



TU Delft

DESIGN FOR VALUES

Saving the Life of Medical Ethics in the Age of AI and Big Data

Jeroen van den Hoven¹, 23 March 2018

Talk given at the WHO Global Summit 'Bioethics, Sustainable Development and Societies (Dakar, Senegal)

Digital technology² impacts health care in all of its dimensions: research & development, clinical practice, policy, innovation, entrepreneurship, insurance and financing. It leaves nothing as it is. It is not a mere enabler, it is a *constitutive* technology. It radically changes the very practices, sectors and institutions to which is applied. We don't need to spend much time on explaining how exactly digital technology manages to disrupt. The picture is clear by now. Incumbents in finance and banking, transport, hospitality, industrial production and retail have all experienced it. You go digital or you disappear. It is also obvious that data and AI can reduce cost in health care, improve patient safety, empower patients, and improve quality of diagnosis, therapy, patient journeys, efficiency in billing and logistics. Smart phones and watches with health Apps, wearables and so-called digi-ceuticals are part of an Internet of Things revolution that is well underway in the health sector. Wearable devices can be used to detect arrhythmia, predict Parkinson (via the accelerometer in the phone), and measure a range of biomarkers such as blood sugar, blood pressure, fat percentage, oxygen, and stress. They can diagnose skin cancer, retina damage, and assist in the management of eating disorders, phobias, depression, chronic pain, and PTSD. They can even gauge the risk of suicide on the basis of social media posts – looking at the time of day of the post, the number of human faces it contains and the colours.

Big Tech companies are moving into Health Care and Biomedicine big time. There is a lot of money to be made and big data to be harvested. Big data in turn will drive the development of more powerful Machine Learning and AI, which leads to superior positions in the market. IBM applies its Watson technology in oncology and uses health data collected via Apple ResearchKIT and HealthKit. Alphabet focuses via its London based Deep Mind Health on AI for health. They have managed to close deals with NHS and individual hospitals to get access to large amounts of patient data in the UK and are trying to achieve breakthroughs comparable to beating human world champions at Chess and Go. Alphabet already claims to be able to predict the death of patients in hospital much earlier than clinicians using traditional methods. Alphabet's Cityblock Health focuses on low income inhabitants of cities, mining data to see where care for this vulnerable group is needed and dispatch health care workers to people's home, circumventing existing health care infrastructures. Alphabet's Verily has embarked upon a base line study for good health by following 10 K patients for 4 years closely also looking at insurance schemes and 'population health management'. Microsoft just started a new healthcare division in Cambridge looking at medical algorithms. Apple has chosen the hardware route and focuses on smart phones and wearable sensors as medical devices in collaboration with Stanford. Amazon has teamed up with Warren Buffet's investment company Berkshire Hathaway and JP Morgan Chase to move into health care. Uber has laid out plans to use its ride hailing platform for transportation of patients to and from the hospital.

¹ Delft University of Technology; m.j.vandenhoven@tudelft.nl; University Professor Ethics and Technology Scientific Director Delft Design for Values Institute

² *The Economist* of 13-9 February 2018 makes the data revolution in health care into the cover story "Doctor You. How Data will transform health care". In the following I draw upon the recent developments reported in the background article, pp. 51-53.

The first forays into Health care in the past have not all been good. Google Health failed in 2011 and Google Flu Trends claimed it could outperform the CDC in predicting the Flu, but completely missed the peak of the influenza epidemic in 2013. Spurious correlations and overfitting were the culprits³. The Ebola crisis response on the basis of Orange Telco calling data was less than adequate⁴ for similar reasons. There are serious limitations to a purely data driven approach and the glorification of statistical correlation. On the basis of a recent critique on Machine Learning by the leading AI researcher Judea Pearl from Berkeley one could say that - especially in medicine and in other fields of great social importance – the data driven approach in order to be clinically useful and morally responsible needs to be complemented by theory driven approaches which aim at uncovering causal mechanisms.

Another set of problems are attitude problems, so to speak. The Digital Industry and the Silicon Valley approach to health care is what Evgeny Morozov has called a *solutionist approach*, which focuses exclusively on problems for which we have nice and clean technological solutions at our disposal. David Lazer has diagnosed their preferred approach as *Big Data Hubris*. The idea that there are simple digital solutions to complex problems in a very complex world of health care, with its very complex institutional settings, multiple stakeholders, plurality of moral values, and cultural diversity. This is irredeemably and culpably naïve.

There are not only *epistemic* failings in the digital usurpation of the health domain. There are also *moral* concerns⁵. We know that there are racial and gender -and many other - biases in health care data and they may get entrenched in algorithms or may be built in to medical systems to deceive, to save cost or make profits. We have seen with Volkswagen and Uber how this can work. How would we know? Algorithms and block boxed Intelligence in clinical decision support systems affects the fiduciary relationships of doctors and patients. Furthermore, there have been massive breaches of security and privacy in health care in last decade⁶. Google Deep Mind Health has been reprimanded by the Information Commission in the UK for their processing of NHS patient data. Google Deep Mind replied that they had "...underestimated the complexity of the NHS and of the rules around patient data, as well as the potential fears about a well-known tech company working in health." There are also worrying developments beyond privacy and security breaches and touch upon human dignity. A US based company called *Aspire Health* tries to save on cost in palliative care by estimating which patients will die soon. A discussion in the New England Journal of Medicine of March 2018 states that "there may be a temptation to teach ML systems to guide users towards clinical actions that would improve quality metrics but not necessarily reflect better care."⁷

Return on Trust

But the key question of course is: Why would we trust Facebook, Uber, Google, Amazon and Microsoft with all of our sensitive medical data? They can't even fix basic problems with fake news, security, filter bubbles, bias, nor can they prevent data of 50 million users being abused to run political campaigns.⁸ Big Tech is essentially about Quartely Revenues. They come to health Care with a Silicon Valley approach to innovation: Innovate in the gray zone, move fast, break things first and apologize later. This is not a very helpful approach in health care.

³ See David Lazer's account in Wired Magazine, <https://www.wired.com/2015/10/can-learn-epic-failure-google-flu-trends/>; see also <https://gking.harvard.edu/files/gking/files/0314policyforumff.pdf>

⁴ <http://www.bbc.com/news/business-29617831>; and <https://anthrosource.onlinelibrary.wiley.com/doi/pdf/10.1111/maq.12440>

⁵ See for an overview e.g. "Implementing Machine Learning in health care – Addressing Ethical Challenges", Chan, Nigam, e.a. *N.Engl. J. Med.* 378:11, March 15, 2018.

⁶ <https://digitalguardian.com/blog/top-10-biggest-healthcare-data-breaches-all-time>

⁷ Chan, e.a., op cit., p. 982. See a detailed report on the consequences of the use of algorithms in health care for individual patients: <https://www.theverge.com/2018/3/21/17144260/healthcare-medicaid-algorithm-arkansas-cerebral-palsy>

⁸ Wired Magazine, <https://www.wired.com/story/embattled-tech-companies-charge-deeper-into-health-care/> and for recent problems: <https://www.theguardian.com/technology/2018/mar/22/they-were-given-an-inch-and-took-100-miles-readers-on-cambridge-analytica-facebook-and-privacy>

This poses one of the most important and hardest problems of the twenty first century in my opinion: Trust or the lack of Trust. We could all benefit if we could only trust others with our data. If we cannot trust or misplace our trust, there will be enormous cost.

Let's look at a simple example to understand what trust is. Suppose your plumber arrives at your house to fix your kitchen sink. Do you trust the plumber? If you would have to place a bet on him fixing your sink in an hour or so, would you be prepared to stake a 100 dollar? If you would have great confidence in him as a professional, his skills and reliable performance, you could take the risk. But now ask yourself: do you trust your plumber with the silver that is in your kitchen drawer? That is seemingly a totally different matter, which is a question that has nothing to do with his abilities, skills, knowledge and tools. In order to decide that question (and the fate of 10.000 euro worth of silver) you would need evidence of a different sort. Nothing related to his skills and expertise. You need evidence of the fact that he is an honest or a morally good person that he has no intention to deceive you whatsoever. A distinction must thus be made between confidence in the reliable performance of some sort, and trust in moral motivation of others.

Now I repeat: can we trust Big Tech and their acolytes and subsidiaries with our health data? Can we ever be sure that their services will not be solicited by foreign failing states, guarantee that they (and our data) will not merge with companies and data bases in the hands of oligarchs who do not feel constrained by the rule of law? It is against this background we need to situate the discussion about sharing and using identity relevant data.

Sweden provides an interesting lesson. It is moving ahead in sharing and pooling of health care data. People have no problem, but Sweden is not a company, it is a very High Trust society, known for its openness, transparency, democracy for a long time. Now it is getting what could be called its "Return on Trust". Swedish citizens trust their government with their medical data. In order to be able to emulate this result we need to understand that Trust is a moral phenomenon and it is tied to the quality of institutions in the context of which ethics is taken serious and is widely known to be taken serious by all..

Ethics and Health care Data

Now how should we do Medical ethics in the age of big data and AI? In a statement on the ruling in the Google Deep Mind NHS case⁹ information and privacy commissioner Elizabeth Denham said: "The price of innovation does not need to be the erosion of fundamental privacy rights"¹⁰.

Denham is right. And there are two new ideas here. The first idea is that Ethics in the digital world should be about designing things, systems, devices, algorithms, governance systems, protocols and the combination thereof. If we do not consciously and carefully design for our shared moral values in the age of high technology, then our conceptions of privacy, accountability, democracy, autonomy, safety and security will simply not materialize. Ethics will not be inserted where the rubber hits the road in our day and age. Our talk of privacy and accountability in health care will in a sense be gratuitous if we cannot provide the specifications for - let us say - 'the role based access matrix of a hospital information system'. If we do not design for our values others will fill in the void and design the digital world for their ideas, and preferences, instead - and often in ways that are not clear to us or unacceptable. It is true what Churchill said: first we shape our houses and then our houses shape us. We become in a sense like the things we design. If we want ethics in the 21st century we must *build* it. This is what I have referred to as the design turn in applied ethics¹¹.

This brings us to our second idea, that of responsible innovation. For if we design for our values in a systematic, transparent and accountable way, we also may be rewarded with solutions to value conflicts and to what

⁹ <https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2017/07/royal-free-google-deepmind-trial-failed-to-comply-with-data-protection-law/>

¹⁰ <http://www.bbc.com/news/technology-40483202>

¹¹ Van den Hoven, e.a. (eds.) *Designing in Ethics*, Cambridge University Press, 2018.

seemed in first instance insurmountable dilemmas. Again a garden variety example: you want to have a walk outside and you want to stay dry. But unfortunately it is raining. So you face a tragic choice: you don't have your walk and stay inside and be dry, or you go out and have your walk outside, but you get wet. The innovative idea of the umbrella would allow you to overcome the conflict and solve your dilemma. The core idea of responsible innovation draws our attention to the fact that we can try to accommodate as many of our moral values as we can - at the same time - by design, by tweaking the world, by innovation, by creativity. We could try to be really smart and have both our privacy and security, health and efficiency, economic prosperity and sustainability. Our research shows that remarkable innovations often achieve this moral sweet spot and that is exactly what strikes us as smart about them. We felt morally overloaded by the demand of all of our values, the dilemma's, but a smart innovation shows us that it does not always need to be a matter of "Either Or", but it could possibly be "And And". We are looking for new functionality and smart solutions that allows us to have our cake and eat it and help us to prevent having to make tragic choices. A responsible innovation changes the world in such a way that it allows us to do more of the things we are morally obligated to do. That is moral progress¹².

This applies to medical ethics in a world of Big Data and AI. We need to accommodate privacy concerns AND make use of Big Data and AI in health care. There is no guarantee that this will always be possible. But because the stakes are high we have the obligation to explore at least whether there are such solutions.

Medical Ethics, Design for Values and Responsible Innovation

In 1982 the philosopher of Science Stephen Toulmin published a paper with the title "How medicine saved the life of ethics."¹³ Toulmin argued that new health care technologies and new developments in medicine had caused a Renaissance in ethics at the moment the discipline was about to make itself irrelevant and obsolete. A Dutch colleague recently revisited Toulmin's paper in her inaugural address as professor Ethics of Biomedical Innovation at the University of Utrecht under the title "Can Ethics Save the Life of Biomedicine?" She argued that now ethics should come to the rescue of biomedicine, since it is beset with so many hard moral problems¹⁴.

I have suggested, that medical ethics needs to shift gears in the age of Big Data and AI if it wants to save the lives and dignity human beings and stay relevant. It needs to engage the digital world in its complexity and entanglement with industry. Without compromising on ethics, human well-being, human dignity and public interest under the pressure of profit maximization. It needs to seek responsible innovations and needs to design for moral values. Medical ethics is for a good part digital construction work in the remainder of the 21st century.

If we do ethics along these lines effectively we may be able to solve the trust problem which is - as we saw - about genuine and obvious commitment to ethics in a digital age.

Chan, e.a. have nicely captured that idea in New England Journal of Medicine in their discussion on AI and Big data in Medicine of march of this year when they argued that Machine Learning systems should "...be built to reflect the ethical standards that have guided other actors in health care – and should be held to those standards" (ibidem, p. 983). Medical ethics will have to go digital, or become irrelevant.

¹² See for this view Van den Hoven "Value Sensitive Design and Responsible Innovation" in Owen e.a. (eds.) *Responsible Innovation*, Blackwell, 2013 (<https://doi.org/10.1002/9781118551424.ch4>)

¹³ <https://www.ncbi.nlm.nih.gov/pubmed/11643854>

¹⁴ <https://mdrresearch.nl/2018/02/26/prof-dr-annelien-bredenoord-will-deliver-her-inaugural-speech-march-2nd>