



Moshe Y. Vardi

DOI:10.1145/2839512

On Lethal Autonomous Weapons

THE IMPRESSIVE PROGRESS in artificial intelligence (AI) over the past decade and the prospect of an impending global race in AI-based weaponry have led to the publication, on July 28, of “Autonomous Weapons: An Open Letter from AI & Robotics Researchers,” with over 20,000 signatories by now, calling for “a ban on offensive autonomous weapons beyond meaningful human control.” *Communications* is following up on this letter with a Point-Counterpoint debate between Stephen Goose and Ronald Arkin on the subject of lethal autonomous weapons systems (LAWS) beginning on page 43.

“War is hell,” said General William T. Sherman, a Union Army general during the American Civil War. Since 1864, the world’s nations have developed a set of treaties (known as the “Geneva Conventions”) aiming at somewhat diminishing the horror of war and ban weapons that are considered particularly inhumane. Some notable successes have been the banning of chemical and biological weapons, the banning of anti-personnel mines, and the banning of blinding laser weapons. Banning LAWS seems to be the next frontier in effort to “somewhat humanize” war.

While I am sympathetic to the desire to curtail a new generation of even more lethal weapons, I must confess, however, to having a deep sense of pessimism as I read the Open Letter, as well as the two powerful Point and Counterpoint articles. I suspect many computer scientists, like me, like to believe that, on the whole, computing benefits humanity. Thus, it is disturbing for us to realize computing is also making a major contribution to military technology. In fact, since the

1991 Gulf War, information and computing technology has been a major driver in what has become known as the “Revolution in Military Affairs.” The “third revolution in warfare,” referred to in the Open Letter, has already begun! Today, every information and computing technology has some military application. Let us not forget, for example, the Internet came out of ARPAnet, which was funded by the Advanced Research Projects Agency (ARPA) of the U.S. Department of Defense. Do we really believe AI can, somehow, get an exemption from military applicability? AI is already seeing wide military deployment.

Rather than call for a general ban on military application of AI, the Open Letter calls for a more specific ban on “offensive autonomous weapons,” which “select and engage targets without human intervention.” But the concept of “autonomous” is intrinsically vague. In the 1984 science-fiction film *The Terminator*, the protagonist is a cyborg assassin sent back in time from the year 2029. The Terminator seems to be precisely the nightmarish future the Open Letter signatories are attempting to block, but the Terminator did not select its fictional target, Sarah Connor; that selection was done by Skynet, an AI defense network that has become “self-aware.” So the Terminator itself was not autonomous! In fact, the Terminator can be viewed as a “fire-and-forget” weapon, which does not require further guidance after launch. My point here is not to debate a science-fiction scenario but to point out the intrinsic philosophical vagueness of the concept of autonomy.

Goose argues that ceding life-and-death decisions to machines on the battlefields crosses a fundamental moral and ethical line. This assumes

humans perform every life-and-death decision in today’s battlefield. But today’s battles are conducted by systems of enormous complexity. A lethal action is the result of many actions and decisions, some by humans and some by machines. Defining causality when discussing composite actions by highly complex systems is nearly impossible. The “fundamental moral and ethical line” discussed by Goose is fundamentally vague.

Arkin’s position is that AI technology could and should be used to protect noncombatants in the battlespace. I am afraid I am as skeptical of the potential of technology to humanize war as I am skeptical of the prospect of banning technology in war. Arkin argues that judicious design and use of LAWS can lead to the potential saving of noncombatant life. Technically, this may be right. But the main effort of military designers has been and will be to increase the lethality of their weapons. I fear that protecting noncombatant life has been and will be a minor goal at best.

The bottom line is that the highly important issue raised by the Open Letter and by the Point-Counterpoint articles is highly complex. Knowledgeable, well-meaning experts are arguing the two sides of the LAWS issue. To the best of my knowledge, this is the first time the computing-research community is publicly grappling with an issue of such weight. That, I believe, is a very positive development.

Follow me on Facebook, Google+, and Twitter.

Moshe Y. Vardi, EDITOR-IN-CHIEF

Copyright held by author.