Statement of:

Subbarao Kambhampati
Professor, CIDSE, Arizona State University
President, Association for the Advancement of Artificial Intelligence

Before the
Standing Committee on Social Affairs, Science & Technology
Senate of Canada

Study on the role of robotics, 3D printing and artificial intelligence
in the healthcare system
https://sencanada.ca/en/Committees/SOCI/NoticeOfMeeting/449379/42-1

March 30, 2017

Thank you all for the invitation. It is quite an honor to get a chance to speak with you today.

My name is Subbarao Kambhampati, and I am a professor of computer science at Arizona State University. I am also the current president of the Association for the Advancement of Artificial Intelligence, AAAI for short, the premier scientific organization devoted to Artificial Intelligence. AAAI has a significant membership and participation from Canada; Prof. Alan Mackworth from University of British Columbia was a former president. Several of our annual meetings have been held in Canada. I would also like to mention that I am an inaugural trustee of the Partnership for Artificial Intelligence, a consortium that was formed last year to focus on the responsible development of AI technologies. The partnership includes both industries, Google, Apple, Amazon, IBM, Facebook and Microsoft; and non-profit organizations including AAAI and ACLU (the American Civil Liberties Union). I thought this may be of interest to you if you have questions about initiatives on the societal impacts of AI.

I have been involved in AI research for over 30 years now. My current interests are on planning and decision-making in human-in-the-loop and human-aware AI systems. I will say a little more about that later. While I have a good understanding of the current state of AI and its applications, as I mentioned to your staff, I am not an expert on the applications of AI in health-care specifically. I did take the time to read the transcripts of your committee's meetings since February, and I thus have some understanding of what you have already discussed.

As you have already heard, Artificial Intelligence, as a discipline, aims to get computers to show behavior that we would consider intelligent. Intelligence, of course, is multi-faceted. In my discussions about progress and impact of AI with people outside our field , I often find it useful to contrast the progress of AI to the way human babies acquire their intelligence. Informally, babies start by showing signs of perceptual

intelligence--how to see, hear and touch the world around them; and physical manipulative intelligence--how to walk, roll, manipulate objects etc. They then acquire emotional intelligence and social intelligence (how to model other agents), and only then get into the realms of cognitive intelligence--doing well in standardized tests, playing games like Chess etc.

It is interesting to note that the progress of AI happened almost in the reverse direction. We had systems showing high levels of cognitive intelligence quite early on. There was the boom of medical expert systems in the 80's, and in the 90's we had Deep Blue defeat Kasparov. However, it was only recently that AI systems have reached impressive levels in perceptual intelligence--how to see and hear the world around them. Surprising as it might seem at first glance, this reverse direction progress is understandable.To a first approximation, AI researchers started first by teaching computers things they consciously know how to do. Perception and Manipulation are things that we do unconsciously, so the only way to get the computers to do them well was to make them learn. For this, we had to await the availability of large-scale data. The so-called "deep-learning" systems existed since 1980s but blossomed only after the internet made data--be it images, speech signals, or text--easily accessible.

This perspective also helps us put in context the recent flurry of interest and commercial applications of AI--advances in perceptual intelligence made it possible for AI to reach a much wider audience; it becomes easy to experience the fruits of AI technology when your cell phone recognizes your voice and the images around you. The next wave of developments for AI are expected to come from harnessing the strides made in cognitive intelligence to connect perception, reasoning/planning and action.

Turning to the health-care applications of AI, by the late eighties, AI systems were being used for clinical assistance. They were in fact the big drivers of expert systems boom. However,they had to be hand-fed inputs. Just as deep blue couldn't recognize a chess piece, medical expert

systems could not “see” the patients they were diagnosing (there is a funny anecdote about Mycin, the first expert system, merrily diagnosing the hand-fed symptoms of a faulty car engine, for internal medicine problems). The new wave of AI applications are able to combine perception with reasoning/diagnosis--for example reading X-rays, cardiograms, and photographs. These new technologies can help us tackle some of the most intractable problems related to health care, including, for example, human errors in hospitals, that, according to some studies are the third leading cause of deaths in the USA!

One long standing tension in our field has been between Artificial Intelligence vs. Intelligence Augmentation--AI vs. IA. Most AI research outside health-care applications focused more on getting systems to operate autonomously and on their own, rather than with humans in the loop. When you have humans in the loop, the AI systems need aspects of emotional and social intelligence. In particular, they need to be able to model the mental states and intentions of their human team-mates, behave in explicable ways, show appropriate emotional responses, and provide adequate explanations of their recommendations, in order to earn the trust of the humans in the loop. This sort of “Human-aware AI” systems are crucial for applications of AI in healthcare--especially when you start looking at systems that interact directly with patients whether to encourage healthy behaviors, or to provide direct home care for elderly or injured.

The advances in AI research, and the widening use of AI technologies have also brought to fore concerns about responsible and ethical use of these technologies. While voluminous amounts of health-related data is available, and can be leveraged to provide better health-care, it is critical to put in place best practices that preserve the privacy and confidentiality of the patients. This requires both technical developments (such as the use of blockchain technology), and policy decisions. Another issue is “data

bias”--in an inclusive multi-racial and multicultural democracy like yours, it is crucial that the predictive models guiding the health-care decisions are learned from data that is truly representative of the population. There are initiatives and organizations, including AAAI and Partnership for AI, which are looking into these challenges, but much remains to be done.

I want to end, if I may, with a personal recommendation to the committee. The so-called “deep-learning” revolution owes a considerable debt to the farsightedness of the Canadian funding agencies in continuing to support research into neural networks, when much of the rest of the world had moved on to other things. Canada richly deserves to reap the benefits of its investments into this basic research, and I am heartened by the astounding entrepreneurial activity in the Montreal and Toronto corridors. As impressive as its recent successes have been, deep learning is only a part of the broader AI enterprise. The good news is that the Canadian AI research community itself is a whole lot broader and multifaceted. I do hope that your government continues to support basic research across the breadth of AI. Paraphrasing one of Canada’s famous sayings, broad-based support for basic research is the only way to ensure that you skate to where the puck is going to be, rather than where it has been.

Thank you for your attention. I am happy to answer any questions.
