FOREIGN AND SECURITY POLICY 14.08.2017 | Jennifer Baker

Controlling the killer robots

We are in danger of sleepwalking into a situation where the decision to kill is ungoverned by our normal combat laws. We need to keep a human in the loop

In the very near future, robots and artificial intelligence (AI) will be able to take on the tasks we humans find too dangerous, boring or difficult. But there is one area where lawmakers might want to take a closer look about how much control we cede to the machines.

Lethal autonomous weapons (LAWS) are deadly weapons that have no human intervention. Nation states around the world – including China, the UK, Israel, the US and Russia – and even private armies are already developing these weapons. We are in danger of sleepwalking into a situation where the decision to kill is ungoverned by our normal combat laws. Neither the United Nations convention on weapons, nor current international humanitarian law is fully equipped to deal with AIs pushing the kill switch. In 2014, the International Committee for Robot Arms Control (ICRAC) asked the United Nations to come up with some sort of regulatory framework.

We are in danger of sleepwalking into a situation where the decision to kill is ungoverned by our normal combat laws

Most developments in artificial intelligence have been largely driven by civilian, not military, demand. But with the technology within grasp, how can the arms industry be held back? It can't. Like it or not, killer robots are on the horizon and we should be thinking now how to deal with them.

The push for ever more autonomous weapons is the inevitable step in a trajectory that has placed combatants at further and further remove from their targets. We've moved from handheld weapons, to guns, to missiles to remotely piloted drones – those killed are increasingly out of sight of the killers. In some respects this may be positive — fewer snap decisions made in the heat of battle, less post-event trauma for the soldiers involved, more precise targeting.

Last August, the UK's Ministry of Defence (MoD) announced a £800-

million fund for cutting-edge weapons technology over the next 10 years – including insect drones and laser guns. The idea is to get industry and academia to work together to assist the MoD, or in its words: "to anticipate the challenges of tomorrow, to gain critical advantage for our defence and security forces."

In the US the Defense Advanced Research Projects Agency (DARPA) is working with robotics company *Boston Dynamics* to produce machines that are the stuff of nightmares according to some commentators. And just last month, it awarded contracts to five research organisations and one company to develop a high-resolution, implantable neural interface – a "brain link" in other words.

While these sort of tools sound like they come from the realms of science fiction, they do still at least have some sort of meaningful human control – unlike the lethal autonomous weapons that might be developed in the future. And both the the Pentagon and the MoD have internal guidelines that require some sort of "appropriate levels of human judgment over the use of force" or "human intervention."

Yet even with remote weapons, there are questions to be answered: who designs the interface? How can the operator understand the environment in which it acts? What safeguards are in place for voice or speech activated controls?

Autonomous weapons mean fundamentally giving up the decision to kill someone to a machine. Do we really want that?

There is a shift from remote piloting to autonomous machines. In some ways, even this also makes sense. Communication links between the device and the remote pilot can be jammed. There are latency issues and remote controlled devices can more easily be hacked – robotics in general are vulnerable to cyber attacks.

But – and it's a very big but – autonomous weapons mean fundamentally giving up the decision to kill someone to a machine. Do we really want that? It's not quite the same as allowing your smart fridge to order milk when you run out.

Machine learning and artificial intelligence are not the same thing. Machine learning allows many of our everyday tools to work without specific programming — like the smart fridge above or the predictive text on your smartphone. Machine learning is required in order to create an AI. The UK has attempted to define this as "human-like intellect."

The big problem is that machine learning is very different to human learning, and even with battalions of engineers, military leaders and legal experts, autonomous weapons can only be tested for the situations we humans expect. As with all self-learning tools, bad data equals bad results. Unconscious bias in those training, testing and designing the self-evolving algorithms in these tools will also result in bad outcomes – that's fine if all you get is the wrong sort of milk, but lethal autonomous weapons have the capacity to kill people, start wars, and upset international stability.

A single autonomous weapon could initiate armed conflict with a risk of escalation. Take a very simple example, where heavy winds blow a device off course in tense border areas, even a small miscalculation in decision-making could have far-reaching consequences.

Lethal autonomous weapons have the capacity to kill people, start wars, and upset international stability

Of course many countries are developing laws for cyber defence and outlawing cyber attacks. Certainly some forms of lethal autonomous weapons, or at least their operating software might fall under some of these new laws. But the problem as Marcel Dickow, head of international security at the German Institute for International and Security Affairs (SWP) sees it is that "military robotics people and cyber people don't talk to each other."

They must start.

In February this year, the European Parliament called for the creation of a European agency to follow the development of robotics and artificial intelligence as well as EU-wide civil law rules. But these focus on civil liability rules, and the impact of robotics on health, employment, transport and privacy.

But so far, as Portuguese MEP João Pimenta pointed out "there is no reference or condemnation of the use of robotics for military and security purposes. The ideas presented do not protect the national interest in the development of the robotics sector, with support from public policies of investment, research and development. A narrow approach is also being taken on the analysis of their impact on jobs."

We need an urgent answer to the question of how society will deal with issues of liability when robots are involved in accidents, but what if the damage caused by a robot was not accidental? What if it was a deliberate strategic strike by a machine designed to inflict damage?

At the very least we need to pin down the terminology – AI, robots, drones, autonomous vehicles, cyborg, self-learning, etc. Generally speaking, bad laws are worse than no laws, particularly when we don't yet fully understand the technology. But lethal autonomous weapons will be created and the international community has a responsibility to start taking it seriously.



Jennifer Baker
Brussels

Jennifer Baker has been a journalist in print, radio, and television for nearly 20 years, the last eight specialising in EU policy and legislation especially in the tech/digital sector. Her aim is to translate EU policy-speak and tech jargon into understandable English. She likes good books and bad films.

-