GLOBAL STRATEGY FORUM

EDITION No. 22 - JULY 2020

The 22nd in our series of expert comment and analysis, by our regular commentator, General Sir Richard Barrons, Commander Joint Forces Command (2013-2016), now Co-Chairman of Universal Defence & Security Solutions, and GSF Advisory Board member. As always, the views expressed are those of the author and not of Global Strategy Forum unless otherwise stated.

The Unpleasantness Of Death By Robot

I doubt many members of the GSF have had the good fortune to watch 'The Mandalorian', the latest Star Wars spin-off, stimulating though it undoubtedly is. We would all then be familiar with IG-11, 'the masculine-programmed IG-series assassin droid' that despatches a very considerable number of hostile stormtroopers with enormous vim and aplomb. IG-11 is enabled by the integration of sensors, weapons, robotics, data, AI, a formidable on-board power source and a magic ammunition supply to operate at a pace, accuracy, precision, and resilience in combat that no human opponent can match. Its subsequent reprise (having finally been shot out of service) as a tea-bearing 'nurse droid' is a mark of endearing, frankly highly inefficient, algorithmic versatility.

But the point is that the only technical elements missing in the real world from IG-11 are the power source and the magic ammo box. We are already living with the advent of 'Lethal Autonomous Weapon Systems' (LAWS) and we need to decide what we will tolerate and how to inhibit what we will not. LAWS are coming for us even if we choose not to have our own.

We cannot be surprised that the potential of new technologies is being exploited in new weapons, that would defy the entire history of human conflict - great ingenuity has always been shown in harnessing new means to threaten, inflict or defeat harm. Denial and regret have never been sufficient responses: what can be made will be and what has been made will not be dis-invented. Albert Einstein and Robert Oppenheimer both had regrets about the development of the nuclear weapons in which they were instrumental, Mikhail Kalashnikov grieved over the deaths his rifle has inflicted on so many, and Ethan Zuckerman knows he will be damned in perpetuity for writing the code for the original pop-up ad on the internet.

Today, the combined application of data (especially, but not only, in the cloud), processing power, the Artificial Intelligence

that allows machines to 'see' and 'hear' their surroundings and respond accordingly, the connectivity that enables the 'Internet of Things' (everything that can be connected to the internet is connected), robotics, driverless car technology, nanotechnology and other forms of miniaturisation, battery improvements, solar power and new materials will all contribute to the potential for LAWS.

This is now, not years away: the robotic lawnmower that relentlessly trims the grass (though not yet at a level that delivers a Proper Stripe) could be the genesis of a machine that also executes that neighbour's cat for abluting on the carrots again, just by equipping it with some cheap sensors, a simple algorithm, and a small firearm. But as your own cat would then need a decent flak jacket, even this example shows where some wrinkles lie.

It is important to distinguish between 'uninhabited (woke)/unmanned (common parlance)' and 'autonomous'. The General Atomics MQ-9 Reaper is a hunter-killer Unmanned Aerial Vehicle with a range of 1200 miles, a maximum speed of 300mph, ceiling of 50,000 feet and flight endurance of 14 hours when fully loaded with sensors and a combination of 3800lbs of precision guided missiles and bombs. It costs around $16m, compared to around $100m for a F35 Joint Strike Fighter - although unlike F35 it cannot survive against big air defence systems. It is flown from anywhere on the planet by a crew of two, so one difference from a manned jet is that a Reaper pilot is sat in front of instruments in a comfy chair free of G-forces, in an airconditioned cabin, sipping the beverage of choice with a full stomach and an empty bladder.

But a Reaper only takes life when the pilot decides. An autonomous Reaper would have all the same capabilities, but when it connected what it sensed with its programming it would fire a missile or drop a bomb without any human intervention at that point. So do advances in technology like voice and facial recognition and their connection to data storage, AI and weapons mean the human pilot can have the week off?

Weapons that kill without a human controlling the action, machines that once unleashed apply lethal force on the basis of programming with no further recourse or control, are already in use. At least 15 navies use the Phalanx Close In Weapon System (CIWS), a radar and computer controlled gun that fires 4500 rounds of 20mm ammunition a minute to destroy approaching missiles. It can be set to fire automatically or invite a human operator to take that decision, but given the speed of an anti-ship missile (the latest shift at 4700kmh) it is most likely to be effective when 'set free'. The system is programmed to make some basic decisions about 'friend or foe', but in the circumstances these are necessarily rapid and do not involve a conversation.

We are generally content with the idea of defensive systems such as these: they destroy armed attack, protect our own lives and assets and operate in well-defined conditions with small capacity for collateral damage. They are cheap to use, a bullet costs around $30 and about 100 are employed in most engagements – the incoming missile might be charged at $500,000 and a CIWS is certainly cheaper in every sense than having an aircraft carrier sunk. The equation was a bit more complicated when these systems were dismounted ashore as the Centurion to protect

bases such as the UK airfield in Basra in Iraq from rockets, artillery and mortars, because of the potential for spent rounds to cause collateral damage. Different ammunition that destructs at tracer burn-out range ameliorated that, though this works better over desert than Kensington.

The need for the speed of decision and action that only AI can achieve also applies in cyberspace, where cyberattacks can be so large scale, so agile and potentially damaging that there is no time to ask a human if now would be a good time to switch off the power grid/water supply/factory/computer network etc. AI is the only way of acting well and in time to head off great harm in micro-seconds. Where this is defensive, protecting our Critical National Infrastructure and the lives that depend upon it, most people are supportive. But should we also adopt the advantages of autonomous cyberattack? We could build AI-enabled cyber payloads that would search for weaknesses in an opponent's CNI and once found exploit them as fast as possible. Controlling the effects of this in advance would be very difficult as so much would depend on where the cyber weapon found a weakness and how, in micro-seconds, this was exploited to cause damage. Such a weapon could not possibly judge the full second or third order consequences of, say, shutting down power to an area of military significance that also happened to include a hospital with no back-up generators. This could render the attack disproportionate, indiscriminate and unnecessary – in other words illegal – but who would we hold accountable?

At a slightly different level, Samsung built the SGR-A1 sentry gun for deployment in the Korean Demilitarised Zone, with integrated surveillance, tracking, firing and voice recognition. The military case for an autonomous sentry is that unlike a human counterpart this relentless machine doesn't nod off, doodle or get distracted by sex or beer every 18 seconds, nor want uniforms, holidays, pay, quarters, medical care, or a pension. It doesn't sneak chocolate into its ammunition pouches, it doesn't need regular training or miss a shot by failing to release the safety catch in all the excitement. An autonomous sentry also raises the ethics of machines killing people without a human pulling the trigger or sanctioning the firing: how would we know what happened, who judged a shot was necessary and on what grounds, who is in charge of the act and who is to be held accountable for the consequences? These challenges apply to fixed systems, and even more so if LAWS move independently at sea, on land, in the air and in space, killing in accordance with their programming on the basis of what they sense.

The prospect of these individual, conventional systems generally provokes a mix of horror and interest, usually in a tangle. For example, it would be possible to build a LAWS for the specific purpose of aspects of urban combat: instead of sending a soldier or a dog as the first through the breach in a wall of a building full of hostile soldiers we could send a machine. The machine would be sent in programmed to send back video footage of what it 'saw', autonomously shoot at anything it recognised as carrying a weapon with the speed of decision and reaction of IG-11. It would not be invulnerable, but unlike the soldier and the dog it would not bleed or be mourned if destroyed. It just goes in a skip. However, there is every prospect that this machine would kill many people, and quite likely some that met

the criteria of carrying a weapon, even if they were at that moment minded to surrender.

The reaction of most people to this scenario is a deep-seated reaction that we should not build these machines, not because they are not really cricket, but because they open the Pandora's box of machines killing people 'out of control' and killing 'the wrong people'. On the other hand, if this means we are back to sending in people or dogs, how do you feel about it being your son or daughter instead, because it will be somebody's? And you are definitely not keen on it being your dog.

By far the greatest concern, and the greatest impetus behind the now mature efforts by many pressure groups such as the 'Stop the Killer Robots' campaign, is not just the ethical dimensions of localised use of LAWS but the prospects of elevation to Weapons of Mass Destruction. The Internet of Things will bring many benefits to our lives, even though the wow factor of what is technically possible will be moderated over time by what is actually useful. Few of us really need a lightbulb that causes our electric toothbrush to charge up if we yawn after 2200hrs. Most of us do not want our every bathroom word, snort and scratch captured for the entire internet to digest before sending us unction for an improved crevice management regime. The Internet of Things, however, will mean that a whole range of innocuous things could directed to cause us harm and a vast number of very dangerous things could be connected to the internet.

If every device of every type is connected, and our physical location is almost always known, certainly in more urban areas, the theoretical possibility exists of an opponent capturing control of domestic systems and making them a danger to individuals and communities. This could be by cyber tools that indiscriminately cut or pollute water supplies, cause power outages or surges, or crash road and rail traffic control systems without warning or consideration of the consequences. It might be more local, such as filling your house with gas and igniting it with your toaster, or capturing your car's computers to cause a high speed accident. The malicious use of autonomous cyber tools or everyday items to unleash unpredictable, uncontrolled and unaccountable death and destruction could easily create as much harm as nuclear, biological or chemical weapons.

Easier to do, and worse in terms of consequences, would be to build millions and millions of small, lethal weapons, each with limited power and range but together a profoundly dangerous system and a highly disturbing development. These might be things like flying, cooperative micro-drones, or connected explosive floating devices programmed to target beaches, or tiny wheeled equivalents that could dominate an entire city with the risk of death or injury for a long time - especially if solar-powered. Weapons such as these would compete with small nuclear, chemical and biological weapons in their potential to cause indiscriminate harm to vast numbers of people.

It is the prospect for LAWS to become WMD that lies at the heart of the developing opposition to these weapons. They are also not immune to the problem of cyber resilience: if a state built such weapons could it be certain that they would not be captured by cyberattack and turned against

the wrong targets? And since these things are not a massive manufacturing challenge, we can assume that they will attract the interest of many flavours of well-resourced and resourceful violent extremism.

The problem, as with galloping surveillance technology, is that there is little public discourse about where this technology may take us, no consensus about what the limits should be, no means of enforcing a treaty or agreement that may be signed and should be universally abided by, and every likelihood that some of those who signed had no intention of actually complying. The UK has no intention of building LAWS, agreeing that they would violate International Humanitarian Law, but UK still opposed (with US and Russia) a move for legal regulation by the UN in March 2019. Part of the problem is the absence of any common definition of what constitutes or should bound LAWS and there is the sensible need to avoid inhibiting how AI and autonomy can be safely useful in defence. Attempts so far to draw LAWS into conventional arms control regimes have foundered.

Even if the UK does not support or build its own LAWS, it will have to expect both that other states will and leakage to non-state actors such as ISIS. It seems that just as the world tumbles into conditions of greater, more existential threat, technology will simultaneously open the door to new and potentially terrible weapons. We might be forgiven for finding this combination of risk and capability heading arm-in-arm down the same slippery slope disturbing, perhaps we will not be forgiven if we do nothing about it.

What might we do? The first thing will be to raise the debate from pressure group level to a wider civil-society and political discourse. The easiest way of telling this story is likely to be to energise our best storytellers – the film and TV industry. This subject is so richly dystopian that the best way to raise the issues to wider society is through a thrilling story. The Mandalorian is a start – but the story needs to be based on science fact, not fiction, to get the traction needed. Perhaps IG-11 could appear outside the smoking ruins of 10 Downing Street, not just in another galaxy?

Greater awareness needs to provoke more than alarm, it should cause us to invest in knowing who in the world is developing what in order to set up our defensive physical and cyber measures and to judge what nature of international consensus or intervention is needed. The clever minds developing AI do not mean to spur the construction of LAWS, but neither did Albert Einstein ever mean to see a nuclear bomb exploded over Japan once a German atomic programme was stopped.

General Sir Richard Barrons
July 2020

Commander Joint Forces Command (2013-2016), now Co-Chairman Universal Defence & Security Solutions