

BY RONALD ARKIN

Let me unequivocally state: The status quo with respect to innocent civilian casualties is utterly and wholly unacceptable. I am not in favor of Lethal Autonomous Weapon Systems (LAWS) nor of lethal weapons of any sort. I would hope that LAWS would never need to be used, as I am against killing in all its manifold forms. But if humanity persists in entering into warfare, which is an unfortunate underlying assumption, we must protect the innocent noncombatants in the battlespace far better than we currently do. Technology can and should be used toward that end. Is it not our responsibility as scientists to look for effective ways to reduce man's inhumanity to its fellow man through technology? Research in ethical military robotics can and should be applied toward achieving this goal.

I have studied ethology – animal behavior in their natural environment – as a basis for robotics for my entire career, ranging from frogs, insects, dogs, birds, wolves and human companions. Nowhere has it been more depressing than to study human behavior in the battlefield. The commonplace occurrence of slaughtering civilians in conflict over millennia gives rise to my pessimism in reforming human behavior yet provides optimism for the prospects of robots being able to exceed human moral performance in similar circumstances.

I have the utmost respect for our young men and women in the battlespace, but they are placed into situations where no human has ever been designed to function. This is exacerbated by the tempo at which modern warfare is conducted. Given this pace and resultant stress, expecting widespread compliance with international humanitarian law seems unreasonable and perhaps unattainable by flesh-and-blood warfighters.

I believe judicious design and the use of LAWS can lead to the potential saving of noncombatant lives. If properly developed and deployed, it can and should be used towards achieving that end, and not simply about winning wars. We must position this humanitarian technology at the point where war crimes, carelessness and fatal human error occur and lead to noncombatant deaths. Unmanned systems will never be able to be perfectly ethical in the battlefield, but I am convinced that they can ultimately perform more ethically than human soldiers.

I am not averse to a ban should we be unable to reach the goal of reducing noncombatant casualties; but for now we are better served by a moratorium, at least until we can agree upon definitions regarding what we are regulating and it is determined whether we can indeed achieve humanitarian benefits through the use of this technology. A preemptive ban ignores the moral imperative to use technology to reduce the persistent atrocities and mistakes that human warfighters make. At the very least it is premature.

Alternative considerations include the following: Regulate autonomous weapons usage instead of prohibiting them entirely; consider restrictions in well-defined circumstances rather than an outright ban and stigmatization of the weapons systems; do not make decisions based on unfounded fears – remove

pathos and hype while focusing on the real technical, legal, ethical and moral implications.

Numerous factors point to autonomous robots soon being able to outperform humans on the battlefield from an ethical perspective:

- They are able to act conservatively, as they do not need to protect themselves in cases of low certainty of target identification.

## Robot imperative

### The moral obligation of using AI to reduce atrocities

- The eventual development and use of a broad range of sensors will render robots better equipped than humans for battlefield observations.

- They can be designed without emotions that would otherwise cloud their judgment or result in anger and frustration with ongoing battlefield events.

- They avoid the human psychological problem of “scenario fulfillment,” which contributed to the downing of an Iranian airliner by the USS Vincennes in 1988.

- They can integrate more information from more sources far faster than a human possibly could in real-time before responding with lethal force.

- When working in a team of combined human soldiers and autonomous systems, they have the potential to independently and objectively monitor ethical behavior in the battlefield by all parties and to report any infractions that may be observed.

- LAWS should not be considered an end-all military solution. To the contrary, their use should be limited to specific circumstances. Current thinking recommends:

- Specialized missions where bounded morality applies, e.g. room clearing, counter-sniper operations or perimeter protection in the DMZ.

- High-intensity inter-state warfare, not counter-insurgencies, to minimize likelihood of civilian casualties.


- Deployment in concert with soldiers, not as their replacement. Human presence in the battlefield should be maintained.

Smart autonomous weapons systems may enhance the survival of noncombatants. Human Rights Watch considers the use of precision-guided munitions in urban settings to be a moral imperative. In effect, there may be mobile precision-guided munitions that result in a similar moral imperative for their use. Such weapons have the possibility of deciding when to fire and – more importantly – when not to fire. They should be designed with overrides to ensure meaningful human control. Moreover, they can employ fundamentally different tactics while assuming far more risk than human warfighters in terms of protecting non-combatants and assessing hostility and hostile

intent. In essence, these systems can more effectively operate on a philosophy of “First do no harm” rather than “Shoot first and ask questions later.” Building such systems is not a short-term goal, but rather part of a medium- to long-term agenda addressing many challenging research questions. However, exploiting bounded morality within a narrow mission context helps to achieve better performance with respect to preserving noncombatant life, and thus warrants robust research on humanitarian grounds. Other researchers have begun related work on at least four continents. Nonetheless, many daunting questions regarding lethality and autonomy remain unresolved. Discussions regarding regulation must be based on reason, not on fear. Until these questions are resolved, a moratorium is more appropriate than a ban. Only then can a careful, graded introduction of the technology into the battlespace be ensured.

The status quo is unacceptable with respect to noncombatant deaths. It may be possible to save noncombatant lives through the use of this technology, and these efforts should not be prematurely terminated by a preemptive ban. AI can be used to save innocent lives where humans may and do fail. Nowhere is this more evident than on the battlefield.

The feasibility of autonomous weapons is also not in question, at least for a broad class of missions that might currently be contemplated.



**RONALD ARKIN** is Regents' Professor and director of the Mobile Robot Laboratory at the Georgia Institute of Technology's School of Interactive Computing in Atlanta.

# The new weapons of mass destruction?

Building a lethal autonomous weapon is easier than building a self-driving car. A new treaty is necessary

BY STUART RUSSELL

Beginning in 2014, the High Contracting Parties of the Convention on Certain Conventional Weapons (CCW) have held meetings at the United Nations in Geneva to discuss possible limitations on the development and deployment of lethal autonomous weapons systems (AWS). In November 2017, the CCW convened a formal Group of Governmental Experts (GGE), chaired by India's Ambassador to the UN Amandeep Singh Gill, with a mandate to “assess questions related to emerging technologies in the area of lethal autonomous weapons systems.” This article reflects views shared by a great many in the artificial intelligence community. These views were expressed in an open letter on July 28, 2015, signed by over 3,700 AI researchers, and in a letter to the Obama administration written on April 4, 2016, by 41 leading American AI researchers, including almost all of the living presidents of AAAI, the main professional society for artificial intelligence. The British AI community sent a similar letter to then Prime Minister David Cameron.

The UN defines autonomous weapons as having the capacity to “locate, select and eliminate human targets without human intervention.” Some have proposed alternative definitions – for example, the UK Ministry of Defence says that autonomous weapons systems must “understand higher-level intent and direction” and “are not yet in existence and are not likely to be for many years, if at all.”

Much of the discussion at the UN has been stymied by claims that autonomy is a mysterious, indefinable property. In the view of the AI community, the notion of autonomy is essentially unproblematic in the context of lethal weapons, which is quite distinct from the philosophical context of human autonomy. The autonomy of lethal weapons is no more mysterious than the autonomy of a chess program that decides where to move its pieces and which enemy pieces to eliminate. The key is that the specific targets are not identified and approved – either in advance or at the time of detection – according to human judgment, but are instead selected by an algorithm based on sensory input the algorithm receives after the mission is initiated by a human.

The feasibility of autonomous weapons is also not in question, at least for a broad class of missions that might currently be contemplated.

All of the component technologies – flight control, swarming, navigation, indoor and outdoor exploration and mapping, obstacle avoidance, detecting and tracking humans, tactical planning, coordinated attack – have been demonstrated. Building a lethal autonomous weapon, perhaps in the form of a multi-rotor micro-unmanned aerial vehicle, is easier than building a self-driving car, since the latter is held to a far higher performance standard and must operate without error in a vast range of complex situations. This is not “science fiction.” Autonomous weapons do not have to be humanoid, conscious and evil. And the capabilities are not “decades away” as claimed by some countries.

UN Special Rapporteur Christof Heyns, Human Rights Watch, the International Committee of the Red Cross and other experts have expressed concerns about the ability of autonomous weapons to comply with provisions of international humanitarian law regarding military necessity, proportionality and discrimination between combatants and civilians. Discrimination is probably feasible in most situations, even if not perfectly accurate. However, determining proportionality and necessity is most likely not feasible for current AI systems and would have to be established in advance with reasonable certainty by a human operator for all attacks the weapons may undertake during a mission. This requirement would therefore limit the scope of missions that could legally be initiated.

Another important component of international humanitarian law is the Martens Clause, according to which “the human person remains under the protection of the principles of humanity and the dictates of public conscience.” In this regard, Germany has stated that it “will not accept that the decision over life and death is taken solely by an autonomous system” while Japan “has no plan to develop robots with humans out of the loop, which may be capable of committing murder.” BAE Systems, the world's second-largest defense contractor, has asserted that it has no intention of developing autonomous weapons, stating that the removal of the human from the loop is “fundamentally wrong.”

At present, the broader public has little awareness of the state of technology and the near-term possibilities, but this will presumably change if the killing of humans by autonomous robots becomes commonplace. At that point, the dictates of public conscience will be very clear, but it may be too late to follow them.

Compliance with international humanitarian law, even if achievable, is not sufficient to justify proceeding with an arms race involving lethal autonomous weapons. President Obama:

“I recognize that the potential development of lethal autonomous weapons raises questions that compliance with existing legal norms – if that can be achieved – may not by itself resolve, and that we will need to grapple with more fundamental moral questions about whether and to what extent computer algorithms should be able to take a human life.”

One of the “fundamental moral questions” is the effect of autonomous weapons systems on the security of member states and their peoples. On this matter, the message of the AI community, as expressed in the letters mentioned above, has been clear: Because they do not require individual human supervision, autonomous weapons are potentially scalable weapons of mass destruction: an essentially unlimited number of such weapons can be launched by a small number of people. This is an inescapable logical consequence of autonomy. As a result, we expect that autonomous weapons will reduce human security at the individual, local, national and international levels.

It is estimated, for example, that roughly one million lethal weapons can be carried in a single container truck or cargo aircraft, perhaps with only 2 or 3 human operators rather than 2 or 3 million. Such weapons would be able to hunt for and eliminate humans in towns and cities, even inside buildings. They would be cheap,

STUART RUSSELL

is a computer science professor and the Smith-Zadeh Professor in Engineering at the University of California, Berkeley.

effective, unattributable and easily proliferated once the major powers initiate mass production and the weapons become available on the international arms market. For the victor they would have advantages over nuclear weapons or carpet bombing: they leave property intact and can be applied selectively to eliminate only those who might threaten an occupying force. Finally, whereas the use of nuclear weapons represents a cataclysmic threshold we have – often by sheer luck – avoided crossing since 1945, there is no such threshold with scalable autonomous weapons. Attacks could escalate smoothly from 100 casualties to 1,000 to 10,000 to 100,000.



The considerations of the preceding paragraph apply principally to weapons designed for ground warfare and anti-personnel operations, and are less relevant for naval and aerial combat. It is still the case, however, that “to entrust a significant portion of a nation's defense capability in any sphere to autonomous systems is to court instability and risk strategic surprise.” Autonomous weapons in conflict with other autonomous weapons must adapt their behavior quickly, or their predictability will lead to defeat. This adaptability is necessary but makes autonomous weapons intrinsically unpredictable and thus difficult to control. Moreover, the strategic balance between robot-armed countries can change overnight due to software updates or cybersecurity penetration. Indeed, a nation's autonomous weapons might be turned against its own civilian population. With no possibility of attribution to an external adversary or individual, one can imagine that the nation's government would be less popular after such an event. Finally, the possibility of an accidental war – a military “flash crash” involving spiraling and unpredictable high-speed interactions among competing algorithms – cannot be discounted.

It seems likely that pursuing an arms race in lethal autonomous weapons would result in a dramatic and probably irreversible reduction in international, national, communal and personal security. The only viable alternative is a treaty that limits the development, deployment and use of such weapons and prevents the large-scale manufacturing that would result in wide dissemination of these scalable weapons.

This argument parallels that used by leading biologists to convince US Presidents Lyndon Johnson and Richard Nixon to renounce America's biological weapons program. This in turn led to the drafting by the United Kingdom of the Biological Weapons Convention and its subsequent adoption. I think we can all be glad that those steps were taken. ■

BY KIM MIN-SEOK

Conventional warfare tactics, traditionally maintained through rifles and tanks, artillery, and fighter jets, is now centered on the rapid innovation and advancement of IT, artificial intelligence (AI), avionics and cyber technology in developed countries. The battlefield is expanding to five dimensions – ground, sea, air, space and cyber. However, the South Korean military faces a triple handicap. Thus, without technical innovation, it will be difficult to maintain its combat power.

The first of Korea's handicaps is its surrounding conditions. The Korean Peninsula has achieved stabilization following the Cold War era, but there are indications of another potential clash between land powers such as China and Russia and sea powers like the United States and Japan.

Secondly, the threat posed by North Korea has increased. The North boasts an immense conventional force of 1.2 million soldiers and is becoming an actual nuclear state.

Thirdly, South Korea faces a steep “demographic cliff,” and a decrease in its number of troops is inevitable. Thus, the South Korean military is expected to shrink from its current level of 620,000 troops to 500,000 by 2022. While the decrease in manpower has little influence on the navy and air force, the biggest burden will fall upon the 483,000-strong army, which will soon be reduced to 365,000 troops. The South Korean military, especially the Army, will not be able to cope with the changing military structure and potential future threats unless it seeks innovation through drones and automation.

Although the South Korean military has pushed for defense reform since the late 1990s, it has repeatedly failed due to changing administrations, the 1997 Asian financial crisis and instability resulting from North Korea's frequent military provocations. But a current sense of crisis is pushing the South Korean government and the military toward speedy reform. The development of drones and IT technology has caused the South Korean army to establish a drone combat unit. General Kim Yong-woo, the South Korean Army chief of staff who assumed the post last year, intends to change the structure of the Army and shift to new warfare tactics based on technology. The Moon Jae-in government is backing this concept.

The US and Israeli armies were the first to deploy unmanned aerial vehicles (UAVs) for military

operations and have experience in developing norms and rules of engagement for their use. The US military has actively used drones in Afghanistan and Iraq in efforts to eradicate terrorist forces including al-Qaeda and the Taliban. Israel was the first country to use a drone militarily and is still plays a leading role in the development of weaponized drones. During a prolonged period of joint military exercises with the US, the South Korean

to replace the Songgolmae. On the division level, an indigenous version of the existing KUS-9 drone must be developed and the RemoEyes replaced. The Agency for Defense Development (ADD) will soon complete development of a MALE (medium-altitude long-endurance unmanned aerial vehicles), which is similar to the US MQ-9 Reaper and, aside from basic reconnaissance capabilities, could carry air-to-ground missiles like the AGM-114 Hellfire.

Miniature drones on par with America's WASP UAV are in the works. They will be deployed to the new “decapitation” special forces unit established last year by the South Korean Army. These micro drones provide intelligence on targets

## Droning on

### The South Korean military is catching up with a new drone army

military has observed the US army's deployment of drones and learned how to use them in combat. It also received technological support from Israel when it first developed its own weaponized drones. Even now, many South Korean companies in the defense industry develop military drones through technical cooperation with Israel.

The South Korean Army currently operates the RQ-101, or Songgolmae (wingspan: 6.4 meters), developed through its homegrown technology in 2000, along with Israeli-made Searcher IIs, while some Army corps have deployed Israel's Heron drones. While the division level commands the KUS-9 drone (2014), the regiment level uses the RemoEye-15 (2004) and the RemoEye-006 UAV (2006) was developed and deployed at the battalion level. Defense reforms emphasizing drone warfare could expand the operational area of corps and divisions by a factor of three to four in the future. This would also enable network-centric warfare based on a drone system.

Furthermore, carrying out the Kill Chain pre-emptive system operations to strike down North Korea's ballistic missiles in near real-time now requires raising the reconnaissance capabilities of forward units in a groundbreaking manner. The Army corps also needs a next-generation drone

When taking into consideration the demographic cliff, threats from North Korea and China, and the harsh reality of the need for defense reform, the build-up of a drone and robot system in South Korea becomes a realistic prospect. Korea's drone and robot technology is nearing 80 to 85 percent of the level reached by the highly advanced countries like the United States. Setting up a drone army will be facilitated by South Korea's industrial prowess in such fields as materials and battery technology as well as electronics and communications technologies. ■

KIM MIN-SEOK

was a researcher at the Korea Institute of Defense Analyses from 1982 to 1994 before becoming a journalist. He was spokesman for the South Korean Ministry of National Defense from 2010 to 2016. He is now a military and security expert as well as opinion columnist for the Korean daily newspaper JoongAng Ilbo.