BLOG @CACM

The *Communications* Web site, <http://cacm.acm.org>, features more than a dozen bloggers in the BLOG@CACM community. In each issue of *Communications*, we'll publish selected posts or excerpts.

Follow us on Twitter at <http://twitter.com/blogCACM>

DOI:10.1145/2771281

<http://cacm.acm.org/blogs/blog-cacm>

The Dangers of Military Robots, the Risks of Online Voting

John Arquilla considers the evolution of defense drones, and why Duncan A. Buell thinks we are not ready for e-voting.

John Arquilla
“Meet A.I. Joe”

<http://bit.ly/1HYARMI>
May 1, 2015

Isaac Asimov's Three Laws of Robotics, first codified in his 1942 short story “Run-around” (<http://bit.ly/1AAkKhW>), sought to steer use of artificial intelligence (AI) in peaceful directions. Robots were never to harm humans, or by inaction allow them to come to harm. Within months of the elucidation of these laws, however, an extremely primitive robot, the Norden bombsight, was being put to lethal use by the U.S. Army Air Corps. The bombsight combined a computer that calculated various factors affecting an aircraft's arrival over a target with an autopilot. Touted for its accuracy from high altitude, the Norden nevertheless tended to miss the aim point by an average of a quarter-mile. Yet pilots routinely turned over control to the machine for the final run to target.

In the 70 years since the end of World

War II, AI has advanced enormously, and the U.S. military has continued to show a steady appetite for acquiring lethal robots. Tomahawk cruise missiles, the great-grandchildren of the Norden, once launched find their own way to far-distant targets over even the most complicated terrain; yet the enemy sometimes eludes the Tomahawks, which end up killing the wrong people. The Phalanx ship-defense system is another important military robot—many missiles move too fast for human reflexes—but Phalanx, limited to close-in, Gatling-gun-like defensive fire, is unlikely to cause collateral damage to noncombatants.

For all the advances in AI, there remains a significant reluctance to embrace the notion of fully autonomous action by machines. Hence the popularity today of remotely controlled aircraft and ground combat systems. Indeed, the growth in the number of drone pilots has been explosive, and soldiers on the ground have become very attached to their machine buddies; there have

been instances when drones have been given medals, and have been ceremonially buried when “killed in action.” It is fascinating to see the great emotional appeal of the machines juxtaposed with the intellectual fear of what robots might do in the future, when they become more able to act independently.

That fear is more than just the residue of the dark tropes of the *Terminator* and *Matrix* film franchises—not to mention *Ultron*—or of the “Battlestar Galactica” TV series reboot; or, for that matter, of the worries about robots most recently expressed by luminaries like Stephen Hawking. There are real and practical concerns about autonomy and flawed machine judgment. How is a robot to determine the difference between enemy soldiers and non-combatants? This is a difficult-enough problem for human soldiers in the irregular wars of our time. Other questions are: When should a robot keep fighting or stop shooting if the foe is wounded, or is trying to surrender? In the case of a robot-piloted attack aircraft, can it discriminate between military and civilian targets adequately?

Even worse, can robots be hacked? The Iranians claim to have hacked an American drone and brought it down safely on their territory back in 2011. However it happened, they have it, and refused to return it when President Obama somewhat cheekily asked for it back. This incident should prompt us to consider the question: What if robots could be taken over and turned on their masters? A coup of this sort would require a high level of technological so-

phistication and an absolute mastery of the radio-electromagnetic spectrum, but the consequences of one side being able to do this would be catastrophic for the side whose robots were taken over.

Strategic concerns aside, on the political front, ethical and practical concerns about robots have been raised at the United Nations, which in an April 2013 report called for an immediate moratorium on the development and deployment of “lethal autonomous robots.” The report is part of a valiant effort to stave off the onset of an AI arms race, but indicators from many places are that the race is already on. In Britain, for example, the Taranis combat aircraft (named for the Celtic god of thunder) has autonomous capabilities, unlike the Predators and other types of drones we have become familiar with over the past decade. The British are also moving toward fielding robot soldiers; their humanoid Porton Man reflects exceptional sophistication of design.

The Russians are not far behind, although their Strelak sharpshooter and Metalliste machine gun and grenade launcher systems, while apparently having some autonomous capability, seem to be under fairly close human control. But don't count on it staying that way.

The Chinese military is making swift advances in both remote-controlled and autonomous systems, on land, at sea, and in the air. Their advances in naval mine warfare include weapons that can sense the type of ship coming along and move stealthily to attack it. There is also evidence of Chinese sea mines with a capability to detect and then attack helicopters flying above them.

Against these threats, the U.S. Navy is developing autonomous mine-clearing technologies, giving the undersea fight an increasingly robot vs. robot flavor.

The same is true in cyberspace, where the sheer speed and complexity of offensive and defensive operations are driving a shift to reliance on robots for waging cyberwars. This is an area about which little can be said openly save that here is yet more evidence that the arms race the United Nations seeks to prevent is well under way. That suggests it is high time for an international conference, and an accompanying discourse, on the prospects for crafting robotic arms control agreements. In this way, we can

keep alive the ideal of peaceful robotics that Asimov introduced so long ago.

There may be no way, ultimately, to stop the spread of killer robots, but at the very least they should be obligated to observe the laws of armed conflict, like their human counterparts.

Duncan A. Buell
**“Computer Security
 and the Risks
 of Online Voting”**

<http://bit.ly/1CeFQ4A>

April 2, 2015

One hears from the science community that scientists need to stand up and explain their science to the public. The usual topics—climate change, evolutionary biology, stem cell research—are outside the normal scope of most ACM members, but we in the computer sciences are the experts on one matter that has enormous impact.

In the wake of the 2000 U.S. presidential election, Congress passed the Help America Vote Act (HAVA). The nation largely turned away from older election technology and moved to computer-based systems, in some states relying entirely on the software and sometimes-arcane procedures that provided no secondary method against which to compare the tally produced by the software.

Now, as the HAVA machines age, many election officials around the country and around the world seem enchanted with the marketing hype of Internet voting software vendors and are buying in to the notion that we could—and should—vote online now and in the very near future.

Never mind the almost-daily reports of data breaches of financial organizations with deep pockets to spend on securing their computers. Never mind that governments, with shallower pockets, are routinely hacked, or that former FBI director Robert Mueller went on the record that the only two kinds of computers are those that have already been hacked and those that no one has yet bothered to hack. Election officials seem in awe of ill-defined vendor terms like “military-grade encryption.”

In a small corner of the ACM world, scientists wonder why officials are not hearing the message of computer security experts. It is time for that small corner to expand and for the membership to find a voice in hopes that election of-

ficials can be made to listen.

ACM members are moderately aware that Alex Halderman of the University of Michigan and his students (legitimately) hacked the test Internet voting software of the District of Columbia. They are probably less aware that Halderman and Vanessa Teague of the University of Melbourne recently demonstrated that vendor software being used for an Australian statewide election was subject to a standard security flaw. Teague and Halderman responded in the best manner of professionals recognizing a flaw: they alerted the Australian CERT, which alerted the New South Wales Election Commission (NSWEC), which patched the software. Then Halderman and Teague went public, but the software patch to prevent votes from being intercepted and changed in flight to election central did not happen until one-fourth of the total online votes were cast.

The response from the NSWEC was both dismissive and frightening. On the one hand, it was admitted the NSWEC had factored the likelihood of corrupted votes into their analysis. On the other hand, they targeted Teague with a formal complaint to her university, accusing her of a breach of ethics and suggesting further that a DoS attack was coming from the University of Michigan. The NSWEC then had the audacity to publish with the vendor a puff piece (<http://bit.ly/1GVp4Nr>) on online voting without mentioning the security flaw.

The Australian election is not an anomaly. Many U.S. states are toying with the notion of online voting, contracting their elections to private companies whose code has never been given a public vetting. As scientists, we would all probably rather be doing science than trying to find ways to convince the public and election officials that security online today is not up to the task of voting online today. Yet the need is critical for disseminating the hard facts about computer security and the huge risks of online voting. We must take this on as a professional responsibility. The nation and the world deserve no less than our full involvement.

John Arquilla is professor and chair of defense analysis at the United States Naval Postgraduate School. **Duncan A. Buell** is professor of computer science and engineering at the University of South Carolina.

© 2015 ACM 0001-0782/15/07 \$15.00