Autonomous weapons systems are drones and other military robots that can recognize and attack targets such as tanks, warships or enemy soldiers on their own. Building such a robot seemed extremely difficult in the past, but at the current level of technology even countries that we do not think of as technological powerhouses could build them.

Using autonomous warbots is a subject of controversy. Could these follow the same moral and legal rules humans are required to follow? Is there something inherently wrong in using them? What would happen if these became standard in all armies? Lawyers, AI scientists and philosophers have been trying to answer these questions for some time. While consensus is fleeting, professionals agree that new regulation needs to be develop to stop autonomous warbots from being intentionally or unintentionally used to target civilians or civilian objects. Some postulate tougher restrictions, or even a complete ban; yet going for a less controversial and more politically palatable set of regulations seems best both as a matter of principle and as a matter of practice.

The history of regulating weapon use, unlike the history of weapon bans, is generally a success story. The way in which most countries conduct war today has significantly changed since the dark days of WWII and the Cold War; the barbaric actions of Russia, Syria or the Islamic State shock us because they are so out of step. Instead of anti-personnel landmines that can kill playing children decades after a conflict, today anti-tank mines no child could trigger are preferred; instead of destructive artillery barrages, smart munitions hitting within meters of their target are used. The reforms in this spirit should also encompass the emerging autonomous weapons technology, a successor to smart munitions.

It is not enough to agree the we need thoughtful regulation – such regulation needs to be created in detail, and they need to be acceptable from the perspective of various stakeholders and multiple values. My project is to take up some of this hard yet badly needed work. I will look into regulatory proposals concentrating on three especially difficult issues: 1) how to design a test verifying that a military robot will behave in compliance with relevant ethical and legal rules? 2) how should surrendering to a robot look like? and 3) how should robots comply with the principle of proportionality in military action?

Unless question #1 is successfully answered, no regulation of autonomous weapons will be possible. Unfortunately, it is simultaneously particularly hard to answer. Autonomous weapons' software is and will be based on machine learning – grown and taught rather than handcrafted. As such, it can be unpredictable in highly unintuitive ways. On the other hand, there can be only so many test conducted. How to prove in a finite number of tests that a robot will behave in an infinite number of scenarios – this conundrum calls for focused philosophical analysis.

International Humanitarian Law mandates accepting the enemy's surrender, while simultaneously adjusting this rule to the realities of combat with different technical means. A supersonic fighter surrenders in a different way than an infantry soldier does. Adjusting the emerging technology to the mandate, and the mandate to the existing technology presents a challenge, but also an opportunity. Today, surrendering is a tense and ambiguous affair. Maybe adjusting the procedure for the use of warbots may lead to general improvements.

Both law and ethics forbid attacking civilians and civilian objects. They also forbid attacking military targets if accompanied by disproportional collateral damage. Algorithms have a hard time reasoning about such matters, as this requires a general understanding of how the world works, which they lack. It's clear humans have to pre-program such answers into the bots. To be able to do this, we need to solidify our presently vague standards on this matter. I will seek a way to do this.

Apart from undertaking these three specific topics I will also monitor and evaluate regulatory proposals made by other authors. It is highly probable that autonomous weapons will soon become ubiquitous in all types of military operations. Regulating them is thus an urgent matter with enormous practical consequences, and I my research contributes to this worthy task.