

THE INTERNATIONAL
C2 JOURNAL

VOLUME 4, NUMBER 2, 2010

*Requirements for Autonomous Unmanned
Air Systems set by Legal Issues*

*Tony Gillespie
Robin West*

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE DEC 2010		2. REPORT TYPE		3. DATES COVERED 00-00-2010 to 00-00-2010	
4. TITLE AND SUBTITLE Requirements for Autonomous Unmanned Air Systems set by Legal Issues				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Defence Science & Technology Laboratory (Dstl), UK				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Unmanned Air Systems (UAS) are an accepted part of the military inventory and it is anticipated that they will become more autonomous in the future. This paper examines the problem of raising autonomy levels whilst still meeting the requirements of the Laws of Armed Conflict (LOAC). A top-down approach is proposed, starting with LOAC and deriving requirements for more autonomous UAS. The method should ensure that technology developments to raise autonomy levels will have acceptable methods of use. It shows that fully autonomous weaponized systems may never be acceptable. The LOAC are used as capability requirements. Engineering requirements for autonomous UAS are derived from them. The next step in the systems engineering process requires an approach that turns qualitative criteria into quantitative ones. A three-component model of the human decision-making process is used to derive sub-system requirements and the essential technologies for autonomous generation of commands within a legal framework are identified. The implications for technology development are discussed. There is not an intention to develop fully autonomous weapon-carrying systems, but the techniques presented here should provide criteria to decide whether a command decision can be made autonomously or by a human.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 34	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

THE INTERNATIONAL C2 JOURNAL

David S. Alberts, Chairman of the Editorial Board, *OASD-NII, CCRP*

The Editorial Board

Berndt Brehmer (SWE), *Swedish National Defence College*

Reiner Huber (GER), *Universitaet der Bundeswehr Muenchen*

Viggo Lemche (DEN), *Danish Defence Acquisition and Logistics Organization*

James Moffat (UK), *Defence Science and Technology Laboratory (DSTL)*

Sandeep Mulgund (USA), *The MITRE Corporation*

Mark Nissen (USA), *Naval Postgraduate School*

Ross Pigeau (CAN), *Defence Research and Development Canada (DRDC)*

Mink Spaans (NED), *TNO Defence, Security and Safety*

Andreas Tolk (USA), *Old Dominion University*

About the Journal

The International C2 Journal was created in 2006 at the urging of an international group of command and control professionals including individuals from academia, industry, government, and the military. The Command and Control Research Program (CCRP, of the U.S. Office of the Assistant Secretary of Defense for Networks and Information Integration, or OASD-NII) responded to this need by bringing together interested professionals to shape the purpose and guide the execution of such a journal. Today, the Journal is overseen by an Editorial Board comprising representatives from many nations.

Opinions, conclusions, and recommendations expressed or implied within are solely those of the authors. They do not necessarily represent the views of the Department of Defense, or any other U.S. Government agency.

Rights and Permissions: All articles published in the International C2 Journal remain the intellectual property of the authors and may not be distributed or sold without the express written consent of the authors.

For more information

Visit us online at: www.dodccrp.org

Contact our staff at: publications@dodccrp.org



Requirements for Autonomous Unmanned Air Systems Set by Legal Issues

*Tony Gillespie and Robin West
(Defence Science & Technology Laboratory [Dstl], UK)*

© Crown copyright 2010. DSTL/JA43191. Published with the permission of the Defence Science and Technology Laboratory on behalf of the Controller of HMSO.

Abstract

Unmanned Air Systems (UAS) are an accepted part of the military inventory and it is anticipated that they will become more autonomous in the future. This paper examines the problem of raising autonomy levels whilst still meeting the requirements of the Laws of Armed Conflict (LOAC). A top-down approach is proposed, starting with LOAC and deriving requirements for more autonomous UAS. The method should ensure that technology developments to raise autonomy levels will have acceptable methods of use. It shows that fully autonomous weaponized systems may never be acceptable.

The LOAC are used as capability requirements. Engineering requirements for autonomous UAS are derived from them. The next step in the systems engineering process requires an approach that turns qualitative criteria into quantitative ones. A three-component model of the human decision-making process is used to derive sub-system requirements and the essential technologies for autonomous generation of commands within a legal framework are identified. The implications for technology development are discussed.

There is not an intention to develop fully autonomous weapon-carrying systems, but the techniques presented here should provide criteria to decide whether a command decision can be made autonomously or by a human.

Introduction

Unmanned Air Vehicles (UAVs) and Unmanned Air Systems¹ (UAS) are now an accepted part of the military inventory. However there is considerable public debate about their use in current operations especially when there is loss of life.

Research programs are developing ways of introducing more autonomy² into decision processes. Examples can be seen in the results from the UK Systems Engineering and Autonomous Systems Defence technology Centre (SEAS DTC) conferences (www.seasdtc.com). It can be postulated that technology will develop to the stage where there is no human intervention between issuing mission-level orders and their execution. An example mission could be a UAS surveying a given land area, looking for a particular target type and reporting them to a commander when found. Subsequent strikes would meet LOAC as a human would take the strike decision. It would be a small technical step to make a UAS which would fire a weapon at the target without reference to its commander. This technical step would represent a large legal change except in very narrow, strictly-defined circumstances. There has been much debate about the ethics of such systems, with no clear agreement about how they might be used. (Borenstein 2008, Quintana 2008, Finn 2008, Lin et al. 2008).

1. A UAS includes the UAV, its control station, operator, ground stations and communication links.

2. Autonomous systems act on results from their own processing of instructions from external sources; without necessarily involving human operators after initiation. Automatic systems are directly controlled by either a human or quantified input parameters with no interpretation by the automaton.

There are two approaches to aligning technology and the law:

1. Finding ways to fit emerging technologies into the law as a program evolves.
2. Expressing the legal and ethical constraints as requirements to drive system design.

The first approach is essential, but the legal considerations can place strong constraints on the operational exploitation of the military capability.

This paper considers the second approach, using well-known systems and capability engineering techniques to provide the methodology. This should reduce the capability constraints on new technologies. We assume that increased autonomy will happen, resulting in more decisions being made by the machine part of the UAS. The existing legal framework is used to establish the constraints on decision-making and where there are limits to autonomy in UAS. This can be summarized as deriving the technologies which must be implemented to deliver *the art of the acceptable rather than the art of the possible*.³

C2 Considerations

The use of force in modern warfare is governed by international laws comprising treaties and customary rules (Boothby 2009). These are collectively known as the Laws Of Armed Conflict (LOAC). Individual nations reflect their policy, legal and operational constraints in Rules of Engagement (ROE), giving clear criteria for decisions made by humans in their command chain. Different ROEs may be set for different phases, times and locations during the campaign.

3. This expression was coined at a meeting of the NATO SCI-186 Research Task Group in 2009.

Current policy is that legal responsibility will always remain with the last person to issue commands to the military system. There are assumptions that the system's principles of operation have already been shown to meet LOAC and that it will behave in a predictable manner after the command is issued. With long-endurance systems and complex scenarios, this person will need to supervise it to ensure that its actions meet the applicable ROEs. This creates a new, more symbiotic, relationship between man and machine.

Autonomy can be considered to be the introduction of machine intelligence into decision making in the Command and Control (C2) system. However, since autonomous systems take decisions beyond the human operator's direct supervisory sphere,⁴ there must be clarity and certainty in the limits of the autonomous systems' roles. We must recognize that machines are not well-equipped to deal with ambiguity whereas humans are. This point has been stressed as a limitation for Unmanned Combat Air Vehicles (UCAVs) (Burrige 2005). The UK Ministry of Defence (MOD) has no intention to develop systems with no human intervention in the C2 chain, but there is the desire to raise the autonomy level of its UAS.

It can be seen that research advances may lead to large changes in military capability (technology push), but the problem of deriving legal methods of use remains.

4. This is not speculation, as the Unmanned systems safety guide for DOD acquisition states in paragraph 1.3:

As UMSs evolve and increase in their level of autonomy, a system operator or human controller may no longer be a valid assumption; control may be completely relinquished to the UMS. Systems may use man-to-machine or machine-to-machine control.

The Problem for Autonomous Unmanned Air Systems

The legality of weapon release, even for manned assets, is often queried, even with a clear hierarchical C2 chain, with humans at the critical points in the decision loop such as that shown in Figure 1 which is taken from (NATO SCI-186).




Figure 1. Architecture View of UAS Closely Coupled with a Manned Kinetic Strike on a Target

The widespread use of networks is changing all aspects of military operations and the C2 architectures. Arns (2008) discusses possible distributed C2 structures in NATO. Figure 2 shows a C2 structure derived from Arns (2008). A requirement for an effect is met by the Combined Air Operation Centre (CAOC) allocating the necessary resources to the local commander's direct control for a specified period. At the end of this time, which may be short, control is passed back to the CAOC. In this concept of use a constant C2 chain for the duration of a mission will not be valid. However, the requirement for

clear legal authority in the military command structure will remain, regardless of the duration of a transient command chain and the level of autonomous, machine decision-making. How this allocation of authority and control will be achieved for an Autonomous UAS (AUAS) is yet to be defined.




Figure 2. Command Chain and Mission Control for UAS

AUAS requirements must encompass release into service as well as operational use. All types of air operations must be considered including use in military, controlled and uncontrolled air space.

In addition to the above technical questions, interpretations of international law evolve over time, which must be taken into account in any consideration of autonomous control.

Legal Criteria

Regulatory Considerations

UAVs are covered by the same regulatory regime as manned aircraft but currently must operate in segregated airspace. This is supplemented by extra regulations to ensure that the separation of pilot and platform does not add any extra hazards, for example by ensuring that the UAV has predictable responses to loss of control by its operator. It may also be necessary for each UAV to have an operator available in case of emergencies. It should be noted that regulatory authorities are working with others to find a way to allow UAS and manned aircraft to operate in the same airspace.

It has been suggested (Myers 2007) that situations can arise where it is unclear whether the legal liability for inappropriate weapon release lies with the pilot, the design authority or the regulatory authority. However, if the logic of the current processes is maintained, then the responsibilities of the designers will have been discharged once the UAS has been certified by the relevant national air authority. The safety case will, however, have to address the risks in an autonomous system. It should be noted that weapon clearance is solely to ensure that weapon carriage and release does not harm the weapon-carrying platform. The responsibility to clear a new weapon system from the perspective of the UK's LOAC obligations rests with the Development, Concepts and Doctrine Centre (DCDC). The responsibility for legal use of weapons and delivering non-kinetic effects will lie with the command chain during the operation. The commanders will have to know that they can confidently predict the performance of the AUAS.

The US Department of Defense (DOD) has published a set of safety precepts (DOD 2007) to guide the more general Unmanned System (UMS) design process. However, the safety issues addressed are dominantly regulatory, ensuring that there is clarity in defining safety-critical and mission-critical functions.

The DOD document does start to look at wider issues and uses a three-part model of the decision process similar to that discussed in Section 5 below. It introduces the concept of an *authorized entity*, defined as: “An individual operator or control element authorized to direct or control system functions or mission.” This term will be used from this point onward as it describes the function without implying a human or machine instantiation.

It should also be noted that engineering design and certification is based on the As Low As Reasonably Practical (ALARP) principle. This recognizes that it is impossible to make a completely risk-free system.

Laws of Armed Conflict

This paper assumes that UK military forces are deployed in a theater with an international armed conflict covered by LOAC or for authorized peace-enforcement and peace-keeping operations, and that the missions are undertaken for legitimate purposes.

There are two further complicating issues: nations usually operate in military coalitions who may not have common ROEs for that conflict; and the local population may have different cultures with consequent differences in their legal framework.

There are four underpinning legal principles in the LOAC: military necessity, humanity, distinction, and proportionality. They are addressed directly by UK in its interpretation of LOAC and stated

in Joint Service Publication 383 (JSP383 2004). Remotely piloted vehicles are considered to be aircraft, so UAS will come under the same sections of JSP383 as manned aircraft.

The next four sections give a brief summary of the intention of each tenet and gives AUAS design requirements for a system to meet that tenet. This is treating the tenets as a set of military capabilities in a systems engineering process. At this stage we use the term “authorized entity” as there is no need to separate human and machine-made decisions. What matters is the responsibility and authority for a decision.

Military Necessity

The main point is that any force used can be, and is being, controlled. Unnecessary force cannot be used, so wanton killing or destruction is illegal. We must define who or what is controlling the AUAS with its effects; and their authorization.

The design requirements derived from this principle are:

- There must be a clear and unambiguous command chain which controls and limits the actions of AUAS at all times.
- The command chain will have discrete and distinct authorized entities which can interpret and act on commands that meet the “proportionality” requirements below.
- The type of information required by each authorized entity to make decisions shall be stated clearly and unambiguously.
- When an authorized entity does not have all necessary and sufficient information for a decision, there must be acceptable alternative decision-making processes. Timings within the command chain must allow sufficient time to follow this process.

- Human commanders must have clear intervention points and criteria for overriding decisions made by the UAS.
- The status of the communication links between nodes must be known and there must be contingency processes for all types of link failure.
- System behavior following an interruption in the control link must be predictable.

Humanity

This criterion forbids the infliction of suffering, injury, or destruction not actually necessary for the accomplishment of legitimate military purposes. It is based on the notion that once a military purpose has been achieved, the further infliction of suffering is unnecessary. It is closely related to the “proportionality” tenet below.

The design requirements derived from this principle are:

- There must be a method of assessing the effect of the applied force in sufficient time to prevent the use of unnecessary force.
- There must be a means to assess whether targeted hostile forces remain a threat as defined at the start of the mission.
- The last opportunity to stop or divert kinetic and non-kinetic effects must be known to the appropriate authorized entities.
- There must be a method to confirm lethal decisions shortly before weapon release using updates to knowledge of the target area between the start of the mission and weapon release.

- There must be one or more acceptable alternative aimpoints for weapons that can be guided to their target by the AUAS after release. This is to be used if the original target is found to be unacceptable. The last time for the change of aimpoint to be effective must be known to the AUAS.

Distinction

Since military operations are to be conducted only against the enemy's military objectives and armed forces, a clear distinction must be established between: the armed forces and civilians; and between military objectives and civilian objects. Provided the criteria for classification of objects as military or civilian are clear, this is an easier problem than the identification of specific military asset types. The commander will of necessity make decisions in good faith on the basis of the information that is available to him at the time that he makes his decision. The important point is that as much relevant information as practical must be available and it must be clear, definable and intelligible as it may be needed later to justify the action taken.

The design requirements derived from this principle are:

- The basis of identification of hostile forces and materiel must be clear.
- Ambiguity in applying distinction criteria to the scene should be identified.
- Levels of uncertainty in identification must be presented to the authorized entity and there must be a way of incorporating this in the decision process.

- There must be an understanding of the normal state of the target area without the presence of hostile forces. It will include identification of objects that must not be damaged or destroyed. This will need amplification for insurgencies where the hostile forces are deliberately merging into the non-military environment.
- There must be a way for an authorized entity to assess potential collateral damage adequately at the target area before it makes its decision.
- There must be a clear link between the level of authority vested at each point in the command chain and the information available to it.
- The timing and integrity of information about the target area must be known by each authorized entity basing a decision on it.
- There must be a definition of which decisions need an audit trail for the information available at the time it was made and the means to create it.

Proportionality

Under this principle, the collateral civilian losses expected from a military operation should not be excessive in relation to the anticipated military advantage. JSP383, paragraph 2.7 is particularly significant for this paper if we consider UAS to be “smart weaponry”:

[The military planner] needs not only to assess what feasible precautions can be taken to minimize incidental loss but also to make a comparison between different methods of conducting operations, so as to be able to choose the least damaging method compatible with military success.

This is a very subjective area involving many parts of the C2 chain with military experience, legal experience, and knowledge being brought to bear during decision-making. It clearly highlights the difference between quantitative and qualitative decisions and the need for human decision-making.

The use of force outside narrow limits may escalate a conflict to a level that is unacceptable for the aims of the campaign. For example, execution of decisions made autonomously, based on assumptions of a high-intensity campaign may provoke an escalated hostile response if the campaign, or that part of the campaign, is at a lower intensity level. Clarity of ROE, the way they are interpreted, and their accurate translation into the decision making logic is essential to avoid this problem.

The design requirements derived from this principle are:

- Proportionality criteria shall be set by human mission planners in a way that can be interpreted unambiguously by each authorized entity in a manner that allows it to make its decisions.
- All authorized entities must know which weapons are available to them and the limits of their authority to use them.
- All authorized entities must know the effects of the weapons available to them.
- Whenever a non-human authorized entity cannot interpret the available information and meet the proportionality criteria at, or above, a pre-determined confidence level, it must refer the decision to a higher authorized entity in the command chain.
- When an authorized entity refers a decision to another one, it must transmit the basis of its referral decision to that entity. There will need to be a recognized format for transmission of this information.

Decision-Making Process

Command and Control Chains

The deduction from the points above is that the legality and public acceptability of the use of UAVs for military purposes will be based on the quality of the precautions and associated decisions taken in the command chain to ensure that the four legal tenets above are followed. This is consistent with JSP383 which places responsibility for attack on the person who chooses the way the attack is carried out. Those without this discretion, but follow orders to attack, also have responsibility to cancel or suspend it, if they believe that the criteria for legal use of force are not met. The requirements above set criteria for the basis of decisions and the authorization required to make them

The problem that now arises is to identify which of the authorized entities in a UAS C2 chain can become non-human and still meet the above requirements. This step requires an understanding of the human decision-making process and a concept of the generic sub-systems which would collectively deliver a more fully autonomous system.

Assumptions about Human Abilities

There is an assumption that if properly trained humans made all the decisions, the process would inherently have clear points of accountability as every authorized entity is a human. It could be argued that the process to release an AUAS to service then becomes one of ensuring that the UAS makes the same decision as a human with the same inputs. The flaws in this argument are that humans may be incapable of making safe decisions in some circumstances and two humans may make different decisions in the same circumstances.

We must recognize that the process is one of seeking to ensure that the precautions required by law are properly reflected in the decision making logic of the UAS.

Humans are well adapted to make subjective, qualitative decisions whereas machines make good quantitative ones. Therefore, we must understand the human decision process before we can transfer human decisions to autonomous systems and still meet legal criteria.

Artificial Intelligence (AI) is clearly relevant and will be of increasing importance to UAS but in a military C2 context AI systems will probably remain as decision aids for a human commander. If an AUAS that can learn whilst conducting a mission were to be built, it is unlikely that its behavior would be sufficiently predictable to meet the military necessity requirements above.

The questions which now arise are:

- How do we make an autonomous system recognize that it does not have sufficient information to make a decision which can be justified under current ROE?
- Can an autonomous system decide what information it needs to make a decision and decide how to find it?
- How can the autonomous system ensure that the humans in the C2 chain have sufficient information for them to make a better, more informed, decision?

Three-part models of decision-making processes are well-known, Thoms (2009) has developed one to capture the human cognitive contributions to delivering the system's purpose. It is used here to derive a method to convert the human role to one provided by an autonomous decision making system. The three cognitive capabilities are:

1. *Awareness* – Perceiving the current operational position and context. This is the assimilation of all available sensor and other information relevant to the UAS and its mission.
2. *Understanding* – Recognizing the relationships in the information, and their significance.
3. *Deliberation* – Choosing between the various options available, based on an understanding of them and their consequences. It includes making the decision within the known constraints and acting on it; i.e., it is a decision point in the command chain.

Figure 3 shows Thoms's three cognitive capabilities. The input data to the Awareness block has been separated into: sensor data on the current situation; context data from pre-mission sources and their updates; and ROE constraints. Dynamic mission changes, such as being targeted by hostile weapons will change the applicable part of the ROEs. Thoms's cognitive capabilities model provides a useful context for military systems as it partitions the decision process into sub-systems which can be specified in engineering terms. This has been carried out in general terms here.




Figure 3. Three-Part Model of Human Decision-Making Process

Appendix A gives engineering requirements for each of the three capabilities as: function, inputs, and outputs. These provide a conceptual basis to assess whether a specific technology can be used to meet the requirements and whether a technology needs more development work for legally-justifiable autonomous decision-making. Each function will need a clear test strategy when implemented in a system that can be linked directly to the safety case.

Discussion of the Requirements

General Requirements

One possible aim of UAS research is to develop an AUAS meeting all the requirements above and in Appendix A without human intervention. The analysis presented in this report takes a step toward this, but also shows that there are severe limitations on the use of any such systems. The derived requirements clarify the following points which must be included in any AUAS that is to enter UK service:

1. Specification of the information needs at each authorized entity. This has three aspects: generation of data; creation of information; and location and timeliness of the information.
2. Recognition that an authorized entity may be unable to make a decision and must transfer responsibility to a properly authorized one.
3. Most command decisions have subjective components. This contrasts with control-system decisions which are mainly quantitative and lend themselves to automation.

The complexity of many scenarios is such that there will always need to be a human at some point in the control chain. However, criteria have been derived which make it possible to ensure that any decision which needs human intervention will be referred to them at the correct time.

It should be noted that interpretation of the law evolves as well as technology.

Rules of Engagement

Overarching ROEs are set for any campaign. They are interpreted in detail at the ground-based mission-planning stage for the assets and weapons being used. Currently, they are written for human interpretation by humans. ROEs have to be interpreted for each UAS taking into account the UAV's differences from manned aircraft and the role of its human operator. In a similar way, the interpretation will have to be refined for each specific AUAS, taking into account the technical implementation of the decision-making entities.

It has been shown above that ROEs are sources of information for the authorized entity. The problem that arises for a non-human authorized entity in the AUAS is having this information in a form that the machine can interpret. This necessitates quantitative interpretation and identification of those subjective factors which cannot be interpreted by the technologies in that AUAS. There is no fundamental problem with carrying this out, but it can be anticipated that it will lead to developments in the way ROEs are drawn up. The results will provide the basis for limiting the authority level of the non-human authorized entities.

Derivation of information requirements with confidence levels for classifications should be achievable. They are specified in Appendix A as Information Exchange Requirements (IERs) as these are a well-known system requirement. They will also give criteria for: making

a decision and issuing appropriate commands; requesting more information in order to make a decision; or referring the decision to another authorized entity and giving it sufficient information for a rapid, informed decision.

As far as the authors are aware, no-one has tried to identify which information in a typical set of ROEs can be quantified and the tolerances on the resulting values. This needs to be carried out to establish feasibility of this approach and determine the level of ambiguity of the different ROE constraints. This will make clear which types of decision will need human intervention, regardless of the UAS's autonomy level.

Implications for UAUS Procurement

Procurement relies on correct requirements definition and the specification of every equipment and service used by the armed forces. System requirements must be produced at a sufficiently detailed level for a contractor to deliver the desired product. The discussion here is intended to highlight areas where these may be different to those traditionally considered. Thoms's three cognitive capabilities in Figure 3 provide the conceptual framework for this discussion.

The requirements above and Appendix A are not sufficient for a contract. However, they can provide the basis for a User Requirement Document (URD). The functions, inputs and outputs for the three cognitive capabilities given in Appendix A provide a sound basis to derive the System Requirement Document (SRD) for an AUAS and give the supplier a potential basis for sub-system design:

- The Awareness Capability takes the data from external sources and supplies it in a useful form to the Understanding Capability.
- Command decisions are made by the Deliberation Capability, supported by information from the Understanding Capability.

- Each decision will be based on the differences between the current situation and that expected to be encountered. Clearly there can be a range of predicted situations to guide the choice of decisions, but there will always be a limited number of possible decisions.

In engineering terms this is: identifying, with measured confidence levels, the differences between expected and actual scene; followed by a prediction of the result of each decision; then assessing whether any of the results meet the pre-determined ROE criteria.

The result of quantifying these steps as a requirement set can then act as a baseline for assessing the need for new technologies or new exploitation paths for existing ones.

Specific Technology Needs

The requirements presented in this paper have implications for the technologies to support decision-making processes. These are discussed below:

Surveillance Sensors and Sensor Fusion

Current sensor research is based on improving the performance of existing sensors and optimization of their use through sensor fusion. The information is presented to a human in order to add to their knowledge of the current scene and assist their decisions. The work reported in this paper shows that an autonomous system requires scene comparison methods combined with difference detectors and classifiers as the primary approach. This represents a change in the exploitation route for current sensor and fusion development programs.

Scene-matching techniques should provide the basis for Battle Damage Assessment (BDA) using the information available at weapon-release time and immediately after impact.

Automatic and assisted target indication, recognition and identification techniques are being developed as targeting aids, but are still regarded as difficult technologies and often require large databases. With the approach given above, classifiers of a wide range of objects with probabilities of correct classification may be a more robust approach. These lend themselves to AI techniques.

Collateral Damage Prediction

The platform will be carrying weapons of known collateral damage radius and likely effect, so pre-mission data can be provided of blast-range etc. The technical problem is then one of matching the scene provided by the sensor and weapon data to provide sufficient information for collateral damage assessment.

Processors

It is axiomatic that there will always be demands for improved processing power. The particular needs arising from the legal considerations are discussed here. Many are common to manned systems as well. None are new, but will represent significant design challenges for specific systems.

Scene comparison techniques will only be possible if the AUAS platform carries a large amount of pre-mission data with good configuration control on updates. They will also need significant processing power to carry out the coordinate transforms between the data sets for meaningful comparisons to be made.

Recording the precise information used for decision-making will bring issues of data storage and transfer to external users with trade-offs to be made.

These processing and storage requirements will give a significant load on the platform power supply and cooling systems. The well-known rapid evolution of hardware and software will necessitate robust upgrade policies.

Security levels of the on-board data will be a key issue, with a need to corrupt or destroy the information in the event of loss of the platform.

Networks

Spectrum availability both for UAV control and data dissemination will always be a severe constraint. Functionally, there is a clear distinction between the two. Of necessity, control is the most important as it forms part of a mission-critical, if not safety-critical, loop.

The C2 links form the basis of legal responsibility so a robust network monitoring process must be in place. The requirements are to ensure there is continuity of the command chain and that any breaks are known and compensated for in a contingency plan. Potential corruption of updates to mission information must be recognized and accounted for in the decision process.

The scene comparison techniques proposed above may also provide bandwidth-reduction possibilities using existing techniques.

Threats to Own Platform

Electronic Surveillance Measure (ESM) sensors do not provide contextual information. However, under defined conditions, such as lock-on by a Surface to Air Missile (SAM) system, anti-radiation missiles can be launched. As modern threat systems may be able to counter the missiles, it may be necessary to provide safe aimpoints in the event of loss of threat information. These would have to be provided at weapon release from the platform in order to ensure that it is in range after release.

Conclusions

The UK's published interpretation of the LOAC has been used to demonstrate an engineering and procurement framework for legally acceptable systems with a higher level of autonomy than currently in service.

It will be possible to remove humans from many parts of an AUAS control chain with clearly defined limits on the levels of automated decisions and commands. A three-part model of decision-making has been used to derive system requirements. These requirements can only be satisfied with scene-comparison and network-monitoring methods which, in principle, are not new. This will allow increases in UAS autonomy levels using technical advances, but in a legally acceptable way.

Complexity and ambiguities will ensure that there will always need to be human intervention. Criteria have been derived for: the authorization of decisions, their allocation to a human or the machine, and timeliness. The criteria also ensure proper consideration of the critical problems and context.

The principles presented in this paper should be applicable to autonomous systems in other environments.

Acknowledgements

The authors would like to thank colleagues for useful comments during the preparation of this paper, especially Bill Boothby of MOD and John Keirl of Dstl.

References

- Arns, R. 2008. "NATO's Future Joint Air and Space Power (NFJASP)." *The Journal of the Joint Air Power Competence Centre*: Edition 8, 14-17. <http://www.japcc.org>.
- Boothby, William. 2009. *Weapons and the law of armed conflict*, Oxford: Oxford University Press.
- Borenstein, Jason. 2008. "The ethics of autonomous military robots." *Studies in Ethics, Law and Technology*: 2, Issue 1, Article 2.
- Burridge, B. Oct 2005. "Post-modern warfighting with unmanned vehicle systems: Esoteric chimera or essential capability." *Royal United Service Institute Journal* 150: 20-23.
- Department of Defense. 27 June 2007. *Unmanned systems safety guide for DOD acquisition*.
- Finn, Anthony. 2008. "Legal considerations for the weaponisation of unmanned ground vehicles." *Int. J. Intelligent Defence Support Systems* 1, No 1: 43-74.
- JSP383. 2004. *The Joint Service Manual of the Law of Armed Conflict*. Chapters 2, 5 and 12D.

- Lin, Patrick, George Bekey, and Keith Abney. 2008. "Autonomous military robotics: Risks, ethics, and design." Paper prepared for US Department of Navy, Office of Naval Research. Accessed through <http://www.isn.ethz.ch/isn/Digital-Library/Publications/Detail/?lng=en&id=104955>.
- Myers, Andy. Spring 2007. "The legal and moral challenges facing the 21st century air commander." *Royal Air Force Air Power Review*: 10, No 1: 76-96.
- NATO Research Task Group SCI-186. "Architectures for Operations with Manned and Unmanned Aircraft." *In preparation*.
- Quintana, Elizabeth. 2008. "The ethics and legal implications of military unmanned vehicles." *Royal United Services Institute Occasional Paper*.
- Thoms, Joanne. 2009. "Understanding the impact of machine technologies on human team cognition." *4th SEAS DTC Technical Conference*: Paper B7.

Abbreviations

ACC	Air Component Commander
AI	Artificial Intelligence
ALARP	As Low As Reasonably Practical
AUAS	Autonomous Unmanned Air System
BDA	Battle Damage Assessment
C2	Command and Control
CAOC	Combined Air Operation Centre
CV	Carrier Vessel
DCDC	Development, Concepts and Doctrine Centre
DOD	Department of Defense
ESM	Electronic Surveillance Measure
IER	Information Exchange Requirement
ISTAR	Intelligence, Surveillance, Targeting, Acquisition, and Reconnaissance
JFACC	Joint Force Air Command Centre
JSP	Joint Service Publication
LOAC	Law Of Armed Conflict
LCC	Land Component Commander

MC	Mission Commander
MOD	Ministry of Defence
NATO	North Atlantic Treaty Organization
NCW	Network Centric Warfare
NEC	Network Enabled Capability
R&D	Research and Development
ROE	Rule of Engagement
SAM	Surface to Air Missile
SEAS DTC	Systems Engineering and Autonomous Systems Defence Technology Centre
SRD	System Requirements Document
UAS	Unmanned Air System
UAV	Unmanned Air Vehicle
UCAV	Unmanned Combat Air Vehicle
UMS	Unmanned System
URD	User Requirement Document

Appendix A: Requirements for Cognitive Capabilities

A1: Awareness Cognitive Capability

Functions

1. Identify threat level to platform
2. Generate operational capabilities of platform such as endurance, weapon load
3. Generate Information Exchange Requirements (IERs) for current ROEs
4. Monitor status of command chain and current authorized entities above and below this one
5. Configuration control of information from external sources and data generated on the platform
6. Generation of anticipated air or surface scenes using updated pre-mission data
7. Generate actual scene in format for comparison with anticipated scene
8. Detect differences between actual and anticipated scene
9. Derive parameters and confidence levels for differences

Inputs

1. Current ROE status
2. Pre-mission data and updates in flight
3. Current situation from sensors
4. Network status
5. Recognized air picture
6. Recognized ground picture
7. Current platform flight path and intent
8. Current vehicle status

Outputs to Understanding cognitive capability

1. Threat level to platform
2. Location of friendly forces
3. Location of hostile forces
4. Location and flight paths of other platforms in area
5. ROE status as IERs
6. Differences between anticipated and actual scene
7. Quantified parameters for differences between anticipated and actual scenes
8. Current weapon status

A2: Understanding Cognitive Capability

Functions

1. Classification of scene differences into object types
2. Identification of prohibited target areas
3. Geolocate potential targets and probability of correct classification
4. Geolocate non-targets
5. Compare quantified scene-difference parameters and potential targets with ROE IERs
6. Identify missing information, if any, between scene and ROE IERs
7. Propose sources of missing information between scene and ROE IERs
8. Contingency planning for changes to platform status or loss of command link

Inputs

1. These will be the outputs from the Awareness capability
2. Data from on-board data bases; e.g., on-board weapon range and blast damage area
3. Network status

Outputs to Deliberation capability

1. Threat level
2. Objects of military significance with:
 - Probability of correct classification
 - Military significance
 - Missing information to give higher probability of correct classification
 - Proximity to friendly forces
 - Prohibited target areas
3. Civilian objects giving restrictions on attack plans
4. Location of friendly forces
5. Command chain status

A3: Deliberation Cognitive Capability*Functions*

1. Generate self preservation commands if threat level is high
2. Predict effect of use of on-board weapons and capabilities
3. Decide if ROE IERs are met and generate options available to platform

4. Identify options that can be chosen autonomously
5. Quantified method of ranking available options
6. Decide if one or more options meet mission aims
7. Decide on potential sources of missing ROE IERs, including tasking another asset
8. Test that current command chain has given AUAS the authority to act
9. Identify if decision must be referred to another authorized entity
10. Generate list of reasons why a command decision cannot be made in AUAS

Inputs

1. These will be the outputs from the Awareness capability
2. Data from on-board data bases e.g. on-board weapon range and blast damage area

Outputs

1. Commands to effectors if allowed
2. Information for transfer of decision to other authorized entity