Addressing ethical questions of modern AI warfare

The use of Artificial Intelligence in war is only a question of time. But it must be regulated in terms of international humanitarian law

It might be hard to believe, but the smartphone in your hand has more computing power than a Eurofighter. This is set to change – through initiatives such as the joint Franco-Spanish-German Future Combat Air System (FACS). Yet the point stands: for a long time now, technological innovation has not been the one-way street it once was from military to civilian use; rather, impetus goes in both directions. Civilian markets, after all, are gigantic and lead their military counterparts. So, in the armed forces, the talk is now of bundling resources: that means coming together as allied nations to pool troops, weapons and reservists. Joining up like this also means using large amounts of data – data which must be produced, captured, exchanged and evaluated in order to draw far-reaching conclusions. This is the kind of mission which can only be completed by networked operational command structures capable of dealing with the complexity of today's battlefield – a set of requirements which, during the war in Ukraine, has repeatedly taken even experts by surprise.

Yet, for all the ways in which war is fought is changing, the classic variables – firepower and mobility – remain key. The former has, following the development of modern explosives and, above all, nuclear weapons, basically reached its physical limits; the latter, however, is currently gaining in importance as digitisation and automation open up wholly new options. This, in turn, is creating new challenges for military intelligence as, on tomorrow's battlefield, enemy positions – and those of its own troops – will be changing constantly. Unless IT systems are closely coupled and used intelligently, armies will find themselves in an impenetrable fog of data; but if things go well, the complete suite of all real-time battlefield data regarding enemy assets will be pooled, categorised and then immediately directed to the weapons systems best suited to engage. Then again, any equal foe will be doing exactly the same on its end, meaning that split-seconds will become decisive. Here, when it

comes to identifying and destroying enemy targets, artificial intelligence (AI) far superior to any human weapons-user will take centre stage. Especially in the early-warning systems deployed in response to nuclear threats, the consequences could be dramatic.

Speed vs. control

As a result, those building and administrating these kinds of 'systems of systems' are facing a dilemma. Either they harden their systems by implementing a high degree of automatic operation, thereby exploiting their potential speed of reaction to the greatest possible degree but sacrificing the potential to control, administer and, if required, reconfigure them; or they opt for weaker, slower systems by including interfaces allowing operators to intervene in processes, gaining power and flexibility in how they are run. This tension between speed and control can only be resolved by taking the term 'system of systems' for what it means: in future, and once again in an ideal set-up, through every stage of conflict, each command level will as early and as far as possible anticipate the functional spectra of processes at the next lower level in a comprehensive manner and inform the system to be configured adequately to respond.

What everyone needs to be aware of, is that the reliability of this kind of system is subject to technical limits and can only ever deliver a high probability of success, not a guarantee.

So, what are the implications? Taking the concrete example: Exactly which hypersonic missiles an adversary has at its disposal and whether there is an armed conflict with this adversary at all is determined at the political-strategic level. At an operational level, parameters to detect those missiles are precisely defined so that, if a state of conflict arises, hypersonic weapons are identified before or, at the very latest, when they approach our airspace and are then neutralised. At the tactical level, systems are trained on these parameters allowing them to identify and, without needing any extra authorisation, down enemy hypersonic missiles. What everyone needs to be aware of, however, is that the reliability of this kind of system is subject to technical limits and can only ever deliver a high probability of success, not a guarantee.

In order to maintain the advantage of the speed of reaction on the tactical level, the operational order to open fire must be wholly assigned down to the tactical level. This can be achieved without risking a loss of

control as long as the technical parameters for opening fire are defined and set prior to the engagement – and as long as these technical parameters can later be adjusted in response to changes in the way the situation is assessed. Analogously, the same coupling should also apply to the manner in which the operational level interacts with the strategic level. The way this would look would be that the military high command, at the top of the system of systems, would potentially be aided by AI in analysing, optimising and stabilising decision-making at theatre and battlefield level on the basis of the operational situation; this situation would also be used to determine the settings for administering and training the AI programs controlling fire at a tactical level – making sure that said fire is reaching its target, i.e. enemy positions and assets. The extent to which this whole system works depends heavily on the ability of its component systems both at the centre and on the periphery to communicate in such a way that the system as a whole reacts optimally to alterations in conditions and influencing factors.

Addressing ethical concerns

In view of this, the development of AI-supported weapons systems by questioning established chains of military command and political leadership already poses some serious ethical issues – issues to which humanitarian concerns and the limits of international law add further dimensions as, thus far, there is no provision for AI. To date, negotiations on expanding the framework of the UN arms control treaty CCW (Convention on Certain Conventional Weapons) to cover AI have never got beyond the issue of defining 'human control'. Yet, considered rationally, this is not an unsolvable task.

For a start, it would be best to avoid the overloaded term 'autonomous': that is because 'weak AI' does not take decisions and/or action but rather assists judgement in the same way as a telescope helps researchers see or a grabber tool helps people pick up things. In research, AI is generally defined as an academic discipline aiming to endow machines with human-like capabilities of perception and comprehension with a view to developing technical systems, especially systems for processing information (definition from *Handbuch der Künstlichen Intelligenz*). Even the highly-complex visual system in the human brain is not 'autonomous', and so there is no point hoping for an international treaty on banning or regulating autonomous weapons systems for as long as it is unclear what 'autonomous' actually means and, consequently, what is to be banned or regulated.

Can the presence of

Secondly, 'human control' can be defined approximately as follows: AI-supported

characteristics constituting 'military objectives' be reliably left to an AI system to decide?

weapons systems must be deployed so that they are permanently subject to the control of those holding office in governments, high commands, army headquarters and battlefield operations centres. In order to benefit from the time saved by automation without sacrificing this control, orders must be given beforehand and, importantly, parametrisation carried out and administration conducted prior to deployment at top-layer command level. If this is adhered to, AI applications stand not only to provide technical support in target acquisition and neutralisation, but also to offer guidance in respect of legal concerns.

Article 36 of the additional protocol I of the Geneva Convention states that new weapons must be subject not just to technical but also legal certification: so, it will not be enough for AI to measure the basic physical parameters of a site (e.g. location, size, temperature) for it to be declared a valid military target. Rather, Article 52 states that 'attacks shall be limited strictly to military objectives. In so far as objects are concerned, military objectives are limited to those objects which by their nature, location, purpose or use make an effective contribution to military action [...]'

Can the presence of characteristics constituting 'military objectives' be reliably left to an AI system to detect? In the case of a hypersonic missile, most probably yes. In many other cases, however, there will be factors which AI cannot identify and process; politicians, strategists and operational staff will need to decide one way or the other at their level. As such, the three countries involved in the FCAS project, France, Spain and Germany, have an opportunity to show a convincing route out of the current CCW logjam by applying the nuanced model for AI-assisted weapons systems and making this known globally, providing ongoing public information on how the system is being developed.

If there is one lesson history teaches us: it isn't the amount of human and material resources deployed which is decisive in victory and defeat but rather the quality of political and military

One objection remains to be dealt with: aren't armies who respect the constraints of international law at a disadvantage against those who ignore them? It's an easy assumption to make, but there are good arguments against it. After all, what potential disadvantage is an army risking if it avoids unnecessary civilian losses? In fact, doing so can turn into an advantage – when the fighting is over, if not before. A second, even weightier argument is derived from

the concerns of public relations and the justification for military action: the leeway accorded by international law is expansive and expandable; as such, in a conflict, each side must assume that the other will use its own interpretation of international law to maximise its room for manoeuvre; and for this reason, above all, international law does not simply apply but rather must be made to apply and then enforced. This is one of the most important areas of diplomacy and strategic communication: 'warfare by lawfare', as the dictum has it, sums up the importance of public relations, of being seen as a law-abiding and reliable partner. In any armed conflict, the conviction that there is a *just* cause is crucial – and more important than the number of troops which can be deployed.

War is the superlative of armed conflict, an extreme state in which everyone is willing to do everything: and that is precisely why modern international law condemns it. If there is one lesson history teaches us – and of which current events in Eastern Europe are now reminding us – it isn't the amount of human and material resources deployed which is decisive in victory and defeat but rather the quality of political and military culture.



Hartwig von Schubert
Hamburg

Dr Hartwig von Schubert is a Protestant theologian and was Military Dean at the German Bundeswehr Command and Staff College for 15 years. Today, he teaches applied ethics at two universities in Hamburg.