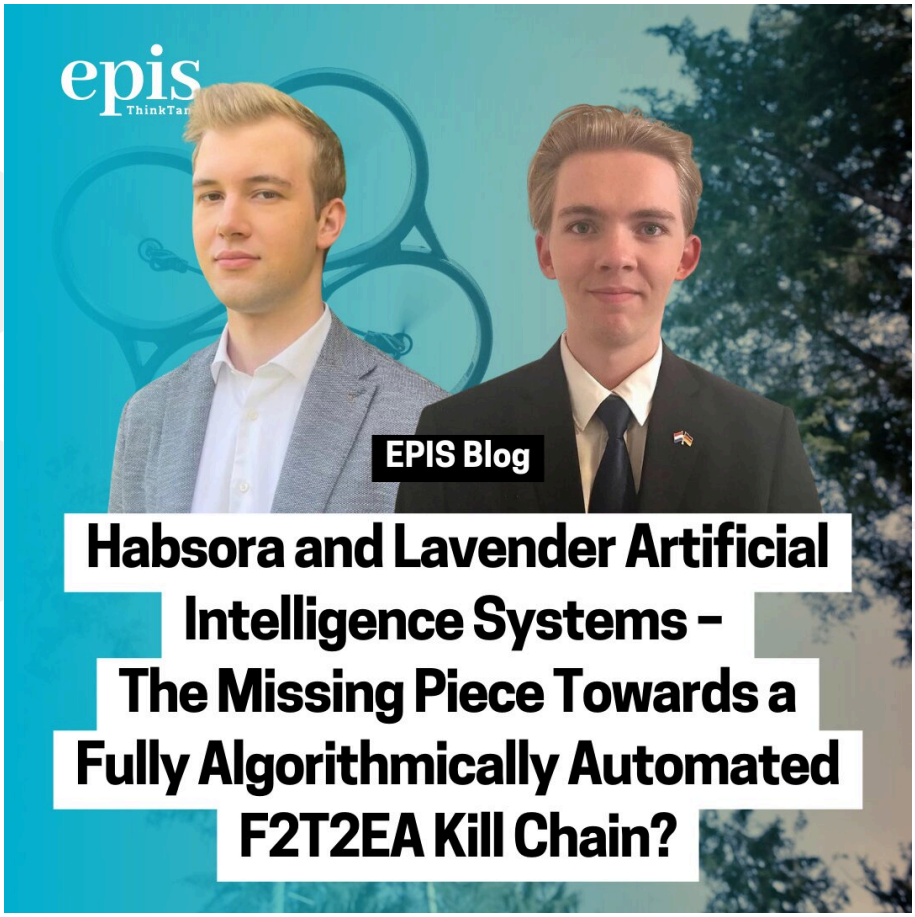
Habsora and Lavender Artfical Intelligence Systems

19 AUGUST 2024

Bjoern Kuehn and Julius Kurek



Habsora (הבשורה) and Lavender (לָבָנדֶר) Artificial Intelligence Systems

The Rise of Habsora and Lavender: Navigating Conflicts and Instability with the Help of Advanced AI Systems

Thus far, artificial intelligence (AI) has been increasingly utilised in modern warfare. Several countries, such as Israel, have started to incorporate AI systems to enhance the process of target selection. This technological

advancement aims to improve the efficiency and accuracy of identifying potential targets by analysing vast amounts of data from various sources. Starting in 2019, the Israeli government announced the creation of a so-called 'targeting directorate' to generate sufficient targets prior to any conflict for the Israeli Defence Force (IDF), especially the Israeli Air Force (IAF). Previously the IDF and IAF faced a shortage of targets during past conflicts in the Gaza Strip, such as "Operation Guardian of the Walls" and "Operation Protective Edge". The targeting directorate comprises hundreds of soldiers, military officials, and data analysts who aggregate data from various sources — drone footage, intercepted telecommunications, open-source information, and data from monitoring the movements and personal behaviour of individuals and larger entities within the Gaza Strip and the West Bank. Both media and IDF sources claim that AI is used by the targeting directorate to process the aggregated data and then generate targets at a much higher pace than human analysts could effectively do under the circumstances of severe hostilities. Similar to the investigations by "+972" and "local call" on the Gospel, investigative journalists also surfaced another AI system in early 2024 that was specifically built in order to target individuals by using personal historical data. However, the exact introduction dates of these systems remain unknown and classified.

Operation Iron Swords: From 50 Targets a Year to 100 Targets a Day

The "Gospel" (Habsora) AI system produces bombing targets for specific buildings and infrastructure in Gaza, working in conjunction with other AI tools. Notably, the specific usage of the term "Gospel" implies a biblical connotation of infallibility and ultimate authority potentially attributed to the Israeli system, reflecting its trusted and authoritative status within the IDF. Thereby, the connotation underscores the system's critical role in justifying and executing military strategies, much like the unquestioned truth of the religious gospel. The gospel itself is divided into two subsystems. The whole process begins with data collection via the Alchemist system. The collected data is then categorised and analysed through the so-called fire factory. The latter categorises the targets into one of four categories. The first category consists of tactical targets, such as armed militant cells, weapons warehouses, launchers, and militant headquarters. The second category includes underground targets like Hamas or Islamic Jihad tunnels beneath the Gaza Strip. The third category, which garners the most media attention and public outcry, includes the homes of Hamas or Islamic Jihad operatives. Lastly, power targets comprise residential and high-rise buildings typically occupied by civilians.

Picture: Mohammed Al-Masri/Reuters. Edit: Björn Laurin Kühn

According to the "Dahiya doctrine", these power targets are to be bombed to exert pressure on the local population, who are then expected to

pressure Hamas operatives. Once all of the available information has been categorised, it is processed by the Gospel, which then suggests possible ammunition, usually so-called dumb bombs, and subsequently calculates potential collateral damage in the case of an airstrike. After the target has been analysed and confirmed, the last step is human approval, which is arguably the process most scrutinised by the international community.

Algorithms of War: Precision Targeting of Individuals in Modern Warfare and the Implications of International Law

While the Gospel is used to generate infrastructural and military targets, Lavender is used in order to specifically target high-level and low-level operatives. About two weeks after October 7th, the IDF started to adopt a list of potential targets provided by Lavender after a random sample check had found an estimated 90% accuracy in identifying an individual's affiliation with Hamas. Ultimately, Lavender is used in order to scan information on individuals that is available through subsystems like the Alchemist. Subsequently, these individuals receive a likelihood score from 1-100 according to the probability that they are affiliated with Hamas and Palestinian Islamic Jihad. Thus, after generating targets, targeting lists are fed into other interconnected automated systems like "Where's Daddy?" which tracks and links certain individuals to homes and then recommends a weapon for the IDF to use on the target, mostly depending on the ranking of the operative. During the first weeks of Operation Iron Swords, Lavender supposedly identified at least 37,000 Palestinians as potential targets. Similar to the processes of the Gospel, the last step is generally human approval to strike the identified target when they have entered their family's residences.

While AI technologies like Lavender offer sophisticated tools for military targeting, it is widely debated whether such systems can uphold the principles of distinction and proportionality if human approval only acts as a "rubber stamp" for the machine's decisions. While international law has increasingly become intertwined with modern warfare, its relation to AI has become somewhat blurry in the past few years. Yet, as already identified, it remains clear that the battlefield is becoming increasingly autonomous, while international regulation stagnates. Whether such Lethal Autonomous Weapon Systems (LAWS) will be regulated in the future will depend on which side prevails in the disputed conflict between ethical concerns and the overall military effectiveness of AI.

Kill Chains, LAWS, and the F2T2EA Model

In October 1996, the U.S. Air Force's chief of staff, Gen. Ronald R. Fogleman, stated "In the first quarter of the 21st century, it will become possible to find, fix or track, and target anything that moves on the surface of the Earth". While utopian in his own time, his promise has certainly become a reality with the advent of advanced AI systems, enabling these capabil-

ities to be achieved even quicker and more accurately. This advancement is particularly evident in the series of tactical processes and decisions involving weapons use, referred to as a kill chain. The "kill chain" serves to conceptually capture the process of combating an enemy entity. It begins with finding the target and encompasses every subsequent step up to its eventual destruction. One model to structure the kill chain internally is the so-called F2T2EA model, which is divided into six steps. Finding the target is a matter of intelligence, which may come in the form of surveillance or reconnaissance. Once the target is identified, its precise location needs to be determined (fix) and kept track of as the appropriate weapon is selected (target). Afterwards, the target can be engaged and, once the attack has been carried out, its effectiveness might be assessed.

Historically, every step in this chain involved humans making decisions, or multiple humans working in unison, with limited and uncertain knowledge, tight timelines, and within highly dynamic environments, often with severe human consequences. AI systems like Habsora and Lavender enhance kill chain operations by reducing uncertainty, increasing the speed of decision-making as well as the volume of data that can be taken into account, improving decision assessments, and taking the human operator out of harm's way.

Generally, discussions of automation in the public domain focus on LAWS. These are systems that are able to perform the last five steps of the kill chain themselves. LAWS fix a suspected target's location and execute the remaining steps independently, but they do not perform broader intelligence work and target generation. In turn, the systems discussed above are not weapons systems as they do not actually engage any targets. They are reportedly able to identify targets by analysing intelligence from a range of sources, determine a target's precise location or home and keep track of it and finally create a collateral damage assessment and recommend a weapon to engage it. They automatically generate targets, thus ticking off the first four steps of the kill chain before passing on the information to a fighter aircraft or a drone, which has to simply service the target.

Towards an Automated Kill Chain?

With so many of the steps of the kill chain already taken, it is easy to imagine the step of engagement being carried out by an automated system that is simpler and more attainable than the lethal autonomous weapons systems that are commonly being conceptualised. This would mean that instead of LAWS performing some steps, there is, for the first time, a genuine potential for a change towards a number of separated automated systems performing separate steps in conjunction. Neither system is itself autonomous, and their coordination results in a fully automated kill chain rather than LAWS.

For now, there still remains the step of human approval of the individual attacks. However, if the target generation systems are indeed as capable as

speculated, then this represents not a hard technical or engineering limit but a legal and moral agenda item that could be scratched today rather than within a couple of years. Consequently, the first fully automated kill chain might be a matter of integrating existing technologies, optimising systems, and adjusting procedures. The involvement of humans and human decision-makers could soon be limited to the mere set-up of initial parameters and systems as well as the maintenance of hardware on the ground.

EPIS