

Embargoed until 6am AEST July 27

World's leading AI and Robotics researchers join forces to call for ban on offensive autonomous weapons

Open letter released today at the leading Artificial Intelligence conference, IJCAI 2015

<http://ijcai-15.org/>

Stakes are high as UN discusses future of "killer robots"

Sydney, Australia, 28 July 2015 - As the world's top experts in Artificial Intelligence (AI) today congregate for the 2015 International Joint Conference on Artificial Intelligence (IJCAI) in Buenos Aires, leading researchers in the field from all over the globe have banded together and released an open letter in an effort to urge the UN to support a ban on offensive autonomous weapons systems. The letter follows the April meeting of the Convention on Conventional Weapons (CCW) held at the United Nations (UN) Palais des Nations in Geneva.

Released today, the open letter is endorsed by over one thousand AI and Robotics researchers and nearly one thousand other experts including SpaceX and Tesla CEO Elon Musk, physicist Stephen Hawking, Apple co-founder Steve Wozniak, Skype co-founder Jaan Talinn and linguist and activist Noam Chomsky. The letter expresses a grave concern with the development of such weapons and urges the UN to consider a ban of offensive autonomous weapons beyond meaningful human control.

The letter argues that the deployment of such autonomous weapons is feasible within years, and will play a dangerous role in driving a third revolution in warfare, after gunpowder and nuclear arms.

Toby Walsh, professor of Artificial Intelligence at UNSW and NICTA said, "All technologies can be used for good and bad. Artificial intelligence is a technology that can be used to help tackle many of the pressing problems facing society today – inequality and poverty, the rising cost of health care, the impact of global warming... but it can also be used to inflict unnecessary harm.

"We need to make a decision today that will shape our future and determine whether we follow a path of good. We support the call by a number of different humanitarian organisations for an UN ban on offensive autonomous weapons, similar to the recent ban on blinding lasers."

The Open Letter as follows:

Autonomous Weapons: an Open Letter from AI & Robotics Researchers

Autonomous weapons select and engage targets without human intervention. They might include, for example, armed quadcopters that can search for and eliminate people meeting certain pre-defined criteria, but do not include cruise missiles or remotely piloted drones for which humans make all targeting decisions. Artificial Intelligence (AI) technology has reached a point where the deployment of such

systems is — practically if not legally — feasible within years, not decades, and the stakes are high: autonomous weapons have been described as the third revolution in warfare, after gunpowder and nuclear arms.

Many arguments have been made for and against autonomous weapons, for example that replacing human soldiers by machines is good by reducing casualties for the owner but bad by thereby lowering the threshold for going to battle. The key question for humanity today is whether to start a global AI arms race or to prevent it from starting. If any major military power pushes ahead with AI weapon development, a global arms race is virtually inevitable, and the endpoint of this technological trajectory is obvious: autonomous weapons will become the Kalashnikovs of tomorrow. Unlike nuclear weapons, they require no costly or hard-to-obtain raw materials, so they will become ubiquitous and cheap for all significant military powers to mass-produce. It will only be a matter of time until they appear on the black market and in the hands of terrorists, dictators wishing to better control their populace, warlords wishing to perpetrate ethnic cleansing, etc. Autonomous weapons are ideal for tasks such as assassinations, destabilizing nations, subduing populations and selectively killing a particular ethnic group. We therefore believe that a military AI arms race would not be beneficial for humanity. There are many ways in which AI can make battlefields safer for humans, especially civilians, without creating new tools for killing people.

Just as most chemists and biologists have no interest in building chemical or biological weapons, most AI researchers have no interest in building AI weapons — and do not want others to tarnish their field by doing so, potentially creating a major public backlash against AI that curtails its future societal benefits. Indeed, chemists and biologists have broadly supported international agreements that have successfully prohibited chemical and biological weapons, just as most physicists supported the treaties banning space-based nuclear weapons and blinding laser weapons.

In summary, we believe that AI has great potential to benefit humanity in many ways, and that the goal of the field should be to do so. Starting a military AI arms race is a bad idea, and should be prevented by a ban on offensive autonomous weapons beyond meaningful human control.

Professor Toby Walsh adds, “Many of the leading scientists in our field have put their names to this cause. With this Open Letter, we hope to bring awareness to a dire subject that, without a doubt, will have a vicious impact on the whole of mankind. We can get it right at this early stage, or we can stand idly by and witness the birth of a new era of warfare. Frankly, that’s not something many of us want to see.

“Our call to action is simple: ban offensive autonomous weapons, and in doing so, securing a safe future for us all.”

ENDS

Note to Editor:

Full list of signatories available upon request.

For those wanting to know more about these issues, Professor Stuart Russell from

Berkeley, one of the researchers behind drafting this letter, will be releasing the letter to the Press at a Press Conference on Tu 28th July that opens the IJCAI 2015 conference. He will also be giving a presentation on Wednesday 29th entitled, "*Killer robots, the end of humanity, and all that: What should a good AI researcher do?*". All media are welcome to attend.

The International Joint Conference on Artificial Intelligence (IJCAI) is the leading conference on Artificial Intelligence. It has been held every two years since 1969, sponsored by IJCAI Inc, a not-for-profit scientific and educational organization incorporated in the U.S. state of California. It attracts the best researchers in AI from around the world. All journalists and media are invited to attend **the daily IJCAI Press Conference**, from 8:00 to 8:30 AM at the Ombú Room, 2nd floor, Sheraton Hotel, Buenos Aires or via livestreaming on Periscope.

To obtain a free Press Pass to attend the IJCAI 2015 conference, please send an E-mail to Ignacio Uman: ignacio@umancomunica.com.ar
Follow @TobyWalsh on Twitter to be notified of the **livestreaming of the daily IJCAI Press Conference on Periscope**. Alternatively, send an E-mail to toby.walsh@nicta.com.au and he will send you the link.

For further information or to interview Professor Toby Walsh, please contact:

Elizabeth Mckenzie
Communications Specialist
Landline: +61 2 9376 2098
Mobile: +61 403 359 105
Email: Elizabeth.Mckenzie@nicta.com.au

Some of the signatures to Open Letter

AI & Robotics Researchers:

Stuart Russell Berkeley, Professor of Computer Science, director of the Center for Intelligent Systems, and co-author of the standard textbook "Artificial Intelligence: a Modern Approach"

Nils J. Nilsson, Department of Computer Science, Stanford University, Kumagai Professor of Engineering, Emeritus, past president of AAI

Barbara J. Grosz Harvard University, Higgins Professor of Natural Sciences, former president AAI, former chair of IJCAI Board of Trustees

Tom Mitchell CMU, past president of AAI, Fredkin University Professor and Head of the Machine Learning Department

Eric Horvitz, Microsoft Research, Managing director, Microsoft Research, past president of AAI, co-chair of AAI Presidential Panel on Long-term AI Futures, member of ACM, IEEE CIS

Martha E. Pollack University of Michigan, Provost, Professor of Computer Science & Professor of Information, past president of AAI, Fellow of AAAS, ACM & AAI

Henry Kautz, University of Rochester, Professor of Computer Science, past president of AAI, member of ACM

Demis Hassabis, Google DeepMind, CEO

Oren Etzioni, Allen Institute for AI, CEO, member of AAI, ACM

Geoffrey Hinton University of Toronto and Google, Emeritus Professor, AAI Fellow

Erik Sandewall, Linköping University, Sweden, Professor of Computer Science, member of AAI, ACM, Swedish Artificial Intelligence Society

Francesca Rossi Padova & Harvard, Professor of Computer Science, IJCAI President and Co-chair of AAI committee on impact of AI and Ethical Issues, member of ACM

Bart Selman Cornell, Professor of Computer Science, co-chair of the AAI presidential panel on long-term AI futures, member of ACM

Joe Halpern, Cornell, Professor, member of AAI, ACM, IEEE

Richard Sutton Univ. of Alberta, Professor of Computer Science and author of the textbook "Reinforcement Learning: An Introduction"

Toby Walsh Univ. of New South Wales & NICTA, Professor of AI and President of the AI Access Foundation

David Parkes David Parkes, Harvard University, Area Dean for Computer Science, Chair of ACM SIGecom, AAI Fellow and member of AAI presidential panel on long-term AI futures, member of ACM

Berthold K.P. Horn, MIT EECS & CSAIL, Professor EECS, member of AAI, IEEE CS

Gerhard Brewka, Leipzig University, Professor for Intelligent Systems, past president of ECCAI, member of AAI

Other Endorsers:

Stephen Hawking Director of research at the Department of Applied Mathematics and Theoretical Physics at Cambridge, 2012 Fundamental Physics Prize laureate for his work on quantum gravity

Elon Musk SpaceX, Tesla, Solar City

Steve Wozniak, Apple Inc., Co-founder, member of IEEE CS

Jaan Tallinn co-founder of Skype, CSER and FLI

Max Tegmark MIT, Professor of Physics, co-founder of FLI

Daniel C. Dennett, Tufts University, Professor, Co-Director, Center for Cognitive Studies, member of AAI

Noam Chomsky MIT, Institute Professor emeritus, inductee in IEEE Intelligent Systems Hall of Fame, Franklin medalist in Computer and Cognitive Science

George Dyson, Historian of Technology

Anthony Aguirre, UCSC, Professor of Physics, co-founder of FLI

Markus Elsner, Nature Publishing Group, Senior Editor Nature Biotechnology