

ASSER
INSTITUTE

Centre for International & European Law



Designing International Law and Ethics
into Military Artificial Intelligence

REPORT

Decision-Support Systems and Human-Machine Interaction

REAIM Summit Break Out Session

15 February 2023

The Hague / The Netherlands 2023

REAIM

Responsible AI in the Military domain **Summit**

Report prepared by:

Taylor Kate Woodcock (Asser Institute, University of Amsterdam)

1. Introduction

a. Background

On 15 February 2023, the Asser Institute's DILEMA Project¹ organised a panel discussion at the inaugural REAIM Summit on Responsible AI in the Military Domain, hosted in the Hague by The Government of the Netherlands, and co-hosted by the Republic of Korea.

The summit provided a platform for all stakeholders to discuss the key opportunities, challenges and risks associated with military applications of AI. It hosted over 2,000 participants from 100 countries, including foreign ministers and other government delegates, as well as representatives from knowledge institutions, think tanks, industry and civil society organisations. The DILEMA panel discussion on 'Decision-Support Systems and Human-Machine Interaction' was one of 35 break out sessions across the two days of the Summit.

The panel was moderated by **Dr Bérénice Boutin**, head of the DILEMA Project, Senior Researcher in International Law, and Coordinator of the Research Strand on Disruptive Technologies in Peace and Security (Asser Institute). The panel consisted of **Dr Neil Davison**, Senior Scientific and Policy Adviser, Arms and Conduct of Hostilities Unit, Legal Division (International Committee of the Red Cross (ICRC)); **Dr Jurriaan van Diggelen**, Senior Researcher in Military Human-Machine Teaming (TNO) and ELSA Lab Defence Project Leader; **Klaudia Klonowska**, Researcher in International Law (Asser Institute), PhD Candidate (University of Amsterdam); and **Dr Gregor Pavlin**, Senior Scientist and Program Manager (Thales Nederland B.V.).

b. Scope and Thematic Focus

The session centred around issues relating to decision-support systems and human-machine interaction with respect to applications of AI in the military domain. Whilst international diplomatic discussions on military applications of AI have tended to revolve around autonomous weapon systems (AWS), other algorithmic tools may also emerge in this sphere and warrant further scrutiny. In particular, algorithmically-supported decision-support systems (DSS) raise many crucial legal, ethical and technical challenges. These tools utilise data processing capabilities to mine vast bodies of data to support and supplement decision making in military operations. Due to their characteristics and functions, as well as the circumstances and context in which they are deployed in military operations, DSS raise important questions regarding human-machine interaction (HMI).

¹ Designing International Law and Ethics into Military Artificial Intelligence (DILEMA). The DILEMA project is funded by the Dutch Research Council (NWO) Platform for Responsible Innovation (NWO-MVI), project number MVI.19.017. Website: www.asser.nl/DILEMA.

As highlighted in the GGE LAWS Guiding Principles: ‘Human-machine interaction, which may take various forms and be implemented at various stages [...] should ensure that the potential use of [...] systems based on emerging technologies [...] is in compliance with applicable international law, in particular IHL’.² As such, this session was well placed to explore and develop insights on what HMI means in the context of DSS. The discussion was guided by the key question of: How to ensure an adequate ‘quality and extent of human-machine interaction’³ in relation to DSS, in particular to ensure compliance with international law?

In order to answer this question, a multi-disciplinary group of experts discussed how the use of DSS can erode the exercise of human control, and reconfigures the role and place of human decision making and human agency. Furthermore, the session explored the implications of the discussion on DSS and HMI for the design and regulation of such systems.

To facilitate the discussion, the following guiding questions were addressed by panellists:

- (1) How are DSS used (or envisaged) in the military context? What are their key characteristics, and differences from the autonomous weapons narrative?
- (2) Why is it important to move the debate on military applications of AI beyond the narrow scope of autonomous weapons and ‘meaningful human control’?
- (3) What is HMI, how to conceptualise it as a notion and research field, how does it relate to MHC and how is it different?
- (4) Why is it important to focus on pre-deployment (design) stages in order to ensure effective human-machine interaction?
- (5) What are risks associated with DSS? How can systems supposed to support human decision making end up replacing decision making of users/operators?
- (6) How can effective HMI design contribute to fostering the ability of human-machine teams to act in compliance with international law, in particular international humanitarian law?
- (7) What aspects of military decision making, or domains of military AI applications, should we pay particular attention to when developing and deploying DSS, and what safeguards can we put in place?
- (8) What are the implications of and recommendations from the discussion for the international regulation and governance of military AI?

² Guiding Principles affirmed by the Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons System (‘GGE LAWS Guiding Principles’), UN Doc CCW/MSP/2019/9, Annex III, <http://undocs.org/CCW/MSP/2019/9>, Principle (c).

³ GGE LAWS Guiding Principles, Principle (c).

2. Summary of the Session

(1) How are DSS used (or envisaged) in the military context? What are their key characteristics and differences from the autonomous weapons narrative?

In order to conceptualise DSS, the first speaker, Klaudia Klonowska, made a key distinction between decision making and decision-support systems. Decision making systems have the capacity to be connected with the physical elements for enforcing the decisions, for instance software embedded in a weapon system or other platform, where the decision making process can be completed. Whether or not there is a human in the loop or other role for humans, the algorithm has the capacity to connect the processing of data with the use of force. Most prominently, decision-making systems include the so-called 'autonomous weapon systems' in current debates.. On the other hand, decision-*support* systems are embedded in the targeting cycle and are part of the process in a chain of actions that leads to certain military conduct. However, unlike a decision making system that leads to physical engagement, these systems require a human to transform the decision into action. The decision-support systems of most interest in the military domain are those that transform raw data into actionable intelligence and that go beyond showing what is happening on the battlefield, but also add an extra layer of analysis into military decision making processes. Whilst decision making systems have monopolised the discussions around military AI and autonomy at the international level, decision-support systems are not as much a part of that discussion.

Gregor Pavlin further added that there are largely two phases in the decision making process when it comes to the collection and use of AI methods for the extraction of actionable information delivered to the user. The first is for enhancing situational awareness, and the second is for reasoning about possible actions and making suggestions to achieve certain goals. As DSS become part of decision making processes where certain elements are automated, it is of key importance to look at the interface between humans and machines. In order to emphasise the human, you have to understand how DSS fit into socio-technical systems.

(2) Why is it important to move the debate on military applications of AI beyond the narrow scope of autonomous weapons and 'meaningful human control'?

Neil Davison explained that when it comes to these technologies, focus on 'meaningful human control', human-centred AI or human agency is critical because it serves as an anchor to legal obligations and ethical responsibilities owed during warfare. A certain level of control is given up with autonomous weapon systems, which select and apply force to targets without human

intervention. With DSS it is a different situation: it is about how the human and the machine interact and how measures can be taken to ensure that human judgment is retained. There are important practical aspects to this, such as cross-checking information, so that the user does not solely rely on the machine output.

Additionally, there are technical aspects to consider in the deployment and use of DSS, such as the need to test the system and understand its limits in order to be able to question and challenge its outputs. A further consideration is the role of the human using the system, whether they have time to deliberate on the output it generates and if they have the relevant information required to make the context-specific judgments demanded by the law.

On the point of meaningful human control, Klonowska added that for broader applications of military AI beyond autonomous weapons, this concept has serious limitations. It is crucial to understand what the interactions between the human and the machine entail. The concept of meaningful human control has been used to trigger certain reactions in policy debates. However, it assumes human superiority, with the notion of 'control' elevating the human to a position where it is assumed that humans are able to control the machine and its algorithmic processes. It also separates the human, reflecting the idea that the human is independently judging the situation, instead of recognizing that it becomes impossible to discern and separate the influences of machines and humans in the decision making processes.

Finally, it also tends to reduce issues to a binary question of whether or not there is control. In policy debates when there is discussion about designing for control, it is treated as if there are certain conditions that once met, will ensure control is exercised. However, in practice we see that many of the interactions between humans and machines are very context dependent. For instance, depending on whether a user has one, three, or five hours to make a decision using a DSS, this will significantly impact the sorts of interactions that can take place. It is therefore extremely important to consider these aspects when designing these technologies.

(3) What is HMI, how to conceptualise it as a notion and research field, how does it relate to MHC and how is it different?

Jurriaan van Diggelen pointed out that the field of HMI has been around for a long time. However, previously we used to interact with tools in a relatively straightforward manner, according to what we wanted to do with them. Now we have AI tools that are involved in processing data in new ways.

A concept commonly used is 'human-machine teaming', where both the human and machine collaborate to act together. Here, you can align the goals of the humans and machines. From the human perspective, it is about whether you can trust the AI system. You can look at it as if you were

getting a new employee. If you give it an assignment that it performs well, you trust it a little bit more and grant it more autonomy, or in other cases you might take more control. These interactions are important and take place over time, when you can see what the AI system is capable of. The human is also learning in this context. Compared to the past, the interface to enable these interactions is much more complex, but is nonetheless necessary in order to use AI safely and responsibly.

Pavlin further added that human factors research is a wide field involved in investigating cognitive abilities and limitations, behaviours and mental processes. As mentioned, previously it was used to optimise our interaction between devices and machines. However, now we see a completely new problem when looking at naturalistic decision making processes. This process is often based on teamwork. If you want to understand how decisions are made by humans, you therefore need to understand how many tasks are carried out by humans. This is critical when developing human-machine interfaces. It is very important to understand where the humans involved have certain tasks, how they would perceive certain outputs of an algorithm, under what conditions they would trust these systems, and where in the process AI is being used. This has huge consequences for the acceptance of the systems, as well as guarantees you can give regarding their use.

As mentioned earlier, AI can be used to enhance situational awareness or to give suggested actions, and the former is much easier for current systems. Typical examples are systems detecting certain objects using satellite imagery. This information then can filter down to human decisions makers, who will have the ability to catch errors. However, if you have a machine that is selecting actions, there is significantly less human buffer. Understanding this is critical to put the debate into context, as understanding what specifically a system will be used for will drive the legal and ethical considerations.

(4) Why is it important to focus on pre-deployment (design) stages in order to ensure effective human-machine interaction?

On the design of AI systems, Davison indicated that there are a number of things that should be kept in mind, including who the human user is, what application the decision support system will be used for, and the applicable legal obligations and responsibilities in that particular context. However, rather than simply emphasise the important role of designers, the design, development and use of these systems should be viewed as a whole, to ensure that the design and the context of use do not prevent meaningful human engagement in legally mandated decisions. We must avoid what could be essentially a human as spectator, or 'rubber stamp', when it comes to decisions that end someone's life. There is significant tension in that machine learning tools, by definition, are designed to replace humans in order to carry out certain tasks. And so the use of AI-based decision support raises

questions about the role of humans and human judgement in warfare, which is extremely important to protect.

Regarding designing for effective human-machine interaction, van Diggelen remarked that for DSS there are many ways to facilitate communication and the provision of advice to the user. Put simply, it is how the human is given the output of the algorithm which has the highest probability for a certain action. Alternatively, when there are multiple actions available, the system could present each of the pros and cons of a particular course of action. This way allows the user to have a more critical stance on what the AI is doing and to also overcome possible mistakes.

Moreover, the notion of explainability is very important for machine learning systems. If the user cannot explain how the system arrived at a decision, they are not properly engaged in the interaction, which could result in the user blindly following the system's advice. The downside here would be if so much information is provided that the user does not have time to consider it all. This trade-off must be managed by the provision of sufficiently brief explanations and allowing the user to interact in a meaningful way.

Klonowska further built on these points by explaining that it is important to keep in mind that all these stages of development, deployment, and use are not separate. We may go back and forth between them to see how the human is interacting and assess the impacts, then go back to developing and redesigning the system.

From a legal perspective, it is crucial that the design stage should involve legal considerations from the very beginning. Legal and ethical considerations should be brought together with military strategic objectives to guide the discussion about the design of specific systems. Thinking about how targets are defined, what parameters are implemented, for instance, based on the geographical area of the battlefield, and many other aspects of how the system is designed have important implications for the legal decisions that are made afterwards.

(5) What are risks associated with DSS? How can systems supposed to support human decision making end up replacing decision making of users/operators?

Pavlin highlighted a number of issues to consider. First, whilst there are always biases with AI-enabled systems, human actors also have their own biases. There will always be biases present. Second, in order to make the system as reliable as possible, it will always depend on how it is used, the doctrines of the organisation and the types of people using it. Context is very important. Third, there are also difficulties with explainability. If there are concerns that the tempo at which the system operates is so high that the human cannot check it, then the question is what sort of information the system should deliver to the user to make them satisfied with the results or to

corroborate something. Context will also play a role here. As systems are used over time, things change and the system may no longer understand the environment. The human user must be aware of this and ready for it. On the machine side, we are trying to input mechanisms flagging when the system identifies it might not be as reliable. Changes in context are also important with regard to legal and ethical aspects.

Van Diggelen further pointed out that care should be given with respect to the level at which AI processes information. There are difficult decisions in the military context that frequently require trade-offs. Either a commander risks the safety of their troops to achieve high military effectiveness, or they can be cautious to protect their people, but not complete the mission as well. It is the human user that must make these evaluative judgments at this level. In order to promote meaningful human-machine interaction, a system can be used to calculate things like safety or mission effectiveness, but should not go into the moral domain to make value-based decisions.

(6) What are the risks and implications for compliance with international law, in particular international humanitarian law, of deploying DSS that do not ensure adequate HMI?

Davison emphasised three risks when it comes to DSS and compliance with international law. First, there is the loss of human judgment. The output of an algorithm is never equivalent to human legal judgments, you cannot code these into a system. You only ever have a technical proxy that is supporting a legal judgment. For instance, a system that identifies that someone is carrying a weapon or that detects the shape of a tank is not itself making a legal judgment about whether or not it is lawful to attack that person or tank at that moment. That is highly context-dependent and hinges on many factors, including the behaviour of individuals and presence of civilians and so on. As such, there is danger in over-reliance on these systems and a false equivalence being made between human legal decision making and machine processing.

Second, speed is a major issue when it comes to military AI systems. It is often stated that the increased speed AI may offer for decision-making will be useful for decision makers and provide a military advantage. But an equally important question to ask is whether this is also an advantage for civilian protection and compliance with international humanitarian law. There are important questions about whether the requirement for human judgment in applying the law places some limitations on the speed at which these systems can operate and facilitate decisions. It is worth asking ourselves whether international humanitarian law sets a speed limit.

Third and finally, bias and lack of transparency are also key aspects to consider. As was mentioned earlier, it is not possible to remove all bias from these systems, so it is important to be aware of this limitation. Based on biases already well-known with AI systems, there are risks that their use in the

military context could expose certain civilians to more danger. Some targeting practices are already concerning, such as targeting men of a certain age who are carrying weapons in a certain area; this is not a lawful basis for targeting someone. If you encode an algorithm that learns to do that, you have essentially encoded an unlawful practice into the machine. These are all issues to think about regarding the use of DSS for military decisions on targeting and the use of force. We should be cautious about assuming that the outputs of these systems will always help to improve military decision making.

Existing international humanitarian law already sets constraints for the use of any technology in warfare, so that is the place to start in terms of understanding where the limits already lie. It is important to think about the applications that pose the greatest risks. Decisions on the use of force are the foremost example, however there are also other decisions in conflict with serious consequences for people's lives. Whilst constraints certainly exist already in the law, it is an open question whether there may be additional specific constraints needed in future. Certainly, going back to discussions around autonomous weapons, the ICRC has been clear that new international rules are needed to address specific concerns. Whilst not all autonomous weapons are controlled by AI, AI-controlled autonomous weapons do compound the problem, raising the prospect of unpredictable systems that should be prohibited.

Klonowska suggested that the issue of the impact of deploying DSS for compliance with international law can be viewed on two levels. On one hand, we can talk about granular context-specific applications. There are many difficult questions to be answered regarding context-specific changes in the development, deployment, and use of AI systems: what is a sufficient level of accuracy? Are false positives the same as collateral damage? If systems generally have a high accuracy rate but this is impacted by difficult military conditions, how to overcome the system missing important objectives? Whilst militaries want to capture as broad a range of military objectives as possible, we must ask where this leaves us in terms of civilian casualties as a result of the use of machine learning systems. When it comes to accuracy rates, we must understand what accuracy really represents for us and what we are comparing it to. Is the accuracy of a system compared to a human with a gun, or something else? If we are not able to answer these questions, designers will have to deal with them.

On the other hand, on the broader level, with international law we are trying to regulate violence and limit the consequences of warfare for civilians. So even if we have DSS systems that can be used in compliance with IHL provisions, we should still consider the broader implications for violence in general. This is a very big question remaining. As was pointed out by Dr Agnès Callamard at the REAIM opening plenary session, it is very important for us to think of these broader implications of AI for conflict situations.

(7) What aspects of military decision making, or domains of military AI applications, should we pay particular attention to when developing and deploying DSS, and what safeguards can we put in place?

Pavlin suggested that how to treat a certain DSS will depend upon what sort of application it is used for. It could be a fully automated loop, but in not such a critical mission or dangerous task. Or we could have a simple situation assessment contributing to the first part of decision making loop, for example, automated analysis of satellite imagery. On the one hand, automation for decision support in the military domain can be more easily made reliable and integrated into decision making processes in a robust way. On the other hand, automated decision making is not just more critical because there are no humans making decisions, but also because it is also more difficult to develop from a technical perspective. This requires far more complicated models and testing. Rather than simply stating that this is not possible, we should instead agree to finding an efficient way of dissecting the decision making problem. This is not just technically, but should also involve the user, who must be aware of how the system may function. This includes questions on how errors from a certain AI component might cascade down to the system and whether or not the user is able to catch it. Moreover, how this application will be understood by the legal and ethical experts should be understood at the beginning. Ultimately, we are talking about life cycles: it starts with inception, design, implementation, testing, and operations. We need some kind of way of talking about the same problem in a way that everyone across all of these elements can understand. Instead of introducing new rules and regulations, we need a standardised way of analysing the problem. Additionally, interdisciplinary teams should be involved in design processes, with all relevant stakeholders, including lawyers and ethicists.

Klonowska further argued that whilst many safeguards will be necessary, one of particular importance is an expansive interpretation of the duty to review weapons means and methods of warfare under Article 36 of the First Additional Protocol to the Geneva Conventions. Such an interpretation would include DSS that are critical in the targeting process in this review. Though they are not themselves weapons, their impact on the means and methods of warfare is substantial.

From a technical perspective, van Diggelen indicated one thing we should continue monitoring is how users end up using these technologies. They can be creative in their use and a designer cannot control that through the design process. In addition to monitoring use, we should also make sure to have metrics in place so that you can measure concerns. We have discussed the risks regarding explainability and transparency, but how do we quantify that? We need to do so to be able to test and improve these systems.

(8) What are the implications of and recommendations from the discussion for the international regulation and governance of military AI?

Pavlin concluded that it is very important to bring all these components together. Again, it is the user, the technical part, as well as the ethical and legal considerations. A major question is whether we should have a flexible way of solving a specific problem in the given context. It may be unlikely that with a set of regulations we can guarantee that we will have ethical machines. There are always loopholes and many of these regulations might also block development. It may be instructive to look at the field of aviation for inspiration, where there are certified systems and a standardised way of talking about them. Whether it is the producers of avionics, the plane authority or the airlines, they talk about the same sets of problems. You have to understand the application and then have a standardised way of talking about it. There may also be some additional factors to consider, for instance requiring legal regulation for certain technologies so that best practices can be followed and all stakeholders are involved.

Klonowska highlighted that a way to bring the discussion forward is to focus on the way humans and machines work together and the socio-technical interactions that result. This way we can begin to understand how they create the military decision making process together, rather than separating and idealising certain capabilities of humans.

Davison urged caution for government, militaries and companies developing these systems, particularly the sorts of applications of AI for decisions on the use of force discussed during this break out session. It will also be necessary to collect and assess information on the current use of these systems, as this is not merely a theoretical discussion – these systems are already being used today. This should include information regarding how these systems are used, what impact they have and the humanitarian and legal challenges we face as a consequence. Ultimately, it should be clear that legal decisions and ethical responsibilities in war cannot and must not be outsourced to software, no matter how computationally sophisticated it becomes.

Finally, van Diggelen stated that this is a problem that will not go away and with no easy solution. The analogy made with safety research in aviation is a good one, as it was introduced in an inherently dangerous domain. As we see in that field, there are many things we can do to address the challenges of technology. In that case, it was about training pilots, making aeroplanes safer, developing infrastructure, having certification of materials and more. It isn't simply that you can regulate once and the issues are solved, there are still ongoing concerns around safety. We should also think about AI the same way and continue to try and make these systems more ethical.

The audience also engaged the panellists on several issues:

Question: First of all, if it is all about context, then how can we approach dealing with these systems if circumstances just change? Secondly, when talking about responsible AI, we really mean that we are responsible for the use of these systems. This includes a range of ethical and legal issues. What do we do if we face an opponent who does not do that, but has very different ethics or a very different legal system? Do we stick to our approach or do we also create a kind of adjustable ethics to respond to opponents who might utilise AI to attack?

Davison responded that a good place to start is the Geneva Conventions, to which all States are party. In this legal framework there are obligations on all States and how they engage in armed conflict, how they conduct attacks, and how they make decisions on who or what to attack. It is important to have a discussion about what interpretations of international humanitarian law mean in practice and to try and build some common understanding in relation to AI applications in decision-making. This it is certainly difficult as with many multilateral discussions at the moment, but it is an important one to pursue. We certainly should not go along with a concept of adjustable ethics or adjustable law. There are certain international rules that have been agreed upon and we must build confidence that these will be respected in the use of new military technologies.

Klonowska added that how many international legal rules are formulated takes into consideration these changing dynamics, so it is not necessary to abandon them because of new challenges. When the Geneva Conventions were drafted there was a big discussion about the open-ended rules like proportionality and reasonableness. There was huge debate about whether States should make strict rules or afford commanders with more discretion within reasonable limits. In the end commanders were given more leeway in order to deal with changes in circumstances. This means that we have a body of rules that can account for context-based differences and will not require us to respond in an unethical way, either.

Question: One thing that seems to be becoming increasingly important from a legal review perspective, but also from a broader AI assurance perspective, is test and evaluation. When you talk about these kinds of really complex decision making systems, I often use the example of the Palantir system called Gotham, which is a very complex systems that not only creates situational awareness, but also identifies and recommends different possible courses of action that will have the least degree of risk. How do you test and evaluate systems like that? And how can we also take into account the human-machine interaction in the test and development phase? Can you test human interaction, and could this require some sort of human-machine certification? We would not only test and evaluate the system and decide whether that is lawful in isolation, but also look at the operators and how they understand system failure modes, how they understand the expected behaviour of the system, etc.

Van Diggelen suggested there is a paradox going on, as you need to test certain situations but you do not know how context might change and lead to new situations. So in testing, there needs to be as many situations covered as possible. After that it will be a matter of trusting that the system is capable of operating. This is a different way of thinking about verification and validation because normally you can achieve high confidence in how a system works, but for AI that will just not be possible. The human in the loop can also cause unforeseen situations. There are metrics for measuring human interactions with a system. For instance, you can also understand what the system is doing and measure how the human learns over time to work with the system. In traditional verification and validation, there are separate disciplines, but for AI we should merge them.

Pavlin added that the question of how to test human interaction can be difficult. If you were able to know everything about how a system functions, like in the aviation industry, that makes testing easier, as you can calculate in advance how things will go. But the problem with AI is that it relies on statistical analysis, which creates uncertainty. If you want to properly evaluate this and have statistically significant results that you say should be correct in 99.9% of cases as an operational requirement, you would have to test the system in almost all possible situations. This is a huge task with huge numbers, particularly when you go to high-level decision making. If the system is used for situational awareness and recognising the environment, it can be even sometimes application dependent. You have many cases when this works. But the moment you go into decision making where you have to recognise things and have an accurate model of the world, understanding the consequences for different courses of action is a huge task.

3. Key Take-Aways

- It is crucial to focus on the interaction between humans and AI systems to properly consider the opportunities and challenges of decision-support systems.
- Decision support systems in the military context raise risks regarding the influence they have on human judgments, the speed at which these operate, challenges relating to explainability and transparency, as well questions around how to deal with system errors.
- The stages of design, development and use must be considered holistically in order to ensure effective human-machine interaction that respects legal and ethical requirements.
- International law provides applicable legal limits to the use of DSS by militaries. It is always humans that must engage in legal and ethical decision-making, AI systems may only ever support these decisions.
- When used for targeting, clarification is required on design issues for DSS around how targets are defined, what parameters are implemented, what constitutes a sufficient level of accuracy and how to deal with collateral damage that results from the use of these systems.
- Systems need to be designed, tested and monitored with the human user, relevant application and applicable laws and ethical norms in mind, with metrics in place to quantify concerns such as explainability and bias.
- Interdisciplinary teams reflecting a variety of relevant stakeholders across the development, deployment and use of algorithmic DSS should work together when algorithmic decision-support systems are embedded in military decision making.