# The Perils of Artificial Intelligence and Big Data

*Max Welling, University of Amsterdam, January 31, 2015*

As technology becomes more powerful, both its positive and negative aspects are enhanced. When the axe was invented we could cut trees more efficiently but also kill people more efficiently. When nuclear fission was discovered we could build nuclear power plants but also nuclear bombs. These days there is a certain perception among eminent people such as Elon Musk, Bill Gates and Stephen Hawking that AI is at the point of improving so fast that soon it will become a danger to mankind. Ray Kurzweil's singularity is a point in time that technology and artificial intelligence propels itself so fast that it literally explodes. *Is this singularity something we need to worry about? Even without a singularity will AI pose a serious threat to humanity? And if so when? What are the ingredients for an uncontrolled acceleration of AI?*

A significant number of esteemed researchers recently signed an recent open letter drafted by "The Future of Life Institute" stating that we should direct our resources to better our society. *What are implications of this letter? Do we as AI researchers need to change our research agendas? By signing the letter, do we implicitly agree to the fact that AI is a potential and immediate danger to society?* One place where advanced technology and therefore AI can do serious harm is in the arms industry. *Should we avoid taking grant money from DARPA or other military agencies for our research?*

Besides the rise of the machines, more realistic short term threats to our society should also be contemplated. The explosion in data gathering about our personal lives by both government and industry poses a real threat to privacy. Public datasets, which are claimed to be anonymized are in effect not privacy preserving at all. A famous example was the identification of the governor of Massachusetts' personal health records around 2000 by Latanya Sweeny who combined public (anonymized) health records with public voting records. On the other hand imposing very strict privacy regulations might hurt our ability to combine, say, health records to achieve the dream of personalized medicine. Or it might obstruct our ability to find and track terrorists. *Can we simultaneously achieve privacy and reap the benefits of big data? If not, what is the right balance?* Many (often young) people don't seem to care much about privacy ("they can know whatever they want about me because I haven't done anything wrong"). *Is this point of view naïve? What (besides an uncomfortable feeling) are the real dangers of a lack of privacy?*

Besides privacy, our ability to predict people's needs and behavior may change our society in ways that are hard to image. Facebook wants to develop a "theory of mind" for every Facebook user. *Would we be comfortable with Facebook understanding our drives and needs better than ourselves?* An insurance company that offers lower fees to individuals who are predicted to have low risk effectively raises the fees for everyone else. While it seems reasonable for a car insurance to factor driving behavior into insurance fees, it seems a lot less reasonable for a health insurance to use genetic information for this purpose. *What information are insurance companies allowed to use for determining personalized insurance fees? What other societal changes can we expect when algorithms can predict our behavior and needs with high precision?*

Further reading on the dangers of AI:

http://www.livescience.com/49625-robots-will-not-conquer-humanity.html

http://futureoflife.org/misc/open_letter

http://www.popsci.com/bill-gates-fears-ai-ai-researchers-know-better

http://www.bbc.com/news/technology-31023741

http://www.cnbc.com/id/101774267#.

Further reading on the dangers of big data:

http://arstechnica.com/tech-policy/2009/09/your-secrets-live-online-in-databases-of-ruin/

http://bits.blogs.nytimes.com/2015/01/29/with-a-few-bits-of-data-researchers-identify-anonymous-people/?smid=tw-nytimesbits&_r=1