



This open letter was issued on 28 July 2015 at the International Joint Conferences on Artificial Intelligence 2015.

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.

#

This open letter is available online at:

http://futureoflife.org/AI/open_letter_autonomous_weapons

According to the Future of Life Institute, a total of 20,171 people have signed the open letter: 2,981 artificial intelligence and robotics researchers and 17,190 other endorsers

Notable signatories include:

Barbara J. Grosz

Harvard University, Higgins Professor of Natural Sciences, former president AAI, former chair of IJCAI Board of Trustees

Demis Hassabis

Google DeepMind, CEO

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

Kathryn McElroy

IBM Watson, Design Lead

Elon Musk

SpaceX, Tesla, Solar City

Allison M. Okamura

Stanford University, Professor of Mechanical Engineering (courtesy appt. in Computer Science), IEEE Fellow, member of IEEE RAS, American Society of Mechanical Engineers, Society for Neuroscience, Society of Women Engineers

Martha E. Pollack

University of Michigan, Provost, Professor of Computer Science & Professor of Information, past president of AAI, Fellow of AAAS, ACM & AAA

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

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

Jaan Tallinn

Co-founder of Skype, CSER and FLI

Max Tegmark

MIT, Professor of Physics, co-founder of FLI

Carme Torras

Robotics Institute (CSIC-UPC), Barcelona, Research Professor, member of Academia Europaea, ECCAI Fellow, Senior Member of IEEE RAS and INNS

Steve Wozniak

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