabled. One wonders, however, if there might have been different attitudes expressed if a wider range of possible candidates for the label "robot" had been given, including such mundane devices as printers and ATMs. At the very least it points to a significant, middle, position between the extremes of our AI/Robotics continua, namely, a view of Robotics as a currently breaking wave, with both positive and negative possibilities. This middle area perhaps provides the most likely source for future robotics and cognitive systems initiatives in the medical and care domains. Perhaps it is one which will be more positively received by the European public when a more mature view of the nature robotics is offered to the public than was offered in this Eurobarometer Survey.

Summing up.

We have concentrated in this discussion summary on the potential and risks for using cognitive technologies in the personal care sector, as opposed to the medical sector. We have stressed that the use of robots and other cognitive technology in care settings has to be viewed in the context of the crisis in human-delivered care within the UK (and probably in other European countries; also similar things can be said about crises and strains in the UK's and other member states' health services.) For reasons given earlier, this makes the social and ethical issues arising from the opportunities for increased deployment of cognitive technologies in health and personal care a lot more complex and sensitive. A broad observation made earlier was that there is no point in using human-delivered care as an ideal to which AI-delivered or AI-enhanced care has to match up if, for vast members of the population requiring care-services, human-delivered care is provided in far thinner and less

[2] For convenience I'm from here collapsing what were previously mentioned as two continua into one.

[3] Respondents' conceptions of what counts as a robot were shaped, in the survey, by their being shown two photos, one of a factory robot-arm (of a sort that is probably in widespread use) and one of a wheeled, perambulating, vaguely humanoid servant robot (of a sort that is not yet found in great numbers, and which may or not ever become a mainstream device – and why this clunky, lumbering device, rather than a Nao Robot or something similar?). See Question QA2 [13: 12].

frequent slices than is required, inadequately resourced, subject to insufficient training provision, and experienced as unsatisfactory by the consumers of care in a variety of ways.

In saying this one must emphasize that professional care staff are, for the most part, highly dedicated, sensitive and *caring* individuals, who, driven by their commitment to strong ethical values, are often working for less money than they would get in other, less demanding professions, and, with their managers, are under enormous pressure to provide the best service for their clients that the shrinking resources and growing numbers in need will allow. Again, a good deal of this – both the negative and the positive points – apply as much in the medical as in the care sector. Within this environment there remain many opportunities for development of robotic and cognitive solutions as enhancement to human-delivered care and support. However cognitive systems developers must always remind themselves that any solutions offered are to be deployed on those who are, whether through physical or mental illness, or through old age or other adverse life-circumstances, among the most vulnerable in our society. The EUCog network and similar initiatives thus need to ensure that AI technologies can be used to improve on, rather than to add further to, the stresses in social health and care provision that increasingly seem to be experienced by those who have to receive them.

Concluding note.
It is worth mentioning a couple of related events and a forthcoming publication.

(a) There is a symposium at the AISB 50th Anniversary Convention in London in April 2014 on 'Machine Ethics in the Context of Medical and Care Agents' (MEMCA-14). 15 papers will be delivered. The symposium, chaired by Steve Torrance, Mark Coeckelbergh and Aimee Van Wynsberghe, benefits from support from the EUCogIII network. http://aisb50.org/machine-ethics-in-the-context-of-medical-and-care-agents/

(b) A workshop, chaired by Matthias Scheutz, on "Autonomous Social Robots: The Challenges Ahead" is being held in May at the ICRA 2014 meeting in Hong Kong. http://hrilab.tufts.edu/robotethics14.html

(c) A collection of papers on *Machine Medical Ethics* is being edited by Simon Van Rysewyk and Matthijs Pontier, to be published by Springer later in 2014.

Acknowledgements
I would like to thank Mark Coeckelbergh, Madeline Drake and Sharon Wood for useful inputs into this report.

April 2014.

References.
[1] K. Friedrich & M. Queisner "Automated killing and mediated caring: How image-guided robotic intervention redefines radio surgical practice". *Proceedings, MEMCA-14: Machine Ethics in the Context of Medical and Care Agents, AISB50 Convention,* Goldsmiths, University of London, 2014 (forthcoming).
[2] J. P. Sullins "Ethical trust in the context of robot assisted surgery" *ibid.*
[3] J. Fox & O. Khan "From decision support systems to autonomous agents: how can we ensure ethical practice?" *ibid.*
[4] R. Sparrow and L. Sparrow, "In the hands of machines? The future of aged care" *Minds & Machines* 16: 141-161, 2006

[5] *The Guardian*, "Council funding cuts force care firms to pay less than the minimum wage", Wed 23 October 2013. http://www.theguardian.com/society/2013/oct/22/council-funding-cuts-care-homes-minimum-wage
[6] Tony Prescott: "Socially-responsible research in Cognitive Systems: No-win or win-win?" EUCogIII Members' Conference. Video, available at: http://www.youtube.com/watch?v=lwBbI-XgEuU
[7] Leonard Cheshire Disability, *Ending 15-minute Care*, http://www.leonardcheshire.org/who-we-are/publications/latest-publications-download#.UxOxNaX7puY
[8] BBC Radio 4 *File on 4: Cut-price care* Broadcast in February 2014: http://www.bbc.co.uk/programmes/b03szh9m
[9] M. Drake, "Cognitive and assistive technologies in care practice: a view from the delivery end." *Proceedings, MEMCA-14: Machine Ethics in the Context of Medical and Care Agents, AISB50 Convention*, Goldsmiths, University of London, 2014 (forthcoming).
[10] A.Gomila, V.Mueller, "Challenges for Artificial Cognitive Systems", *Journal of Cognitive Science* 13: 453-469, 2012
[11] D.R.Hofstadter, *Gödel. Escher, Bach: An Eternal Golden Braid,* New York: Vintage Books, 1979.
[12] J.B.Pollack, "Seven Questions for the Age of Robots", Yale Bioethics Seminar, Jan 2004. http://jordanpollack.com/sevenlaws.pdf
[13] European Commission DG INFSO, Special Eurobarometer 382, *Public Attitudes Towards Robots.* (2012) .http://ec.europa.eu/public_opinion/archives/ebs/ebs_382_en.pdf