



Australian Government
Department of Defence
Defence Science and
Technology Organisation

The Generation of Situational Awareness within Autonomous Systems – A Near to Mid term Study – Analysis

Patrick Chisan Hew

Defence Systems Analysis Division
Information Sciences Laboratory

DSTO-TR-1896

ABSTRACT

This study aims to clarify and capture the nature of electronic situational awareness and its interface with electro/mechanical systems. It argues that “autonomous situation awareness” is about the sufficiency of awareness for autonomy in the situation at hand. The approach is calibrated through historical case studies, and the study then considers the potential from near to mid term technology.

RELEASE LIMITATION

Approved for public release

Published by

*DSTO Information Sciences Laboratory
PO Box 1500
Edinburgh South Australia 5111 Australia*

*Telephone: (08) 8259 5555
Fax: (08) 8259 6567*

*© Commonwealth of Australia 2006
AR-013-702
July 2006*

APPROVED FOR PUBLIC RELEASE

The Generation of Situational Awareness within Autonomous Systems – A Near to Mid term Study – Analysis

Executive Summary

The aim of this document is to clarify and capture the nature of electronic situational awareness and its interface with electro/mechanical systems. This interface is critical in the construction of autonomous systems (robots) with the capacity to perceive their environments, to process that perception and to physically act on the results. This document thus aims to clarify issues, and hence improve the understanding of whether and how autonomous systems could impact upon warfighting.

The core theme of this paper is that *autonomous situation awareness* can be considered through the continuums of:

1. *Autonomy*, measured by the time between references to a human being for input.
2. *Awareness*, measured by the system's usage of information about its environment. Importantly, the system will only use a subset of the total information available.
3. *Situation*, driven primarily by the severity of consequences of the system making decisions and taking action. Severity of consequences is the critical link to the human decision to employ the autonomous system.

The central result is that it is not a question of whether a system "is" autonomous or "has" awareness, but whether the system has *sufficient* awareness to be *sufficiently* autonomous for the situation at hand. Historical review identified cases where "dumb" technologies with low awareness had been deployed at high autonomy, while "smart" technologies with high awareness had been held at low autonomy. Historical experience also suggests strong precedents for concept development, notably from mine warfare, beyond visual range combat and mission command.

In the near to medium term future, fundamental delays in communications (line of sight, speed of light) will force certain systems to higher autonomy. However, if the delay is acceptably low, "human virtual presence" with lower autonomy may be preferred. While it is unlikely that technology will be sufficient if raw data flows are used, the volume and rate of data may be manageable given compression and simulation matched to human perception. Mobile phone networks might be used, given sufficient assurance.

Also in the near to medium term future, smart materials and grid computing may ease the problem of monitoring robot status. However, for movement in the physical environment and interaction with other entities, robot utility will be bounded by the capacity to simplify the environment (by either physical or information means), and by the demand for integration into total battlespace management (deconfliction).

Moreover, for mobility in complex terrain, robot system designers are still seeking workable processes for mapbuilding, with enduring problems that either require heuristic insights or intrinsically parallel computing (DNA or quantum computing).

In strategic terms, given the precedents in concept development, and the known bottlenecks and progress in technology, robotics has *passed* the point of being a new strategic threat, to one that *broadens* the threat at the operational and tactical level. The key feature is *comodification*, enabling different actors to utilise formerly-specialised technology. The threat space from autonomous systems thus builds on advances and comodification of enabling technologies, notably including: insertion into space/orbit, civilian communication networks, and computer hardware and software.

Technology advances aside, this study noted potential assumptions about the human agencies employing robots. The options for robot use are shaped by social background (casualty aversion), expected environment (expeditionary forces) and tempo of decisions (combat intensity). It is to be emphasised that low technology, low awareness robots have already been deployed at high autonomy in the past, and could well be used again.

Authors

Patrick Chisan Hew

Defence Systems Analysis Division

Patrick Hew joined DSTO in 2000, having gained a PhD in Mathematics and Intelligent Information Processing Systems from The University of Western Australia. His first work was in support of HQAST and NORCOM-Coastwatch, in the analysis of maritime surveillance operations and developing concepts for command support tools. In 2002, Patrick relocated to the DSTO team embedded in Australian Defence Headquarters, providing advice on capability development, force structure and the implications of future technology and environments. Since then he has led and conducted work including future technology analysis for AIR7000, thinking in support of the Future Strategic Environment, and the Amphibious Task Group 2005 cross-capability study.

Contents

LIST OF FIGURES	III
LIST OF TABLES	IV
GLOSSARY	V
1. INTRODUCTION	1
2. AUTONOMOUS SITUATIONAL AWARENESS DEFINED	2
2.1 Model and Conceptual Framework	2
2.2 Insights from Case Studies.....	6
3. POTENTIAL FROM FUTURE TECHNOLOGY.....	10
3.1 Analytic Approach	10
3.2 Insights from Bottleneck Issues	11
4. CONCLUSIONS	14
APPENDIX A: BACKGROUND	15
APPENDIX B: METHODOLOGY	16
B.1. Mapping Systems of Interest into the Framework	16
B.2. Using the Framework	18
B.3. Compatibility with Previous Definitions.....	19
B.4. Awareness and Reasoning.....	21
APPENDIX C: CASE STUDIES.....	23
C.1. Engagement	23
C.2. Target Criteria	31
C.3. Adversity and Choice	36
C.4. Communications and Programmability.....	45
C.5. Controlling Consequences	50
APPENDIX D: BOTTLENECK ISSUES	52

D.1. Communications and Programmability.....	52
D.2. Locating Robot in Environment	57
D.3. Status of Robot.....	59
D.4. Locating Other Entities in Environment.....	60
D.5. Target Modelling.....	61

List of Figures

Figure 1: Analytic Model – Human and Robot.	2
Figure 2: Autonomy – Situation – Awareness Space.	4
Figure 3: Hardware Autonomy and Software Autonomy.	8
Figure 4: Technological Autonomy and Operational Autonomy.	8
Figure 5: Autonomy – Situation – Awareness Space and Technological Sufficiency.	10
Figure 6: Study of Autonomous Situation Awareness.	16
Figure 7: Definition of an Autonomous System.	16
Figure 8: Definition of Autonomy, Awareness and Situation.	17
Figure 9: “Push-Pull” Land Mine (Conceptual).	22
Figure 10: Case Studies – A Journey of Engagement.	23
Figure 11: Cut-away of the M14 Anti-personnel Land Mine [7].	24
Figure 12: Notional Automated Counter-Sniping System.	26
Figure 13: Phalanx Close-In Weapon System [18] [19].	27
Figure 14: ZSU-23 Shilka self-propelled air defence gun system [22].	27
Figure 15: Sentry guns in the movie Aliens [23].	28
Figure 16: Aegis air warfare system.	30
Figure 17: Case Studies – A Journey of Target Criteria.	31
Figure 18: Strike missiles – offensive robots at infinite autonomy.	32
Figure 19: V-1 Flying Bomb launch ramp, oriented towards target [32].	33
Figure 20: Early model Tomahawk Land Attack Missile guidance [36].	34
Figure 21: Case Studies – A Journey in Adversity and Choice.	36
Figure 22: Gas system automatic fire mechanism.	37
Figure 23: Hubble Space Telescope.	38
Figure 24: Mars Rover (Artists Impression) [53].	40
Figure 25: DARPA Grand Challenge 2004.	42
Figure 26: Final Data from the DARPA Grand Challenge 2004 [58].	43
Figure 27: Case Studies – A Journey in Communication and Programmability.	45
Figure 28: Recovering a Firebee RPV.	46
Figure 29: Global Hawk UAV [63].	47
Figure 30: Predator UAV and Remote Pilot Station.	48
Figure 31: Concept demonstrator for the Joint Strike Fighter cockpit.	48

List of Tables

Table 1: Spectrum of Autonomy - Examples.	5
Table 2: Spectrum of Awareness - Examples.	5
Table 3: Technological Sufficiency for Situation against Autonomy and Awareness.	6
Table 4: Touchstone Case Studies: Land Mine and Aegis.	18
Table 5: Situation Awareness: Endsley and A-S-A Definitions compared.	20

Glossary

CIWS	Close In Weapon System.
DARPA	Defense Advanced Research Projects Agency.
EM	Electromagnetic (radiation).
GPS	Global Positioning System.
IFF	Identification Friend or Foe.
ISTAR	Intelligence, Surveillance, Target Acquisition & Reconnaissance.
JASSM	Joint Air-to-Surface Standoff Munition.
JSF	Joint Strike Fighter.
LORAN	Long Range Aid to Navigation.
NASA	National Aeronautics & Space Administration.
RPV	Remotely Piloted Vehicle.
RSAF	Royal Saudi Air Force.
SAM	Surface to Air Missile.
TDRSS	Tracking & Data Relay Satellite System.
UAV	Uninhabited Aerial Vehicle.
UV	Ultraviolet (radiation).

1. Introduction

‘All you would need to do is tell him to say “What”, “I don't understand”, and “Where's the tea?”. Who would know the difference?’

– Zaphod Beeblebrox on building a robotic replacement to Arthur Dent's brain
[Douglas Adams, *The Hitchhikers Guide to the Galaxy*].

The quest to construct machines that can act autonomously in an “intelligent”, “human-like” manner has, in turn, raised questions as to what constitutes “intelligence”, “autonomy” or similar qualities. While often interesting and certainly philosophical, the ambiguity of the area can discourage pragmatic discussions about the potential of future technology, especially to warfighting. This is a situation that invites clarification in thinking.

Aim

The aim of this document is to clarify and capture the nature of electronic situational awareness and its interface with electro/mechanical systems. This interface is critical in the construction of systems with the capacity to perceive their environments, to process that perception and to physically act on the results.

Level

This document is aimed at DSTO staff studying the implications of future technology.

Scope

The core theme of this paper is that “autonomous situation awareness” can be considered through the continuums of:

1. *Autonomy*, measured by the time between references to a human being for input.
2. *Awareness*, measured by the system's usage of information about its environment. Importantly, the system will only use a subset of the total information available.
3. *Situation*, driven primarily by the severity of consequences of the system making decisions and taking action. Severity of consequences is the critical link to the human decision to employ the autonomous system.

The basic result is that it is not a question of whether a system “is” autonomous or “has” situation awareness, but whether the system has *sufficient* situation awareness to be *sufficiently* autonomous for the application at hand.

The *Autonomy-Situation-Awareness* conceptual framework is used to study systems from historical and contemporary use. The discussion then considers the implications within the near to mid term, based on known and potential developments in technology.

2. Autonomous Situational Awareness Defined

This section introduces a scheme for studying autonomous situational awareness, summarising Appendix B. It then presents insights from historical and contemporary use.

2.1 Model and Conceptual Framework

This paper asserts that *autonomous systems* may be modelled in the manner illustrated by Figure 1. The salient features are a *human* who *deploys* a *robot* to some purpose, with the robot potentially *referring* back to the human for input and assistance.




Figure 1: Analytic Model – Human and Robot.

Notable aspects of this model:

- Discussions of “autonomy” can be confused by terms used in other contexts, including “human in the loop”, “command guidance” or “launch and leave” / “fire and forget”. It is a core theme of this paper that all such systems can be plotted on a continuum, as ordered by the tightness of the *deploy-refer* loop.
- This paper uses the term *robot* to mean all those portions of the system *excluding* the human. This term will avoid confusion as to what constitutes the “autonomous system” in any particular context, especially given the theme of systems lying on a continuum. It should be noted that the robot need *not* be a single physical entity, but may in fact be the virtual assemblage of both hardware and software components.
- The stated aim of this paper is to “clarify and capture the nature of electronic situational awareness and its interface with electro/mechanical systems”. The above model will, however, help show that “electronic” is only a means of information processing; the crux is in the *utilisation* of information about the robot’s environment. Similarly, “electro/mechanical” is only a means for inhabiting a physical environment; the crux is in the *consequences* of the robot’s actions in and on its environment.

The phrase “autonomous situational awareness” is actually rather fortuitous, for the three words provide an adequate framework for its assessment. The enabling constructs are:

1. *Autonomy*, defined to be the robot’s independence from the human once deployed, and measured by the time between references to a human being for input or assistance. *Autonomy* thus ranges from zero for robots run on “remote control” up to infinity for “fire and forget”, with intermediate values running the continuum of “human in the loop” systems.
2. *Awareness*, defined to be the robot’s appreciation of its environment, and measured by the robot’s usage of information from that environment. Importantly, the robot will only use a subset of the total information available. *Awareness* thus ranges from zero for robots that gather no information, through kilobytes, megabytes and gigabytes for robots using audio or vision, up to infinity for robots with the “human” capacity to ask questions.
3. *Situation*, defined to be complexity of the environment in which the robot acts, and notionally measured by the severity of consequences of the robot’s actions. As the robot’s actions can and often will have effects beyond the immediate (physical) environment of the robot, the severity of consequences is the critical link to the human decision to employ a robot.

The continuums of *autonomy*, *situation* and *awareness* form up as dimensions for the 3D space in Figure 2, and Table 1 and Table 2 draw illustrative examples from Appendix C. It is notable that two of these dimensions (*autonomy* and *awareness*) are readily quantified in technological terms^{1 2} (time and information use), while the third (*situation*) is harder to quantify³. This observation is central to understanding how and why a technological solution can be sufficient in one context but insufficient in another. It is this understanding that, in turn, will clarify the true difficulties and opportunities in generating situation awareness in autonomous systems⁴, and applying them to real-world tasks, warfighting included.

¹ One criticism of this measure for autonomy is that it favours slow systems over fast ones. Consider a robot *A* that executes a task and comes back to the human. Then take robot *B* which is identical to *A* except for dawdling an additional 2 min. Robot *B* would be measured to have 2 min greater autonomy than *A*, even though *A* is functionally superior (With all factors equal, *A* would achieve more than *B*).

² One criticism of the measure for awareness is that it is a measure of data *inputs*, not of cognitive *outcomes*; for instance, knowing the position of a friendly and hostile tank is less “aware” than knowing that one is within firing range of another [1]. This, however, is still an open problem within human command and control research; see [2] for example and Appendix B for further discussion. Nonetheless, for machine systems the measurement of inputs is at least an upper bound on awareness; moreover, it is possible to inspect how this data is transformed within the robot.

³ Harder but not impossible, with potential measures including dollar cost, Gallop Poll popularity ratings and “pucker factor”. The point is that universal measures are unsatisfactory, with measures necessarily sensitive to their context.

⁴ It is acknowledged and emphasised that the *Autonomy-Situation-Awareness* framework introduces new definitions for “autonomy” and “situation awareness”, under a particular model of autonomous systems. This, however, was necessary and desirable to study the issues at hand. Moreover, given that there are any number of definitions of “autonomy” [3] [4] and “situation awareness” [1] [5], the provision of definitions appropriate to the scope and context of this paper is deemed acceptable. See Appendix B for further discussion, notably on compatibility with Endsley in particular.





Figure 2: Autonomy – Situation – Awareness Space.

Table 1: Spectrum of Autonomy – Examples.

“Find-Replace” in a document editor.	< 5 seconds.	The user specifies (<i>deploy</i>) some text to Find and Replace, and the robot verifies with the user at each instance (<i>refer</i>).
Mars Exploration rovers.	> 20 minutes (average).	Communication from Earth to Mars has an unavoidable delay, limited by the speed of light. The robot must be autonomous within this communications loop.
Hubble Space Telescope	> 1 hour.	The particular case of entering <i>safemode</i> . Safemode is a programmed response to an unknown (and potentially dangerous) condition. The robot reverts to a known “safe” mode, and calls for human help. Early in its career, Hubble suffered multiple entries into safemode.
“Replace All” in a document editor.	∞	Once the user has specified the text to Find and Replace, the robot can proceed <i>without</i> referencing back.
Land Mine	∞	Once a land mine is deployed, it will explode if triggered, without referencing back to the human who placed it.

Table 2: Spectrum of Awareness – Examples.

Land Mine	1 bit (On/Off)	Physical pressure pad, in “not pressed” or “pressed” state.
V-1 Flying Bomb	Bytes	Azimuth preset by launch ramp, held constant by weighted pendulum and gyro-magnetic compass. Range embedded in a propeller-like anemometer completing a set number of rotations.
GPS Tomahawk missile.	Kilobytes	Sequences of waypoints for navigation.
Aegis Weapon System	Megabytes	Radar tracks, assembled across a naval task force.
Joint Air-to-Surface Standoff Munition	Gigabytes	Imaging infrared seeker, pattern matching on targets.

2.2 Insights from Case Studies

The historical and contemporary case studies of Appendix C yield the following insights:

1. *Technology can improve autonomy and awareness, but the sufficiency of technology is driven by the situation.* (Table 3)

That case studies of the pressure-activated land mine and the Aegis air warfare system show that the “generation of autonomous situation awareness” is a question that goes beyond technology. In particular, the land mine shows that infinite autonomy has already been achieved, and with the bare minimum of awareness. Equally, the comparison of awareness between the systems shows that improved utility does *not* follow from improved awareness, but is driven by situation factors expanding to the social and political.

More generally, it may be argued that the improvements to technology for autonomy and awareness over the 20th Century have lagged behind the rise in complexity, with consequences, from situations over this epoch. This, however, may only hold for conventional Western militaries, and severity of consequences may be controlled by means other than technology

Table 3: *Technological Sufficiency for Situation against Autonomy and Awareness.*

System	Autonomy	Awareness	Situation	Tech Sufficient?
Land Mine.	∞ (When armed by soldier.)	1 bit. (Pressure pad.)	1917 World War I.	Yes. (Combat use.)
			1997 Conventional military forces.	No. (Signatories to the Ottawa Convention.)
Aegis.	∞ (When set to “Auto Special” by ship captain.)	Megabytes. (Radar tracks.)	1984 Blue-water combat, mass Backfire raids.	Yes. (Deployed but not combat tested.)
			2004 Littoral combat, peacekeeping.	No. (Vincennes incident.)

The distinction of technology from its utility also accounts for the capacity to take a robot system from one application and migrate the technology to another. For awareness, the physics of sensors and effectors (mechanical, electromagnetic, optical, acoustic, chemical) apply universally across physical space (aerospace, maritime, land/littoral). For autonomy, the logic of data processing sits independently in information space. The resulting utility sits, however, in the social space of effects and consequences – aspects of situation.

The overall result, therefore, is that any assessment of technological sufficiency of a robot needs to be coupled to assessments of the complexity of the environment and the severity of consequences. This would entail a social analysis of the consequences of

robot failure(s), with corresponding acceptability within its socio-political political environment.

2. *Hardware autonomy should be distinguished from software autonomy.* (Figure 3)

For a given situation, it may be easier to build “brute force” hardware than to generate “smart thinking” software. Trade-offs will thus drive the demands on awareness, and hence the sufficiency of technology.

Hardware and software also yield different technology goals. Increased hardware autonomy is unequivocally desirable and almost implicit, in replacing human physical presence with hardware, and reducing the need for that hardware to be supported. In contrast, it is by no means clear that increased software autonomy will be desired by the humans employing robots, nor is it clear that the robot's effectiveness will improve. Communications technology enables trade-offs between robot “smart thinking” (autonomy) solutions and human “virtual presence” (communications) solutions.

3. *Technological autonomy should be distinguished from operational autonomy.* (Figure 4)

Hardware autonomy is finite, driven largely by endurance of hardware. In principle, however, software autonomy is infinite, in that software can be written to “do nothing” indefinitely, as a valid robot action. Hence, it is straightforward to build a robot to “do nothing” for arbitrarily long periods, achieving infinite autonomy in a technologically trivial but operationally useless manner. Generalising, the capacity for *technological autonomy* is likely to exceed the desired *operational autonomy*, in that a robot “doing nothing” needs to be “redeployed” to do something else.

The word “redeploy” is put in quotes, in that the hardware need not be moved but the software can be reprogrammed. The capacity to reprogram the robot's software once deployed reduces the decision-making load on the robot itself, and hence potentially reduces the technology necessary to support its awareness.

4. *A given robot may have functions operating at different levels of autonomy.*

There has been an unequivocal push to have the robot and support systems handle low-level issues, freeing up a human to think about higher-level issues. It should be noted, however, that a low-level issue for a human may not transform into a technologically simple issue.




Figure 3: Hardware Autonomy and Software Autonomy.




Figure 4: Technological Autonomy and Operational Autonomy.

5. *Concept development for the use of autonomous systems can draw on strong precedents in mine warfare, beyond-visual-range combat and mission command.*

Any system that links sensor to shooter and releases lethal effect without reference to a human is equivalent, technologically, to the simple, pressure-activated mine. To the objection that “mines are indiscriminate” against more advanced systems, the argument is that “indiscriminate” is just a point on the continuum of awareness. Given that the use of mines is, for certain applications, governed and restricted by international law and agreement, the precedents for use may be substantial. Equally too, the historical observation that mines can and have been disarmed, lifted and turned against their original owners has implications for robot tactics and doctrine development. On the positive side, concept development can draw upon the known strengths of mine warfare: the capacity to engage at tempos faster than human reaction, and thus the deterrence/denial of battlespace volumes.

For human beings, beyond-visual-range combat is qualitatively different to within-visual range combat, for it requires machine sensor data to be assimilated by human sensors to yield cognitive awareness. For robots, however, no such distinction exists – it is all just data from sensors. The historical issues that have restricted the use of long-range weapons could thus impact on concepts for using robots. Symptomatic of this is the potential demand upon ISTAR systems.

The question of whether an autonomous system can be used in a given application is at least as difficult as that faced by a commander in delegating – providing autonomy – to a subordinate. The body of knowledge surrounding mission command is thus of potential utility, with the concept of “reach back” (“interfere forward”) mapping directly to the model of autonomy used in this paper.

6. *Controlling the consequences of robot actions is an equally valid dimension for robot systems development, complementary to technology advances.*

Rather than seeking to improve the capacity for a robot to discriminate, (human) robot system designers/users can seek to control the consequences of robot actions, deliberate or otherwise.

A notable particular application is in arms control conventions, notably the Hague Convention VII on automatic submarine contact mines, and the Ottawa Convention on anti-personnel mines. The Hague Convention VIII sought to control the consequences of robot failure, for instance, in requiring automatic safing if becoming untethered. The Ottawa Convention, by contrast, makes its ban in the dimension of autonomy, in banning systems that trigger on human presence. In both cases, the intent of the Conventions was to minimise the collateral effect on non-combatants, particularly from former conflict zones being littered with explosive debris. In targeting this consequence directly, the Hague Convention VIII will arguably be more resilient and lasting than the Ottawa Convention.

3. Potential from Future Technology

This section considers potential future developments in autonomous situation awareness. It argues that this can be considered through *bottleneck issues*, marking step changes in technological sufficiency. Insights are then presented from a number of such issues.

3.1 Analytic Approach

The previous section demonstrated that, for the study of implications, it is *not* enough to look at autonomy or awareness in isolation; the focus is the *sufficiency* of technology against a situation. Put another way, and as illustrated by Figure 5, technology can generate an autonomy-awareness capacity that may, or may not, cross a threshold for being sufficient to the situation.




Figure 5: Autonomy – Situation – Awareness Space and Technological Sufficiency.

It is tempting to approach the analysis by seeking the boundary between sufficient and insufficient technology. This, however, is unlikely to succeed: while measures for situation can be proposed, it is difficult to argue that they are universally applicable, and any boundary on sufficient/insufficient technology is likely to be quite hazy, drawn from factors that are tangible yet not easily articulated. What *can* be done, however, is to think about journeys from (notionally) clear technological insufficiency to (notionally) clear technological sufficiency, and the questions that may have to be addressed along the way⁵. Addressing this *bottleneck issue* generates a step change in technological sufficiency⁶.

⁵ A notional journey only - the zones of clarity are as difficult to define as the intermediating boundary.

⁶ A step change in sufficiency need *not* imply a step change in the underlying technology; it could arise from incremental growth and refinement.

3.2 Insights from Bottleneck Issues

The exploration of near to mid term future technology in Appendix C identified the bottleneck issues listed below, and makes the following predictions:

1. *Communications and Programmability.*

The bottleneck issue centres on the capacity to communicate with the robot once deployed, and thus change its behaviour. This issue weighs into the balance of engineering for robot autonomy versus control by a human.

In the absence of new physics, robot system developers will continue to rely on known communication media (electromagnetic radiation, electricity, acoustic radiation). They will take on the continuing advances in microelectronics and optics, notably in analog-digital conversion and phased array technology, but will otherwise run up against fundamental physical limits (speed-of-light, horizon line-of-sight, blocking).

The chronic shortage in satellite time will be a problem, and will not be solved without substantially cheaper access to space and orbit. Partial solutions are possible through relay aircraft or by leveraging mobile phone networks, though they require an understanding of how the robots are used within the wider force. On point, suitable relay aircraft are valuable assets and may be on other missions; the use of mobile phone networks presupposes assumptions about where the force will be operating.

On using civilian networks in general, it is noted that unconventional (terrorist) forces have a *social* advantage over conventional forces, in being able to parasitically benefit from the confidence built into a civilian network. In conventional warfighting, it may be possible to suppress these networks. However, for peacekeeping this could be counter-productive, since civilian communications could be an important part of (re)establishing normal conditions.

Returning to physics, round-trip delays will limit the feasible range for effective remote-control (zero autonomy) robots. Moreover, the data flows required for “human virtual presence” will not be trivially achieved. The combination of compression and simulation technologies may, however, provide human operators a sufficient feeling of presence; if not, robot system designers will have to provide autonomy to the robot to make up for the gaps in human input.

2. *Locating Robot in Environment.*

The bottleneck issue centres on providing the robot the capacity to locate itself within the physical environment, in both absolute and relative terms. . Given that a robot is being built to inspect, move within or otherwise interact with the physical environment, the physical location of the robot may have both direct and collateral effects. This issue thus weighs into the severity of the situation against the awareness that can be achieved, and hence the autonomy that may be acceptable.

Where possible, robot system designers will continue to simplify the demands on awareness. This can be achieved physically (structuring the environment) or by information means (supplying navigation cues readily detected by machine sensors, pre-mapping the environment). These measures, however, cannot be universally applied, urban terrain being a key example.

Present-day technology is sufficient to gather appropriate depth map data. The challenge lies in the accumulation of depth data as the robot moves. The problem is of combinatorial complexity, and thus will not be solved by ongoing advances in conventional computing. The intrinsically parallel computations enabled by DNA or quantum computing may open options, over and above heuristic insights made by robot system designers.

3. *Status of Robot.*

The bottleneck issue centres on providing the robot the capacity to monitor its physical status. Environmental drivers may compromise the robot or its subsystems, with follow through to mission capability, potentially requiring human intervention. The awareness that the robot has of its status, and its capacity to deal with it within the context of situation, will drive the autonomy that can be achieved.

In the near to medium term future, this issue will be eased by ongoing progress in microelectronics/optics and electronic computing. Significant improvements could also arise from smart materials and grid computing.

4. *Locating Other Entities in Environment.*

The bottleneck issue centres on providing the robot the capacity to locate other entities in the environment. The interaction the robot has with such entities will have direct and collateral effects. This issue thus weighs into the severity of the situation against the awareness that can be achieved, and hence the autonomy that may be acceptable.

All of the issues that hold for physically locating the robot within its environment carry forward, compounded by the dynamic element. Where possible, robot system designers will tag entities by means readily detected by machine (electronic) means, however, the battlespace will include neutrals that cannot be readily tagged and hostiles actively avoiding detection. The scope of the problem thus encompasses the space of sensors and the exploitation of sensor data.

In the near to medium term future, progress will come with the general push to improve battlespace management (notably airspace and friendly fire deconfliction). Conversely, robots will need to be able to report their location and intended movements, at tempos demanded by battlespace management systems. This may compromise robot freedom of movement, and will add to communication loads.

5. *Target Modelling.*

The bottleneck issue centres on providing the robot the capacity to model targets, an issue that is particularly acute when the robot's action generates a lethal effect. The impact is on the severity of the situation against the awareness that can be achieved, and hence the autonomy that may be acceptable.

To date, technology has enabled approaches based on "Engage Unless Friendly" and "Compare with Supplied Model". For conventional militaries, drivers from the political environment, combined with the potential from communications technology, may orient robot system designers to the latter approach.

A third approach, "Training by Similarity", will continue to be an ongoing goal of artificial intelligence research. This research involves the search for workable processes, and has not progressed to the implementation and refinement of processes

that are known to work, so no predictions can be made as to when success may occur. However, the intrinsically parallel computations enabled by DNA or quantum computing may open options, over and above ongoing improvements in conventional computing hardware and scientific software.

4. Conclusions

This document set out to clarify and capture the nature of electronic situational awareness and its interface with electro/mechanical systems, and to explore the potential of technology in the near- to medium- term future. From a model of autonomous systems that centred on deployment by, and input from, a human being, *autonomous situation awareness* was captured in linked but distinct continuums of:

1. *Autonomy*, measured by the time between references to a human being for input.
2. *Awareness*, measured by the system's usage of information about its environment.
3. *Situation*, driven primarily by the severity of consequences of the system making decisions and taking action.

The central result is that it is not a question of whether a system "is" autonomous or "has" awareness, but whether the system has *sufficient* awareness to be *sufficiently* autonomous for the situation at hand. Historical review identified cases where "dumb" technologies with low awareness had been deployed at high autonomy, while "smart" technologies with high awareness had been held at low autonomy. Historical experience also suggests strong precedents for concept development, notably from mine warfare, beyond visual range combat and mission command.

In the near to medium term future, fundamental delays in communications (line of sight, speed of light) will force certain systems to higher autonomy. However, if the delay is acceptably low, "human virtual presence" with lower autonomy may be preferred. While it is unlikely that technology will be sufficient if raw data flows are used, the volume and rate of data may be manageable given compression and simulation matched to human perception. Mobile phone networks might be used, given sufficient assurance.

Also in the near to medium term future, smart materials and grid computing may ease the problem of monitoring robot status. However, for movement in the physical environment and interaction with other entities, robot utility will be bounded by the capacity to simplify the environment (by either physical or information means), and by the demand for integration into total battlespace management (deconfliction). Moreover, for mobility in complex terrain, robot system designers are still seeking workable processes for mapbuilding, with enduring problems that either require heuristic insights or intrinsically parallel computing (DNA or quantum computing).

In strategic terms, given the precedents in concept development, and the known bottlenecks and progress in technology, robotics has *passed* the point of being a new strategic threat, to one that *broadens* the threat at the operational and tactical level. The key feature is *comodification*, enabling different actors to utilise formerly-specialised technology. The threat space from autonomous systems thus builds on advances and comodification of enabling technologies, notably including: insertion into space/orbit, civilian communication networks, and computer hardware and software.

Technology advances aside, the exploration noted potential assumptions about the human agencies employing robots. The options for robot use are shaped by social background (casualty aversion), expected environment (expeditionary forces) and tempo of decisions (combat intensity). It is to be emphasised that low-technology, low-awareness robots have already been deployed at high autonomy in the past, and could well be used again.

Appendix A: Background

Task STR 02/211 (Implications of Future Technology) aims to provide ADF clients an improved ability to define future capability in response to developing technologies and asymmetric and disruptive effects. Autonomous systems technology certainly has the potential to yield such effects, and with advances in information and electro/mechanical systems technology, monitoring of this area is warranted.

This paper was initiated as a focussed complement to a larger study of autonomy currently being conducted by Mr Mark Ellis and Dr Robert Winter. There was also interest in deepening, as possible, the thinking on technology in context, drawing on ideas floated under Task 02/170 (Value-Centric Warfare). Other closely relevant DSTO work includes the thrusts in Automation of the Battlespace (Dr Anthony Finn) and Robotics (Dr Mark Anderson).

Appendix B: Methodology

Section 2 introduced a framework for the study of robotic systems. This appendix formally presents this framework, and the way that it can be applied to generate insights. The discussion is organised to follow Figure 6.




Figure 6: Study of Autonomous Situation Awareness.

B.1. Mapping Systems of Interest into the Framework

The methodology begins with the following definition for a *autonomous system*

Definition: An *autonomous system* is any system that can be modelled in the manner corresponding to Figure 7, with a *human* who *deploys* a *robot* to some purpose, and the robot potentially *referring* back to the human for input and assistance.




Figure 7: Definition of an Autonomous System.

Points that should be noted in this definition:

- The definition is designed to enable the study of autonomy and situation awareness, under definitions to follow. No claim is made about the consistency of this definition with others that may exist in the field, or fitness for purpose beyond the scope of studying autonomous situation awareness.
- The robot need not be a single physical entity, but may be the virtual assemblage of distributed hardware and software components.

The definitions for *autonomy*, *awareness* and *situation* are then cast as follows:

Definition: *Autonomy* is defined to be the robot's independence from the human once deployed, and measured by the time between references to a human being for input or assistance.

Definition: *Awareness* is defined to be the robot's appreciation of its environment, and measured by the robot's usage of information from that environment.

Definition: *Situation* is defined to be complexity of the environment in which the robot acts, and notionally measured by the severity of consequences of the robot's actions.




Figure 8: Definition of Autonomy, Awareness and Situation.

Notable points in these definitions:

- There is no requirement for the robot to refer back to the human. The loop from human to robot may thus loosen from tight control to completely open, with *autonomy* increasing from zero to infinity.
- The definition of *awareness* is a measure of data *inputs*, not of cognitive *outcomes*. However, the measurement of cognitive outcomes is an open problem, complicating its study. For machine systems the measurement of inputs is at least an upper bound on awareness; moreover, it is possible to inspect how this data is transformed within the robot.
- The definition for *situation* does not attempt to quantify complexity or severity, only to recognise that it exists.
- Compatibility of the terms *awareness* and *situation* with existing definitions for *situation awareness* will be discussed later.

The procedure for mapping a system of interest into the *autonomy-situation-awareness* framework is thus:

1. Identify the *human*, and hence the *robot*.
2. Identify the *deploy-refer* loop, and hence the *autonomy* of the robot from the human.
3. Identify the information being used by the robot, and hence its *awareness*.
4. Recognise that the robot's actions in the *situation* will have consequences, and seek to articulate the severity.

B.2. Using the Framework

It is recalled that the aim of this paper was to “clarify and capture the nature of electronic situational awareness and its interface with electro/mechanical systems”. The underlying goal was a framework for considering technological sufficiency, for the study of the implications of future technology. It is the author's contention that any framework for doing so *must* be able to handle the touchstone cases of a land mine compared to the Aegis air warfare system. Under the *Autonomy-Situation-Awareness* framework, it is observed that (Table 4):

- Both systems can be released to infinite autonomy; that is, to act without referring back to a human being.
- The low-technology land mine was fully sufficient for the situation of 1917, of open warfare in open terrain. In comparison, the high-technology Aegis is not regarded as sufficient for 2004, for “Auto Special” operations in littoral airspace crowded with civilians.
- Equally too, (anti-personnel) land mine technology is unacceptable post-1997, for conventional military forces of nations signatory to the Ottawa Convention. Other actors, however, may not feel so restricted.

Table 4: Touchstone Case Studies: Land Mine and Aegis.

System	Autonomy	Awareness	Situation	Tech Sufficient?
Land Mine.	∞ (When armed by soldier.)	1 bit. (Pressure pad).	1917 World War I.	Yes. (Combat use.)
			1997 Conventional military forces.	No. (Signatories to the Ottawa Convention.)
Aegis.	∞ (When set to “Auto Special” by ship captain.)	Megabytes. (Radar tracks.)	1984 Blue-water combat, mass Backfire raids.	Yes. (Deployed but not combat tested.)
			2004 Littoral combat, peacekeeping.	No. (Vincennes incident.)

These case studies yield the following key result:

Observation: While autonomy and awareness can be assessed from technology, assessment of sufficiency requires an understanding of the situation.

That is, it was straightforward to measure the autonomy and (to an extent) the awareness of the land mine and of Aegis. However, to assess the sufficiency of the technology, it was necessary to look at the situations where they were used. In the case studies, the framework was applied historically, but it equally applies to futures analysis.

The above observation was made in using the definitions of the *Autonomy-Situation-Awareness* framework. A different framework might characterise “autonomous situation awareness” in another manner. However, the given the experience from the land mine versus Aegis case, the following prediction is made:

Prediction: The assessment of the utility/sufficiency of (future) robotics technology is impossible without a (future) environment.

This prediction cannot be tested or otherwise proved, since it is impossible to characterise all possible frameworks for capturing technology, even when restricted to “autonomous situation awareness”. However, any particular framework can be tested, by using the land mine and Aegis cases.

The question remains as to whether and how the *Autonomy-Situation-Awareness* framework can be used to consider a set of (future) technology, given a (future) environment. The following threads are suggested, and were used in the historical case studies to follow:

- Given that a robot has the technological potential to operate at given level of autonomy, what factors might cause it to be used at a lower level?
- What is the information supporting robot awareness, and how is it obtained?
- Given the demands of the environment, how does the capacity for the human to intervene raise or lower demands on the robot?
- What is the nature of the deploy-refer loop between human and robot, and how is this supported by technology?

B.3. Compatibility with Previous Definitions

The *Autonomy-Situation-Awareness* framework introduces new definitions for *autonomy*, *awareness* and *situation*. Given the preexisting body of knowledge for *situation awareness*, and the Endsley definition in particular, the question of compatibility of the new definitions with the old invites examination.

Endsley defines *situation awareness* thus [6]

Definition (Endsley): *Situation Awareness* is the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future.

The attraction (and longevity) of this definition derives from its utility in the cognitive sciences. In particular, the notions of perception, comprehension and projection well describe the cognitive processes of military operators at the conceptual level, and map on a one-to-one basis to the data fusion concepts of object refinement (situation picture

compilation), situation refinement (resulting in a situation assessment) and threat/impact refinement (resulting in a threat/impact assessment)⁷.

It is argued that the Endsley and *Autonomy-Situation-Awareness* definitions for *situation awareness* are congruent when (for A-S-A) the words *situation* and *awareness* are taken together, as illustrated by Table 5.

Table 5: Situation Awareness: Endsley and A-S-A Definitions compared.

Endsley		Autonomy-Situation-Awareness
Situation	the comprehension of their meaning	the complexity of the environment in which the robot acts, and notionally measured by the severity of consequences of the robot's actions
Awareness	the perception of the elements in the environment within a volume of time and space ... and the projection of their status in the near future	the robot's appreciation of its environment, and measured by the robot's usage of information from that environment

The basic argument of compatibility is that the A-S-A definitions:

- Uses the word “environment” to encompass what Endsley’s “elements in the environment”, “volume of time and space” and “projection of their status”.
- “complexity of environment” and “severity of consequences” is the necessary support to Endsley’s “comprehension of their meaning”.

The Endsley definition for *situation awareness* can, in fact, be fully applied as a definition of *awareness*, as a “checklist” in terms of inspecting a given robot; the A-S-A definition just generalises the aspects to *awareness* for capture and makes quantification explicit. Crucially though, the A-S-A definition draws out the robot as being something within an environment, with *awareness* being “internal” and *situation* “external”. This is important for future technology studies, when advances in technology can change *both* the “internal” and “external” spaces. This can be seen in studying the landmine, in that both Endsley and A-S-A are congruent for the landmine’s internal space (one binary bit), but the A-S-A definition draws out a crucial parameter on whether the landmine is an acceptable weapon system.

Overall, the Endsley and A-S-A definitions of *situation awareness* are compatible and congruent. However the A-S-A definitions have additional utility here, for the study of (future) robots in (future) environments.

⁷ Commentary from Dr Martin Oxenham (DSTO), whom was one of the review readers to an earlier draft of this document. The author is grateful for this input.

B.4. Awareness and Reasoning

In introducing the definition of *awareness* used in this document, it was noted that usage of information was a measure of data *inputs*, not cognitive *outcomes*. The key question is whether and how “reasoning” is captured within “appreciation of its environment” and “usage of information”.

A stress case⁸ is the following: Consider two robots which have identical sensor and weapons systems that have been designed to fire on detected targets automatically. If Robot 1 fires indiscriminately, while Robot 2 employs additional reasoning to be more selective in its firing, for example firing only on targets which are not friendly, then it should be the case that Robot 2 has more awareness than Robot 1.

Within the proposed definition for *awareness*, this can be accounted for, albeit under a liberal definition of “appreciation of its environment”. Aspects of the above can actually be seen in the land mine. As discussed in Appendix C, a land mine can be modelled as having an awareness of one binary bit, with respect to a pressure threshold – indeed, this is observable as an actual physical (mechanical) switch. However, a more careful review would ask whether and how this threshold pressure counts into the measure of awareness. That is, instead of the land mine having an awareness of one binary bit, it might actually be more accurately modelled as having an awareness of two floating point values (*Sensed* pressure and *Threshold* pressure).

This broadening of information to the robot accounts aligns with the intuitive notion of greater awareness. However, there is a danger of infinite regress, specifically whether the decision logic “if (*Sensed* > *Threshold*) then ...” should *also* count within the awareness measure. This can be resolved by regarding the land mine as being *hard coded*. That is, in a land mine, *Threshold* is fixed at manufacture, so the only information the land mine robot can gather is *Sensed* pressure. In contrast, if a land mine *was* built with a selectable *Threshold*, then this would constitute additional information gathered by the land mine robot. The extension to a digital camera robot – for instance, where *Sensed* pressure is replaced with a digital camera, and *Threshold* pressure is replaced by a reference image – can then be seen. This resolution accounts for the digital camera robot having greater awareness than the pressure sensing land mine robot.

A complementary approach is to look at *internal states*. Consider a “push-pull” land mine – one where either pushing the land mine beyond one threshold, or pulling it beyond another, will cause detonation. The key point is that this could be implemented as one pressure sensor connected to *two* binary switches, one for push and the other for pull, as illustrated by Figure 9. The wording of “appreciation” and “usage” of information from the environment then invite care – both the “push-pull” and conventional land mine have one pressure sensor (“appreciation”?) but the “push-pull” robot records this into two binary switches (“usage”?).

⁸ Provided by Dr Martin Oxenham (DSTO), whom was one of the review readers to an earlier draft of this document. The author is grateful for this input.




Figure 9: “Push-Pull” Land Mine (Conceptual).

For calibrating the definition of *awareness*, the preferable answer is that the “push-pull” land mine has greater awareness than the conventional one, on the intuitive grounds that it has a more sophisticated rule set – “If pushed or pulled” compared with the “if pushed”. This intuition can be captured by aligning the definition of awareness with the number of internal states that the robot can take on. The conventional land mine robot can take on only one of two states, from $Pressure > Threshold$. The “push-pull” land mine robot can take on one of three states: “neutral”, “pushed” or “pulled”. So arguably, though both the conventional and “push pull” robots take in the same sensor data, the “push pull” robot manifests its additional reasoning in terms of additional internal states.

Overall then, the proposed definition of *awareness* invites further examination, but has withstood at least “additional reasoning” stress case. Two complementary yet distinct avenues were examined:

- Broadening the scope of information that could be supplied to a robot, with infinite regress handled through the notion of hard coding.
- Looking at the number of internal states that a robot can take on.

Appendix C: Case Studies

The case study analysis takes groups of systems, identifies the *robot* and its interaction with the *human*, and discusses the surrounding tactical, environmental and strategic issues. In each case, the *robot* is readily plotted on the *autonomy-awareness* technology plane, but it is the issues from the *situation* that enable a technology combination to be used in one context but not another. Indeed, “dumb” technologies may be allowed in one epoch and environment, but “smart” technologies may not be allowed in another – it is all about the *sufficiency* for the application at hand.

C.1. Engagement





Figure 10: Case Studies – A Journey of Engagement.

The case studies in this section aim to show that autonomy is trivially achieved, and that the critical aspect is the generation of sufficient awareness for the situation. The systems under consideration are those that are intended to engage a target of interest, with potentially lethal effect.

The starting point is the pressure-activated land mine [7]. The robot in this instance is the mine, deployed by the human through emplacement and arming. Once armed, the land mine robot has infinite autonomy, in that it requires no further human input, and its awareness is exclusively focussed on the presence or otherwise of a sufficiently heavy object on its pressure activator. On sensing that presence, the land mine will explode.



```

if(Is armed) {
  if(Pressure > Threshold) {
    Explode
  } // if
} // if

```

The land mine, regarded as a robot, has awareness exclusively focussed on the state of the pressure plate. This awareness is adequately captured by the state of the firing pin. This example shows that autonomous situation awareness does not require electronics or other sophisticated information technology.

Figure 11: Cut-away of the M14 Anti-personnel Land Mine [7].

Recalling that the aim of this paper was to “clarify and capture the nature of electronic situational awareness”, the land mine robot has a salutary lesson for technology. The information that the robot is using (“Is there an sufficiently heavy object pressing me?”) is a binary state, and adequately represented by the state of a mechanical switch. Alternatively, an electronic switch could be used, but this does not change the information being used, only its representation. The insight, therefore, is that it is not “electronic situation awareness” that is of interest, but the information being used (awareness) and the context from where it is taken (situation).

A land mine operated in command-detonated mode is essentially the same technology for the robot, but operating at zero autonomy; that is, the robot is continually referring to the human being for input. Other case studies will discuss robots operating at intermediate values of autonomy; the observation here is about autonomy limited by awareness of the situation. To be specific, the land mine robot is fully capable of being autonomous in its direct function – it is just gathering pressure information. However, for the human being, the severity of consequences prevents deployment at higher autonomy⁹.

It is now worth ramping the discussion up from the land mine, by looking at situations where the robot has greater awareness and the situation is potentially simpler. The first systems considered are those used in counter-sniping, such as VIPER (Vectored Infrared Personnel Engagement and Return-fire) [10] [11], SADS (Sniper Acoustic Detection System) [12] or PDCue [13]. The robot in this situation is the virtual assemblage of sensors (electro-optical, acoustic), the robot being emplaced to monitor a zone of interest. The robot does so at infinite autonomy (no human intervention required, though intervention is possible), reporting back to the human if a sniper signature is detected.

Two observations may be made at this point. The first is that the chemoluminescent flash [10] or acoustic report [14] from a sniper rifle yield a distinct and unambiguous target signature amenable to electronic processing. This is true even within complex terrain, where false signals may be present. The second observation is that it is straightforward to expand the robot to include a weapon [15] [16], and have the sense-and-process subsystem cue the weapon subsystem. Automated engagement of a target signature thus falls out as a natural consequence – a robot that is functionally equivalent to a land mine, save that it is somewhat more sophisticated in what it chooses to engage.

It is, however, this subtle distinction in awareness *relative to* situation that determines the acceptability or otherwise of the weapon system. Arguments can and have been mounted against the acceptance of landmines as a legal weapon system [17] the critical aspect being that landmines are indiscriminate, and that the effects of their being discriminate are out of proportion to their military value¹⁰. The thread being followed here is that these are arguments about landmine *awareness within situation*, *not* about landmine *autonomy*, and that to block weapon systems on the basis of autonomy would be (potentially) close off robotics technology in general.

⁹ Australian is a signatory to the Ottawa Convention, banning the use of anti-personnel mines [8] [9].

¹⁰ The article [17] concludes “landmines are incompatible with the fundamental principles of IHL [International Humanitarian Law]. The limited military utility of landmines is far outweighed by the appalling humanitarian consequences of their effects.” The thesis to [17] was built on landmines contradicting the principle of proportionality, that “human suffering caused by a particular weapon must not exceed military necessity”, and that “As weapons landmines are indiscriminate, cause unconscionable harm and excessive injury to civilians”.



```

if(Detector Subsystem is directly connected to Weapon Subsystem) {
  if(Sniper signature is detected) {
    Shoot at location of signature
  } // if
} // if

```


Notional assemblage of a SADS Acoustic Antenna [12] cueing a TRAP T-250 mounting an M82A1 rifle [15]. In principal, the combination of these two systems would generate a robot capable of detecting and automatically shooting back at a sniper. This is functionally equivalent to a land mine, albeit one that is more aware.

Figure 12: *Notional Automated Counter-Sniping System.*

The potential and issues that arise from automated counter-sniping may be seen in operational use of the Phalanx Close-In Weapon System (CIWS). The robot in this situation consists of the “R2D2” housing the search and tracking radars, an electronic computer processing unit, and a 20 mm cannon for engagement [18] [19]. Once released for use, the Phalanx operates at infinite autonomy, engaging anything within its search zone – preferably an incoming missile or other threat, but potentially a friendly helicopter or other false positive. The releasing of a Phalanx for use is thus weighed by the severity of the threat (the need to rapidly and decisively engage ship-killing missiles), with the risks handled by air battlespace management and deconfliction¹¹. The USS Stark incident in 1987 showed that this is not an easy question [21].

It is worth emphasising here the incremental change in technology, on the awareness axis, from land mine through Phalanx to counter-sniping. The awareness is not of the maritime or land environment as such, but of the information signatures that are gathered: mechanical pressure for the land mine, electrooptical or acoustic signatures for counter-sniping, radar returns for the Phalanx. To turn this point around, so long as a signature can be gathered from the environment, it can be used to cue an autonomous system. This point can be seen in practice with the ZSU-23 Shilka air defence system [22] and in fiction with the sentry guns used by the Marines in the movie *Aliens* [23].

¹¹ For comparison, [20] describes an exercise where a helicopter carrying a number of senior commanders was notionally “killed”, a “friction” effect of the overall air warning status of the task force.




```

if(On "Auto Special") {
  if(Radar track in search zone) {
    Shoot at radar track
  } // if
} // if


```

Figure 13: Phalanx Close-In Weapon System [18] [19].



The ZSU-23 uses the same technological principles as the Phalanx CIWS, but in a land anti-aircraft instead of a maritime anti-missile environment.

Figure 14: ZSU-23 Shilka self-propelled air defence gun system [22].



The UA 571-C Remote Sentry Weapon System, deployed to guard a passageway in the movie *Aliens*. In contrast to the Phalanx CIWS, the UA 571-C uses multi-spectral infrared or UV light, rather than radar. However, the same technological principles are involved – the system will engage anything meeting its target criteria, be it swarming aliens or a garbage can hurled into the line of sight for test purposes.

Figure 15: Sentry guns in the movie *Aliens* [23].

The Phalanx system is specifically designed to encapsulate everything within a single physical unit. This, however, is a design choice, and the systems could be linked virtually. Indeed, this concept can be seen in practice and in aspiration in the Aegis air warfare system [24] [25] [26] [27], the robot in this instance being the virtual assemblage of the sensors, controllers and weapons across the maritime task force¹². This robot has awareness of the maritime taskforce battlespace, assembled as track data and associated information¹³. The level of autonomy is selectable, and it is notable that this is a command decision based on threat, environment and consequences. In particular, the captain¹⁴ of an Aegis ship could release the system to operate at infinite autonomy (“automatic special” mode [28]), automatically engaging targets meeting criteria. The actual operational use of this mode is exceptionally context sensitive – acceptable in a bluewater scenario against a multi-regiment Backfire raid, less so in a littoral scenario crowded with neutrals. The issues that arise were tragically brought into relief by the shooting down of an Iranian Airbus in 1990 by the USS Vincennes [29] ¹⁵.

In summary, the conceptual journey from land mine, through automated counter-sniping, to the Phalanx and thus Aegis draws out the following key points:


- It is trivially easy to set a robot to engage targets of interest at infinite autonomy. The real question is the acceptability of engaging false positives, or of not engaging false negatives. Increased robot awareness can improve the human acceptability of deploying the robot, as judged within the context of the situation.
- Robot awareness is a function of what can be collected from the environment, and is otherwise independent of the environment itself. Examples include electromagnetic, acoustic and chemical signatures, all amenable to technological collection. The transferability of signatures across environments enables the potential reuse of robots across environments.
- The networking of systems is a lubricant to the construction of autonomous systems, by enabling the autonomous decision element (robot awareness) to be hooked in as a software/hardware element.

¹² The shipboard Aegis elements of the SPY-1 search radar, the SPG-62 target illuminators and the SM-2 Standard missiles can be directly compared with the search radar, tracking radar and cannons on the Phalanx, recalling that Phalanx deliberately packages its systems together while Aegis elements are distributed around the ship. Moreover, the aspiration for Aegis is that it will network elements on other platforms, notably other ships (mounting sensors) and aircraft (AEW&C in particular).

¹³ A comment on Figure 10 was that *awareness* may be better quantified to the *number of dimensions* of the data being gathered, versus the raw volume of data. The point was that if the measure of awareness was tied to the number of elements detected (Aegis tracks say), then results would be inconsistent (“no elements $\Rightarrow$ no tracks $\Rightarrow$ zero awareness”). However, the contention here is that awareness is actually tied to the battlespace volume being searched, and the maximum capacity to store element details – that is, if a volume is swept and found empty, *this is part of the awareness data*. On Aegis in particular, an upper bound can be stated: In Baseline 4 Aegis, the command/control and weapon control systems each use four AN/UYK-43B computers, with two central processing units each with a memory of 2.5 million words [28].

¹⁴ In consultation with, and possibly “double-hatted” as, the Anti-Air Warfare coordinator of the task force.

¹⁵ It is similarly worth noting an incident from the 1991 Persian Gulf War, which resulted in a US F-15 pilot being awarded the Distinguished Flying Cross. The pilot was under direction from an E-3 Airborne Warning and Control aircraft, was cued to a target and *ordered* to shoot, but declined and closed to visual range to identify. It turned out that his “target” was a RSAF Tornado returning from a strike mission [30, pp496–7].



```

if(On "Auto Special") {
    if(Target track in defence zone, not responding to IFF, ...) {
        Shoot at target track
    } // if
} // if

```

Elements of the Aegis air warfare system: SPY-1 search radar [31], SPG-62 target illuminators [31] and SM-2 missile launch [27]. Compare these elements with those marked in Figure 13 for the Phalanx CIWS.

Figure 16: Aegis air warfare system.

C.2. Target Criteria




Figure 17: Case Studies – A Journey of Target Criteria.

The previous set of case studies established the crucial interplay between awareness and situation. The journey here examines the way that this interplay is raising the demand for ISTAR, and draws out issues in offensive versus defensive applications. The systems discussed here are all robots consisting of aerial vehicles with an explosive payload, deployed “fire and forget” on infinite autonomy.



V-1 Flying Bomb [33]



Harpoon Anti-Ship Missile [35]



Tomahawk Land Attack Missile [36]



Joint Air-to-Surface Standoff Munition [38]

Figure 18: Strike missiles – offensive robots at infinite autonomy.

The starting point is the V-1 Flying Bomb. Target awareness was generated by mechanical means: the launch ramps were oriented towards their targets¹⁶ [32], range information was embedded in a propeller-like anemometer completing a set number of rotations [33] [34, Chapter 3], and target altitude was assumed to be negative [33]. In effect, the V-1 robot had no “actual” awareness of the target other than as an azimuth, range and altitude.

¹⁶ The V-1 maintained its orientation through a weighted pendulum and gyro-magnetic compass [33]. This point draws out the need for precision in discussing *awareness* – strictly, the V-1 was not aware of the absolute azimuth to target, only that it need maintain a zero change in azimuth after launch.



Figure 19: V-1 Flying Bomb launch ramp, oriented towards target [32].

Early models of the Harpoon anti-ship missile [35] may be regarded as a step up in awareness. The awareness is best considered in two parts: mid-course and terminal. For mid-course, the Harpoon is supplied with an azimuth, altitude¹⁷ and a coarse range. When this range is reached¹⁸, the active seeker radar is turned on for terminal guidance, the missile homing on objects with an acceptable radar return. The Harpoon robot's awareness is thus an azimuth, altitude, coarse range and radar signature. In effect, the Harpoon robot may be regarded as a "V-1 carrying a land mine": the "V-1" being the flying element with azimuth, altitude and coarse range, and the "land mine" being triggered by the radar signature.

The earlier models of the Tomahawk Land Attack Missile [36] are in turn, a further step up in awareness, again considered in mid-course and terminal phases. Mid-course guidance was supplied by terrain comparison¹⁹, and terminal guidance by comparison of an electro-optical image of the target with a stored image. The Tomahawk robot's awareness is thus a terrain path to follow, together with an electro-optical image to hit.

¹⁷ The radar altimeter is used to maintain a flight profile, but implicitly embeds a target altitude at sea level.

¹⁸ Text in [35] implies that the radar seeker is turned on at some point after launch. Even, however, if the seeker was active from launch, the point would not change the discussion; it is equivalent to setting the mid-course guidance "coarse range" to zero.

¹⁹ Again note the need for precision in discussing *awareness*. Although the Tomahawk had no awareness of its launch location, in order for terrain comparison to work, the launch locations had to be selected in advance. In this way, after launch, the Tomahawk would detect the terrain that it was supposed to follow. Conversely, if launched from a different site, the Tomahawk would either fail to detect terrain matching its database, or could conceivably follow terrain that coincidentally matched its programming. For illustration, ship-launched Tomahawk missiles launched in the 1991 Gulf War had to be routed over Iran, which had mountainous terrain with sufficient definition for terrain comparison navigation [37].



Early model Tomahawk Land Attack Missiles used Terrain Comparison (TERCOM) for mid-course guidance and Digital Scene Matching Area Correlation (DSMAC) for terminal homing.

Figure 20: Early model Tomahawk Land Attack Missile guidance [36].

That earlier-model Tomahawk missiles used terrain comparison draws out a point of some value, namely the impact of GPS. GPS significantly simplifies the target criteria question, in that it is a straightforward and inexpensive²⁰ means of making a robot aware of its absolute geospatial location. This, in turn, enables target criteria to be expressed directly as a geospatial location, rather than indirectly (through azimuth, range, terrain path ...). In particular, in moving from terrain comparison to GPS navigation, information requirements for mid-course awareness actually *drop*, from hundreds of kilometres of terrain-elevation data to a sequence of geospatial coordinates.

In contrast, the quest for precision continues to boost the information needs for terminal awareness, as seen for instance in the Joint Air-to-Surface Standoff Munition (JASSM) [38] [39] [40]. Whereas the Tomahawk robot was supplied with 2D awareness of the target²¹, the JASSM robot is supplied with 3D awareness²², with a commensurate increase in data.

²⁰ Excluding the astronomical cost of deploying and supporting a GPS satellite constellation.

²¹ Inferred from, but not explicitly stated in, the reference material at [36]; digital scenes imply 2D.

²² Not explicitly declared in reviewed material, however, it is implied in [41] through references to “target wireframe construction” and “wireframe model view manipulation” in the Precision Targeting Module of the JASSM Mission Planning software. Wireframe models are associated with 3D but not 2D target modelling.

From V-1 to Harpoon, Tomahawk and then JASSM, the robot's awareness increases, arguably by some 9 decimal orders of magnitude. There is, however, no increase in the autonomy – each robot is already working at infinite “fire and forget” autonomy²³. The technology push to increase awareness is to satisfy precision targetting within the context of the situation: from total war where hitting the city was sufficient, to high-intensity blue-water conflict where the goal is to hit a ship within an uncluttered vicinity, to major regional land combat where guaranteed elimination of a land target is critical, to casualty-averse littoral combat requiring precise effects on target with limited collateral damage. The target awareness held by an offensive robot is independent from the effect that it delivers; to use an extreme example, the Tomahawk missile has the same awareness irrespective of its warhead – conventional or nuclear.

The other point to draw out is that a 9 decimal orders of magnitude increase in data requirements has clear implications to ISTAR support, at least for those potential robot users that are concerned about achieving precision targetting. Some users will incur necessary and unavoidable costs to achieve this support. Others could judge that lower precision would be acceptable within their context²⁴.

In summary, the conceptual journey from V-1 to Harpoon, Tomahawk and then JASSM draws out the following key points:

- The capacity to create an offensive robot with infinite autonomy is independent of the robot's awareness. The real question is about the acceptability of direct and indirect effects, within the context of the situation.
- Strictly, the robot is not even aware of the target as such, but only of a set of target information criteria to satisfy. The conceptual link from “target” to “information criteria” lies with the human being deploying the robot.
- Increasing the awareness of offensive robots increases the demand for ISTAR.

²³ The capability to retarget a “fire and forget” weapon after launch, as seen in both JASSM [42] and SLAM-ER [43] (a missile system in the same generation to JASSM), is a tactical difference over the other systems discussed in this section. However, for the purposes of discussion, a weapon that is retargetted after launch can effectively be regarded as having being “deployed” – it merely happens to already be in flight. The issues surrounding post-launch reprogramming will be explored in the next section.

²⁴ Or adopt non-conventional means of performing ISTAR. See the discussion of Scud attacks on Tel Aviv during the 1991 Gulf War [44, Chapter 3], and the potential for leaks from the mass media to fine tune towards “fire for effect”.

C.3. Adversity and Choice





Figure 21: Case Studies – A Journey in Adversity and Choice.

The case studies in this section study the influence of adversity and choice in achieving desirable levels of autonomy. Adversity arises from multiple sources, important ones being the complexity of the environment and situation, and limits of technology. Choice principally arises from the capacity that the system has to generate and select options, and is primarily a question of technology. The journey explores how adversity and choice shape the need for human intervention, and hence raise or lower autonomy.

The starting case for discussion is the firing mechanism in a fully-automatic weapon. Examples weapons include machine guns, assault rifles and submachine guns, and [45] provides descriptions and animations of some of the particular firing mechanisms. For the purposes here, it is enough to regard the firing mechanism as a robot, sitting within the environment of the weapon²⁵, and deployed by the human being to load and fire rounds. Also for the purposes here, the discussion assumes that the human has pulled and held the

²⁵ Conceptually equivalent to production machinery in a factory; indeed, the firing mechanism can be directly compared to a machine for stamping labels on milk bottles coming down a conveyor. The firing mechanism robot was chosen for discussion to compare and contrast the timescales of autonomy.

trigger for continuous automatic fire, and that he/she will not release it. The question is to look at what can stop the firing mechanism robot from performing its task of loading and firing rounds, forcing it to refer back to its human for guidance.



Cross-section of the gas system automatic firing mechanism [45]. On full automatic, holding the trigger [1] releases the bolt [2] to feed a round into the breech [3], where it is fired by a pin in the bolt. Gases from firing are diverted into the gas cylinder [5] driving a piston [4] and hence pushing the bolt back against its spring, ready for the next round.

Figure 22: Gas system automatic fire mechanism.

The awareness held by the firing mechanism robot is the state of the mechanism itself, as it cycles from taking in a round, firing it, and recovering. Crucially, the robot has no direct awareness of the inbound ammunition, other than the presence or otherwise of a round to feed and fire. If the ammunition runs out, the firing mechanism robot is unable to continue its autonomous “programming”, and must refer back to the human to fix the situation. This gives the firing mechanism robot an upper-limit autonomy of about a minute, given fully-automatic rates of fire and usual ready-ammunition volumes.

The firing mechanism robot will also need to refer back to the human if either itself or its environment enters an out-of-tolerance state; that is, the weapon jams from wearing of parts, dirt or moisture in the mechanism, faulty ammunition or other problem. Precision manufacturing and effective upkeep can increase the autonomy to the order of hours or even longer; if not, autonomy can decrease to mere seconds.

The firing mechanism robot example yields two crucial observations. The first is how “in principle” infinite autonomy can be severely constrained by real-world adversities. A perfectly constructed firing mechanism robot, fed from an infinite belt of perfectly manufactured ammunition and within a weapon that was itself inside a perfectly clean

environment would, in principle, load and fire indefinitely at infinite autonomy. Dirt, faults and practicalities – adversities – cut the practical autonomy to minutes. Engineering and upkeep can control adversities and hence improve effective autonomy.

The second observation is that the engineering of a robot to adapt to adversities – providing it choices – can be scaled back, provided that it is possible for a human being to make up the gap. Through training, equipment, and depot support, this is eminently feasible for an automatic weapon.

When the capacity for human beings to make up the autonomy gap is limited or non-existent, the demands on technology go up markedly, in order to provide the robot with resilience and choices to cope with adversity. These demands are strongly seen in the Hubble Space Telescope. The robot in this instance is the space platform, and the discussion here is concerned with the autonomous functions for orienting, stabilising and otherwise setting the space platform to employ its payload of science instruments.



Images from The Hubble Project Servicing Mission 3B / Shuttle Mission STS109 (Space Shuttle Columbia), March 01-12 2002 [47]. The Hubble Space Telescope before retrieval, retrieval on Day 3, and during replacement of the starboard solar array on Day 4. The previous servicing mission was Servicing Mission 3A / Shuttle Mission STS103 (Space Shuttle Discovery) December 19-27 1999 [48].

Figure 23: Hubble Space Telescope.

Launched in 1990, the design intent for Hubble was for Hubble to have a nominal lifetime of 15 years [46, Prologue], and this was later extended by 5 years to 2010 [47]. Hubble was also designed for periodic maintenance and upgrade by Space Shuttle servicing missions [47] [48], and while the concept of retrieval to Earth and then relaunch was discarded [46, Prologue], being “astronaut-friendly” was a design criteria [49]. At best, however,

servicing missions could only be mounted at a periodicity approaching years²⁶, imposing a commensurate autonomy requirement on hardware; indeed the loss of the Space Shuttle Columbia in February 2003 raised questions about continued Hubble operations [50].

On the software side, the provision of high-gain communications and access to the US Tracking and Data Relay Satellite System enables software instructions to be uploaded to the robot. Given, however, that communications may be unavailable, or situations can arise faster than ground-based remote control can react, the Hubble robot needed its own capacity to preserve spacecraft integrity. To achieve this, the robot was extensively instrumented to monitor the state of onboard systems²⁷ (temperature, motion, ...), and provided with sensors (gyroscopic, photosensitive, ...) by which it could establish its position and orientation. Onboard computers processed the data, and it is worth noting that the actual robotic awareness thus generated was very low by modern standards²⁸.

Of key interest here is the concept of *safemode*. The concept behind safemode is that, if the Hubble robot detected itself entering a condition that could damage either itself or its onboard instruments, then the robot would stop, move to a state that was known to be safe, and call for help from its human controllers²⁹. Figure 21 plots the Hubble robot in the early part of its career, when numerous difficulties during deployment lead to repeated entries into safemode [46, Chapter 1]; an effective autonomy of hours.

The Hubble robot safemode is directly comparable with the previous example of the firing mechanism robot becoming jammed – with one important difference. In both cases, the robot has entered a situation that requires a human being to intervene. However, for the Hubble robot, the robot actively refers back to the human being for input, whereas for the firing mechanism robot this is only implied by the robot not working. The first observation is that this potentially places yet another demand on human beings filling the gap³⁰. The second observation is, however, that the provision of safemodes eases the urgency with which the human controllers had to act; that by providing the Hubble robot with

²⁶ Discussion in [46, Chapter 1] implies that the primary – mandatory – requirement for servicing missions is to boost Hubble above the minimum altitude of 325 miles, below which atmospheric drag would pull the satellite out of orbit. Although spacecraft can and do orbit at lower altitudes, 325 miles was regarded as the “line in the sand” in terms of timing, risk and the influence of solar atmospheric heating.

²⁷ Some 6200 specific items of information (“telemetry points”), each with a safe range of operation and triggering an alarm otherwise [46, Chapter 1].

²⁸ This paragraph is based on discussion in [46, Chapter 1], which was written in the early part of the Hubble Space Telescope career, and does not necessarily reflect the operations after its multiple refits. In particular, the onboard NSSC-I computers were described (only partly facetiously) as “bug-hardened, cosmic-ray-proof antique[s]” derived from those used for Apollo.

²⁹ Hubble robot had three safemodes of increasing severity: “inertial hold”, in which the robot terminates any further motion; “software sun point”, a perpendicular orientation that prevents sunlight shining directly on internal systems; and “hardware sunpoint” where the telescope aperture door is closed and systems are powered down.

³⁰ The distinction between a robot “actively” versus “passively” calling for help is somewhat fine, and more of a design intent than anything. For the Hubble robot, entry into safemode conveys to its human operators that something is wrong, *not* by active transmission of bad news but by *non*-transmission of good news.

safemode choices to deal with adversity, the robot buys the humans time to think through remedial options³¹.

Space is an adverse environment for the engineering of electromechanical systems, with adversities including extremes in temperature, bombardment by radiation and high-speed particulates and (for the foreseeable future) tight constraints on mass and energy use. In information terms, however, space is relatively simple – astronavigation and Newtonian dynamics are generally sufficient for location and orientation, and space is acceptably uncluttered³². For the combination of adversity in both the physical and information environments, examples lie in the Mars Rover programmes of 1999 (Sojourner) [52] and 2004 (Spirit, Opportunity) [53]. The discussion here centres on the roving vehicle aspect of the respective Mars exploration missions, particularly the autonomous navigation and mobility of the robot vehicles.



Figure 24: Mars Rover (Artists Impression) [53].

In comparison with the Hubble Space Telescope, the Mars Exploration missions had no capacity to upgrade or adjust hardware – once deployed, no physical human intervention was possible. In addition to this, software intervention is severely limited by the time lag from Earth to Mars from the speed of light – on average about 20 min [53]. This time lag imposed a commensurate autonomy requirement, compounded by a preference for moving under sunlight, and hence in a 4 hr window around high noon [53]. The design objective for Spirit and Opportunity was to move 100 m per Martian solar day (1 sol),

³¹ The Hubble Space Telescope is accompanied by an Observation Support System, manned around the clock by human operators monitoring Hubble and its scientific output. The key point is the need to buy time – Hubble could enter a terminal state faster than human operators could detect and react; moreover, communication to Hubble required use of the Tracking and Data Relay Satellite System (TDRSS), access to which had to be booked.

³² At least for now, with space junk being an acknowledged problem [51].

though a “safe” traverse limited this to about 40 m to enable the Earth-bound humans to see the path ahead of time [53].

The robot vehicles were thus engineered with the capacity to select their own paths, and to adapt to hazards as detected. Safemodes were also designed, for instance, to re-establish communications if blocked or out of range [52]. Robot awareness for navigation was gained through motion (odometers), inertial (gyroscopes, accelerometers), and imagery (landscape analysis, map building, sun tracking) means. It should be noted that the information management had to handle both the robot vehicle needs (navigation, mobility, ...) *and* the mission goals of collecting scientific data. The navigation awareness that might thus be possible in engineering terms is thus necessarily reduced for mission purposes.

The Mars Rover robots show the engineering trade-offs of “acceptable” adversity, and in particular, the shifting of adversity from the physical domain to the information one. If the robot vehicle entered a hardware state where it could not proceed (stuck, overturned, or otherwise immobilised), the mission may well have ended. In contrast, having the robot enter a software state where it could not proceed (safemode, lost or otherwise unable to compute) was far more benign, with the controllers able to intervene. This, in particular, made the “do nothing” choice eminently feasible.

The Mars Rover robots also show the interplay of sufficiency of autonomy and technology state-of-the-art. Infinite *hardware autonomy* was mandatory – once the hardware was bound for Mars, no further intervention was possible³³. However, *software autonomy* on the order of minutes was acceptable, since the Earth control team had sufficient communications access into and out of the robot software systems³⁴.

It is to be further noted that different levels of autonomy were in play. Minute-to-minute / metre-by-metre “tactical” mobility was handled by the Mars Rover robots locally. However, day-to-day / target-to-target “operational” mobility was decided by the Earth control team. The tempo of the situation compared to the necessity or desirability of human control shaped a choice for “tactical” autonomy, but not “operational”.

In this light, the DARPA Grand Challenge 2004 [55] [56] is of interest. This situation can be viewed as one where the physical environment was adverse, with an information environment that deliberately blocked human input, and is thus a step up from the Mars Rover case study. The point is *not* that no contestants succeeded, but to understand the adversities, and see how the engineered choices enabled the autonomy that was achieved by the various robot vehicles on the day.

The DARPA Grand Challenge set a 200 mile course of on- and off-road terrain, specifying that “Competitors’ entries must be unmanned, autonomous ground vehicles, and cannot

³³ Of notable side interest to warfighters – deployed hardware can be captured and used by the enemy. Take for example the experience of 8RAR, in taking casualties from mines lifted from an Australian-constructed barrier minefield. The doctrine for minefields was that they be covered by observation and fire, however, there were insufficient forces to do so [54, Chapter 7].

³⁴ There is an implicit assumption that the robot software systems can be reprogrammed, and that the robot’s computing hardware can take the changes. For comparison, when NASA engineers were seeking ways to correct for jitter in the Hubble Space Telescope, one solution considered was a software patch to drive the space platform’s stabilisation system in directions opposite to the oscillations. Unfortunately, the software could not fit into the 48 kilobyte memory of the onboard computer [46, Chapter 4].

be remotely driven.” [55, p6]. The aspiration was thus for infinite autonomy in both hardware *and* software – the only capacity for human intervention was to terminate the robot participation [55, p6].



“David,” the entry of Team ENSCO, rolls over and ends its run at the DARPA Grand Challenge [57].

Figure 25: DARPA Grand Challenge 2004.

In one sense, the information adversity was somewhat controlled, in that the overall route was known in absolute terms tied to GPS navigation. This, however, allowed attention to focus on the more pressing issues of local navigation and mobility, in adverse terrain. Similarly, the clearing of the course from human vehicles also simplified the information environment, but in a manner orthogonal to the thrust of the Challenge. It should also be noted that there was a 10 hr time limit [55, p6], so in comparison to the Mars Rovers, the “do nothing” choice was curtailed.

Figure 26 presents the final results from the day [58]. While this data is too sparse to draw detailed conclusions, some general points are worth discussing. The first is to note the cases where the robot entered difficulties due to physical limitations of the vehicle, notably Vehicles 21, 7 and 9. The question to ask here is whether the route selection forced the robot into a situation where it ran out of physical choices; that is, if a human driver had been in place, whether he/she would have become equally stuck.

Final Data from DARPA Grand Challenge

As of 5:00 p.m. PST, March 13, 2004

Vehicle 22 - Red Team - At mile 7.4, on switchbacks in a mountainous section, vehicle went off course, got caught on a berm and rubber on the front wheels caught fire, which was quickly extinguished. Vehicle was command-disabled.

Vehicle 21- SciAutonicsII - At mile 6.7, two-thirds of the way up Daggett Ridge, vehicle went into an embankment and became stuck. Vehicle was command-disabled.

Vehicle 5 - Team Caltech - At mile 1.3, vehicle veered off course, went through a fence, tried to come back on the road, but couldn't get through the fence again. Vehicle was command-disabled.

Vehicle 7 - Digital Auto Drive - At mile 6.0, vehicle was paused to allow a wrecker to get through, and, upon resuming motion, vehicle was hung up on a football-sized rock. Vehicle was command-disabled.

Vehicle 25 - Virginia Tech - Vehicle brakes locked up in the start area. Vehicle was removed from the course.

Vehicle 23 - Axion Racing - Vehicle circled the wrong way in the start area. Vehicle was removed from the course.

Vehicle 2 - Team CajunBot - Vehicle brushed a wall on its way out of the chute. Vehicle has been removed from the course.

Vehicle 13 - Team ENSCO - Vehicle moved out smartly, but, at mile 0.2, when making its first 90-degree turn, the vehicle flipped. Vehicle was removed from the course.

Vehicle 4 - Team CIMAR - At mile 0.45, vehicle ran into some wire and got totally wrapped up in it. Vehicle was command-disabled.

Vehicle 10 - Palos Verdes High School Road Warriors - Vehicle hit a wall in the start area. Vehicle was removed from the course.

Vehicle 17 - SciAutonics I - At mile 0.75, vehicle went off the route. After sensors tried unsuccessfully for 90 minutes to reacquire the route, without any movement, vehicle was command-disabled.

Vehicle 20 - Team TerraMax - Several times, the vehicle sensed some bushes near the road, backed up and corrected itself. At mile 1.2, it was not able to proceed further. Vehicle was command-disabled.

Vehicle 15 - Team TerraHawk - Withdrew prior to start.

Vehicle 9 - The Golem Group - At mile 5.2, while going up a steep hill, vehicle stopped on the road, in gear and with engine running, but without enough throttle to climb the hill. After trying for 50 minutes, the vehicle was command-disabled.

Vehicle 16 - The Blue Team - Withdrew prior to start.

Highlighted portions of the final data relate to adversities in the physical environment, adversities in the information environment, physical limitations of the robot and information limitations of the robot. The highlighting is not definitive, but does provide a sense of how physical and information adversities interact.

Figure 26: Final Data from the DARPA Grand Challenge 2004 [58].

The dual question is to ask whether and how the problem could have been shifted from the physical domain (with its terminal consequences to the vehicle) to the information one (with softer costs in terms of time, speed and performance). Vehicle 17 is of particular interest here, in that the (potentially) simple physical challenge of driving over a bush was flipped to the (evidently) difficult information challenge of avoiding it.

The capacity or otherwise to shift adversities³⁵ holds further. Grand Challenge 2004 contained an implicit challenge of budget and the demonstration of nearer-term technologies. To take this point to its ridiculous extreme – if the goal had been to demonstrate mobility between two points, solutions using air transport might have been submitted. The point is that the flipping of “land to air” is conceptually equivalent to flipping from “physical to information” in the engineering of systems, autonomous or otherwise³⁶.

In summary, the conceptual journey from automatic firing mechanism to Hubble, the Mars Rovers and the Grand Challenge 2004 draws out the following key points:

- Although a robot will face direct physical adversities, the systems engineering to address these is in turn shaped by information, technological, economic and other wider adversities. This itself presents opportunities in robot design engineering, through avoidance rather than confrontation of adversities.
- Acceptance of lower autonomy and a lower capacity to deal with adversities eases the engineering problems, notably the equipping of the robot with effective choices. The cost is an increased load on the human being to make up the gap. The robot may not even know that it needs human intervention.
- Of particular interest is the transfer of adversities from the physical domain to the information domain. Doing so enables *hardware autonomy* to be increased at the expense of *software autonomy*. This trade-off may be satisfactory, given that mission needs could demand that hardware be deployed outside of human reach, while reaching software through communication links could be acceptable.
- Different functions of the robot can be engineered to different levels of autonomy. The driver is, again, the capacity or desirability of having the human in the loop compared to the tempo of the situation and severity of outcomes.

³⁵ It should be noted that adversities can be deliberately magnified to advantage, with their disadvantages handled by other domains. Modern fighter aircraft, for example, are designed to be physically unstable, and hence more agile and manoeuvrable, the requisite control being regained through electronic flight control systems. This is an example of emphasising an adversity in the physical domain, knowing that it can be corrected for in the information domain.

³⁶ The shifting or avoidance of adversities by a change in domain is one approach to the concept of “multi-dimensional manoeuvre”, as introduced in the thinking of effects-based operations and related concepts.

C.4. Communications and Programmability




Figure 27: Case Studies – A Journey in Communication and Programmability.

The case studies in this section consider the influence of improved communications technology and programmable software on autonomy. These technologies are of particular interest due to a point raised in the previous case study set, namely the separation of *hardware autonomy* from *software autonomy*. In all the cases considered, the robot hardware is being deployed into a physical domain beyond human reach – indeed, this is the whole point of the robot. In contrast, the journey shows that improved communications and software technology actually improves the human reach into the information domain.

With the focus on the impact of communications and software technologies, Figure 27 specialises the awareness axis to the robot's awareness of *human input*. This is moderate abuse of the intent behind the definition of awareness, which aims at the robot's appreciation of the environment, but remains consistent if the human being is regarded as part of the information environment. Definitions aside, the framework still enables discussion and insights.

The starting case is the Firebee RPV, a reconnaissance drone that saw operational service during the Cuban Missile Crisis and over Vietnam [59] [60] [61]. The Firebee had the capacity to fly a specified flight plan and thus reconnoitre a set of targets, the flight plan being (literally) hardwired into place³⁷³⁸ [61]. However, through a communications link from a ground station or control aircraft, human operators could also take direct control [61]. The Firebee thus shows that a communications link and capacity to override onboard programming can flip a robot from infinite to zero autonomy.



Firebee drone being recovered by helicopter. The drone had the capacity to fly autonomously to the capture zone, with a human controller providing final adjustment.

Figure 28: Recovering a Firebee RPV.

The Global Hawk UAV [62] [63] may be regarded as a technological evolution of the Firebee, notably in the information domain. This is primarily evident in the flight control and mission management systems, which moved much of the “scut work” of flight programming from the human to the robot or support systems. Moreover, while the Global Hawk followed the Firebee in being designed to fly a specified flight plan and thus reconnoitre a set of targets, information technology had progressed to the point where the full flight could be handled from takeoff to landing using conventional airbases [64].

³⁷ The programming was “long and tedious”, with direct programming of flight aspects (speed, altitude, course, sensor use, ...) through jumper leads [61].

³⁸ It is worth noting that the Firebee could be fitted with systems, that enabled reaction to, and thus potential evasion of, intercepts by fighters or SAMs [61].



Figure 29: Global Hawk UAV [63].

The notable aspect of Global Hawk is that it does *not* have a joystick or other flight controls for real-time remote piloting [65] [66]. Inflight reprogramming is specified by waypoints [66], yielding autonomy on the order of minutes, in comparison to zero for the remotely-piloted Firebee. The capacity to have the aircraft manage its own flying³⁹ eases the direct load on the human, and also the communications load to the robot.

The comparison of inflight reprogramming of Global Hawk with Firebee is thus a comparison of autonomy solutions with communications solutions. To restate: Firebee reprogramming was handled by remote-piloting at zero autonomy, whereas Global Hawk reprogramming was handled by waypoints at autonomy measured in minutes. In this particular case, increased autonomy appears to have been desirable and perhaps even necessary, but this is by no means clear for all aspects of aircraft operation.

The particular case of air combat is worth discussing further, for two reasons. The first is that the idea of robot aircraft that fight on their own may be put alongside the idea of robot aircraft that provide a physical projection of a remote human. Both are valid, complementary yet distinct technological paths, and identifying and fusing the best aspects of each is not easy⁴⁰. If the Predator UAV is an indicator of the (publicly announced) state-of-the-art, significant progress could be made on virtual presence technologies.

³⁹ The direct flying of the aircraft is separate and distinct from ensuring that the aircraft is safe. In particular, the management of the aircraft within airspace is and remains a significant issue [65] [66] [67].

⁴⁰ The paper [68] considers this point and subsequent issues, in the context of air combat and AIR6000. While the present author does agree with many of the precepts of [68] (notably the implicit assertions that a “wholly autonomous ‘robot fighter’” will require “true AI [Artificial Intelligence]”, or that the aspiration is to “match or exceed the skills of a competent combat pilot”), the perspective and reasoning chains invite consideration.



The Predator UAV and Remote Pilot Station [69] are an example of fusing autonomy and communications solutions. The robot can fly a preset course for a reconnaissance mission, though the pilot can take manual control. Pilots have likened flying the Predator in this way to “flying an airplane while looking through a straw.” [69], suggesting that virtual presence has a long way to go. However, overall awareness of the battlespace can be better than in contemporary aircraft, from provision of networked data [66].

Figure 30: Predator UAV and Remote Pilot Station.

The second point for air combat is to note that, at beyond-visual range, the pilot is working off data provided by machine sensors (radar, infrared, emissions detection, datalink, ...), rather than his/her own biology. Significant research is going into the fusion of this data to improve the situation awareness of the pilot, but this research is necessary precisely *because* the pilot as human combatant has to go through a “two-step” – data from a machine sensor into a human sense and hence cognitive outcome. For the robot, the data from a machine sensor can be directly fed into its awareness, where “robot awareness” follows the definition introduced in Section 2.1 of this paper. To put it another way, to a robot, there is no difference between “beyond-visual range” combat and any other kind – it is all data provided by sensors, with no distinction on their physical location.



The fusion of data for the pilot is a significant proportion of the JSF project. Data from onboard and offboard sensors is fused into a single, overall picture of the battlespace, while the integrated helmet projects an image onto his retina so that navigation, targetting and terrain information appear to float in space, keyed to the direction he/she is looking in [70] [71].

Figure 31: Concept demonstrator for the Joint Strike Fighter cockpit.

It is worth emphasising the overall issue, that the comparison of autonomy with communications solutions is exceptionally strong for combat and the use of lethal force. It is recalled from Appendix C.1 that robots with lethal capability may be regarded as a conceptual scaling of a land mine, so the release of lethal force presents no direct technological issues, the devil being in the collateral effects. To put it another way, one could engineer the robot with a capacity to decide whether a target is valid for engagement, but one also has the choice of tightening the decision loop back to a human to authorise release⁴¹. In parallel, and recalling the discussion of adversity and choice in Appendix C.3, there is the question of whether it is more effective to equip the robot with the capacity to combat its adversaries, versus bringing in the human in the loop. Put together, the choice of designing towards a “smart land mine” (autonomy solution) versus a “human virtually present” (communications solution) admits no general answer, and even in a particular context it is still difficult.

In a separate point, the case studies point to a distinction between *technological autonomy* and *operational autonomy*. In terms of the model introduced in Section 2.1, the human being may choose to “redeploy” the robot, proactively sending input to the robot rather than reactively filling its gaps. Hence, while the robot may have the *technological* capacity to operate without human input at one tempo, the actual *operational* tempo of references to a human is tighter. That operational autonomy can and will be tighter than technological autonomy is both a boost and encumbrance to robot system developers: it enables the human decision-making capacity to be pulled into the robot⁴², but it also means that humans can and will push inputs⁴³.

In summary, the conceptual journey from Firebee to Global Hawk and Predator draws out the following key points:

- The capacity to communicate with and thus reprogram the robot once deployed makes *software autonomy* significantly different from *hardware autonomy*. In particular, it separates *technological autonomy*, driven largely by endurance of hardware, from *operational autonomy*, driven by the human adaptability.
- Communications and information technology enable the complementary yet distinct paths of autonomous solutions (“smart independent robot”) versus communication solutions (“virtual human presence”). This issue is seen most starkly in combat and the use of lethal force, with mine warfare, beyond-visual-range combat and directive command all providing precedents.

⁴¹ “The limiting factor ... is not the technology, but the bureaucracy – getting the necessary permission to engage a target.” – Attributed to Frank Pace, executive vice president of General Atomics Aeronautical Systems Inc (the manufacturer of the Predator UAV) [72].

⁴² Note that this includes the trivial case of having the robot “do nothing”, leaving it to the human being to push instructions forward to perform a task, before returning to an inactive state.

⁴³ This issue is directly comparable to human command and control, notably the concepts of “reach back” (or “interfere forward”). An 18th century ship captain *had* to be trusted with delegation and authority, for there was no way to adequately communicate with him by the available means. In the 21st century, ship captains are now potentially in high-capacity real-time contact with higher, land-based command. In theory, this eases the requirement on the modern ship’s captain to be autonomous, by “reaching back” for input and resources; the same, capability, however, enables higher command to “interfere forward”.

C.5. Controlling Consequences

The previous case study sets considered robots within the dimensions of *autonomy* and *awareness*, within the question of sufficiency. The case studies in this section are about how the dimension of *situation* can also be utilised. The central aspect is in controlling the consequences of undesirable robot actions, rather than in seeking to make the robots “smarter” or otherwise more capable. This section is also an opportunity to consider examples of how robots have already been considered in arms control thinking.

The first case is the Hague Convention VIII relative to the Laying of Automatic Submarine Contact Mines [73], of which the following articles are of particular interest here:

Article 1. It is forbidden --

1. To lay unanchored automatic contact mines, except when they are so constructed as to become harmless one hour at most after the person who laid them ceases to control them;
2. To lay anchored automatic contact mines which do not become harmless as soon as they have broken loose from their moorings;
3. To use torpedoes which do not become harmless when they have missed their mark.

Art. 2. It is forbidden to lay automatic contact mines off the coast and ports of the enemy, with the sole object of intercepting commercial shipping.

Art. 3. When anchored automatic contact mines are employed, every possible precaution must be taken for the security of peaceful shipping.

The belligerents undertake to do their utmost to render these mines harmless within a limited time, and, should they cease to be under surveillance, to notify the danger zones as soon as military exigencies permit, by a notice addressed to ship owners, which must also be communicated to the Governments through the diplomatic channel.

Art. 4. Neutral Powers which lay automatic contact mines off their coasts must observe the same rules and take the same precautions as are imposed on belligerents.

The neutral Power must inform ship owners, by a notice issued in advance, where automatic contact mines have been laid. This notice must be communicated at once to the Governments through the diplomatic channel.

Art. 5. At the close of the war, the Contracting Powers undertake to do their utmost to remove the mines which they have laid, each Power removing its own mines.

As regards anchored automatic contact mines laid by one of the belligerents off the coast of the other, their position must be notified to the other party by the Power which laid them, and each Power must proceed with the least possible delay to remove the mines in its own waters.

The intent of the Hague Convention VIII was to protect neutral commerce and uphold the principle of immunity of enemy merchantmen from attack without warning [73, Introduction]. Whether this geostrategic goal was achieved is beyond the reach of this present study. What can be seen is how the immediate consequences of robot failure are controlled, where the robot is an automatic submarine contact mine: for instance, if the robot is “unanchored” then it is to “become harmless one hour at most after the person who laid them ceases to control them”, or if the robot is “anchored” then it is to “become harmless as soon as they have broken loose from their moorings”. Control of consequences can also be seen in the articles for Convention signatories on the notification about mine deployment, or their post-conflict removal.

By contrast, the Ottawa Convention on the Prohibition of the Use, Stockpiling, Production and Transfer of Anti-Personnel Mines and on Their Destruction [9] has the following articles of note:

Article 1

General obligations

1. Each State Party undertakes never under any circumstances:
 - a) To use anti-personnel mines;
 - b) To develop, produce, otherwise acquire, stockpile, retain or transfer to anyone, directly or indirectly, anti-personnel mines;
 - c) To assist, encourage or induce, in any way, anyone to engage in any activity prohibited to a State Party under this Convention.
2. Each State Party undertakes to destroy or ensure the destruction of all anti-personnel mines in accordance with the provisions of this Convention.

Article 2

Definitions

1. "Anti-personnel mine" means a mine designed to be exploded by the presence, proximity or contact of a person and that will incapacitate, injure or kill one or more persons. Mines designed to be detonated by the presence, proximity or contact of a vehicle as opposed to a person, that are equipped with anti-handling devices, are not considered anti-personnel mines as a result of being so equipped.
2. "Mine" means a munition designed to be placed under, on or near the ground or other surface area and to be exploded by the presence, proximity or contact of a person or a vehicle.
3. "Anti-handling device" means a device intended to protect a mine and which is part of, linked to, attached to or placed under the mine and which activates when an attempt is made to tamper with or otherwise intentionally disturb the mine.

In terms of robots, the Ottawa Convention contrasts with the Hague Convention VIII in seeking to control the technology of autonomy and awareness: "designed to be exploded by the presence, proximity or contact of a person". The intent of the Ottawa was humanitarian, in the post-conflict impact of anti-personnel mines on non-combatants. It can thus be argued, therefore, that the Ottawa Convention should have sought agreement in the dimension of situation, in controlling the consequences of anti-personnel mines littering battlefields post-conflict; for instance, in seeing that such mines are removed, or are otherwise automatically rendered inert. Similarly, it can also be seen that the Ottawa Convention potentially presents difficulties for any proposals for land combat robots that would automatically engage an enemy, in using a black-or-white definition of robot awareness ("presence, proximity or contact of a person") and taking a narrow definition of the effects employed ("exploded"). Overall, it is argued that in targeting consequences directly, the Hague Convention VIII will be more resilient and lasting than the Ottawa Convention.

In summary, the examples of the Hague Convention VIII and the Ottawa Convention show that:

- Arms control thinking has already had to consider robots.
- The control of robot consequences (within situation) is a viable means of enabling robot use, even under limited robotics technology (autonomy and awareness).

Appendix D: Bottleneck Issues

The potential from future technology is studied through *bottleneck issues*, enduring problems in the construction of autonomous systems. For each issue, the discussion presents the reasons for its being a bottleneck, and then postulates technologies with the potential to improve technological sufficiency. The discussion includes bottleneck issues where technological sufficiency has been attained, in order to illustrate and calibrate the analytic process.

D.1. Communications and Programmability

The bottleneck issue centres on the capacity to communicate with the robot once deployed, and thus change its behaviour. This issue weighs into the balance of autonomy solutions against communications solutions.

Communication with robots started with the harnessing of the following mechanisms:

1. *Modulated electromagnetic radiation*, notable examples being radio and laser. Modulated EM radiation has the strengths of not requiring physical media connecting human and robot. Conversely, while EM radiation communications can be made to pass through solid material (land, water), the trading of frequency for wavelength reduces data carriage. Similarly, higher data carriage currently requires higher frequency radiation, forcing shorter wavelengths and hence limits on the radiation path (line-of-sight, fibre-optic as a channel). Finally, EM radiation travels at the speed of light, a fundamental upper bound on communications speed.
2. *Modulated electricity*. Modulated electricity requires a conductive medium between the human and robot. However, given that metals are conducting, ductile and deformable (copper, aluminium), appropriate mediums can be constructed. Communication by modulated electricity is, again, limited by the speed of light. The frequencies are also bounded by energy loss and attenuation issues.
3. *Modulated acoustic radiation*. Modulated acoustic radiation requires a solid, liquid or gaseous medium between human and robot. Communication is then limited by the speed of sound within the medium, dispersion and hence energy loss, and channelling issues on wavelength choice.

It is also notable that all three mechanisms allow the use of modulated sinusoids⁴⁴, enabling cross-transfer of concepts and technology.

Reprogrammability essentially started with electronics, in that instructions could be encoded as electronic states (in silicon), and electronics could also be used to change those instructions. It is possible to imagine doing this using Industrial Age mechanical technology, but not with hardware that could be readily deployed.

The extent of technological sufficiency based on these technologies can be seen in the robots and systems in or entering service. In aerospace, radio communications with no reprogrammability were sufficient for the Firebee RPV, while reprogrammability of 48K onboard computers was sufficient to overcome communication limitations to the Hubble

⁴⁴ Allow but do not compel. For instance, a heliograph modulates solar radiation, on or off.

Space Telescope. In the underwater environment, wire-based communications have enabled operators to reprogram the heading and targetting of torpedoes⁴⁵.

In considering technology in the near to mid term, the starting point is to consider the potential for new mechanisms for communication. Such mechanisms would entail new science, and while they cannot be excluded, it can be safely stated that the mechanisms would need to be in areas of science that are presently not well understood. Candidate areas to monitor include:

- *High-energy physics*, for the potential to discover new particles or forces that might be tapped as media⁴⁶.
- *Quantum entanglement*. The quantum teleportation experiment [75] [76] potentially points to the transfer of data encoded at the quantum level, using laser communications but not encoding the data within a modulated sinusoid.

In the absence of new communication physics, the future for modulated EM radiation, electricity and acoustic radiation may be considered. These mechanisms are currently benefiting from sustained improvements in electronic-based data processing, notably:

- *Analog to Digital conversion*, at speeds compatible with the sinusoids into which data is being modulated⁴⁷. This, together with the raw computational power of modern VLSI microprocessors, are enabling specialised processing hardware to be replaced by generic hardware running specialised software (software-based radio), yielding savings in weight.
- *Phased array technology*, reducing the footprint made by antenna hardware on the robot. These yields savings in physical design (weight, bulk, integration into the platform, energy use, scalability). The value can be seen by comparing with, for instance, the impact on the Global Hawk design to accommodate the dish-type satellite transceiver.

The analysis here postulates that, in the near- to medium-term future, robot system developers will take advantage of the unrealised gains still to be made from these technologies, that are in turn consolidating progress in microelectronics. Under this assumption, a number of predictions can be made.

The first prediction is that EM radiation communications will continue to be dominant, except in those situations where wire-based or acoustic communications are acceptable (short range) or physically necessary (underwater, underground). It follows immediately that system developers using EM radiation communications will run up against path (line-of-sight) issues, and in particular, a chronic shortage in satellite time. As noted, high-capacity EM radiation communications entail high frequencies and hence short wavelengths, so communications to a robot at over-the-horizon ranges requires some kind of relay. Known, potential solutions to this issue include:

⁴⁵ So long as the wire is in place. Once cut, the torpedo depends solely on its own guidance system.

⁴⁶ It is worth noting that, at one time, both electricity and radio waves were regarded as high-energy phenomena. See [74, Interlude B].

⁴⁷ The payoff from this can be seen in the wire-based communication of data over telephone networks. There were no fundamental changes to the physics or frequencies, yet digital encoding and evolution thereof (modems to ADSL) enabled data rates from the order of 72 bits/second to 56000 bits/second.

- *Substantially cheaper access to space and orbit.* Insertion into space and orbit⁴⁸ are, at present, of such a cost that space platforms need to be long-life, high-capacity nodes. This, in turn, drives up the demands on the space platform design, in a vicious cycle of mass, energy and cost. To break this cycle, and enable relay satellites⁴⁹ with shorter lifetimes that are “used and replaced”, costs need to become comparable to that of using an air-launched precision-guided munition in anger. For indicative numbers: about \$100/kg to low-earth orbit⁵⁰ for a payload of up to 100 kg⁵¹, by a system with availability on the order of 1 week with a total acquisition cost on the order of \$10 billion and an annual operations cost on the order of \$500 million/year (2004 Australian dollars)⁵².
- *Relay aircraft.* Possibilities here encompass both lighter-than-air and heavier-than-air, and unmanned and manned aircraft; the latter notably including aerial refuelling tankers on holding patterns. The key difficulty is that long-endurance aircraft are, in themselves, high-value assets oriented to other missions. The situation may change in the near term as such aircraft take on high-capacity communications to achieve their missions, potentially reducing (but not eliminating) “scab impact” of robot system developers seeking to integrate relay gear.
- *Leveraging commercially deployed networks.* With extensive deployment of mobile phone and other wireless networks, robot system developers could choose to use the bearers that are present, to the extent that they can. For robots in a military context, the key question is an understanding of whether and how the military force is expected to be expeditionary. By way of comparison, the Vered Harim (Mountain Rose) communications system being developed for the Israel Defense Forces [82] [83] is based on cellular phone technology, which can be placed in the context of the IDF conducting its ground operations in Israel and its immediate region⁵³.

It is worth noting that unconventional (terrorist) forces have a *social* advantage over conventional forces, in being able to parasitically benefit from the confidence built into a civilian network. A notable example is the use of mobile phones as remotely-activated bomb triggers, as occurred in the Bali Bombing [84]. Although the activation call was eventually traced⁵⁴, it was just one within the vast volume of regular traffic. In

⁴⁸ The distinction is worth noting. Light payloads can be lofted into space by helium balloon, however if released, they will fall back to Earth. To stay in space, payloads need to be accelerated to orbital speed.

⁴⁹ Not necessarily in geostationary orbit, though this does simplify the satellite acquisition problem. Compare with the Iridium satellite constellation, which operated at lower altitude.

⁵⁰ Compare for example with a Delta IV launch being on the order of US\$10000/kg (based on values at [77]), with booster bookings being a major capital item [78]. Note that Delta IV configurations are intended to loft up to 24 t into low-earth orbit or 10 t into geostationary transfer orbit.

⁵¹ 100 kg is the upper limit of systems termed “microsatellite” under the classification scheme at [79].

⁵² For comparison, the AIR6000 programme to replace the F-111 and F/A-18 aircraft is slated at \$12 billion [80], with current annual operating costs being roughly \$244 million and \$300 million respectively [81, Chapter 9]. The suggestion is that of comodification from a superpower-capability to that affordable by a small- to moderately-sized nation state with a high-technology economy.

⁵³ On this basis, the presence of cellular mobile network hardware can be pre-mapped and upgraded as necessary, and frequency deconfliction can also be managed ahead of time [83].

⁵⁴ The SIM card was not destroyed, and the time window was narrowed by using seismic data [84].

conventional warfighting, it might be possible to suppress the network, however for peacekeeping this could be counter-productive, since civilian communications could be an important part of (re)establishing normal conditions. The further application of civilian communication networks to robot control is thus neither avoidable nor easily suppressed.

The second prediction is that robot system designers will run up against the speed of light as an absolute upper limit to round-trip time⁵⁵. This has already occurred in the space program, as Russian controllers found with the Lunakhod remotely-controlled Moon explorer in the 1970s [85], and also shows up on terrestrial scales, in that communications via a satellite in geostationary orbit will have a one-way trip time of about 2 seconds. Relaying through nodes at lower altitudes can shorten the path length, though the cost may be processing delays within nodes.

Communication delays impact primarily on systems that seek to put a human operator in real-time control of the robot – that is, the robot is deployed at zero to low autonomy. This aligns with an earlier thread, that autonomy (software autonomy in particular), is used to make up for deficient communications. Continuing on this line, the question is whether, within deployment ranges that are not compromised by delays, communications technology could be reasonably expected to support “human virtual presence”.

Insight can be gained from looking at vision as the case in point, moving from “viewing through a drinking straw” [69] to providing visual data to a remote human, of equivalent quality to their being present⁵⁶. A nominal goal is to image a hemisphere ($180^\circ \times 180^\circ$) such that a 1 metre object at 100 km yields a single pixel⁵⁷, to 24bit colour and at a refresh rate of 240 Hz⁵⁸. This is a raw data flow approaching⁵⁹ 600 Tbit, some 15000 times that of the 40 Gbit satellite laser communication systems under development [88]. However, the merit of this value is *not* that it is 4 decimal orders of magnitude larger than what can be presently achieved, but that the multiple degrades markedly with lower thresholds on acceptability for “virtual presence”. For instance, if the resolution is reduced to that a 1 metre object at 10 km, the raw data flow rate drops to 6 Tbit, or a factor of 150 over envisaged systems. The application context will thus re-enter the debate of sufficiency – for instance, 100 km may be necessary to match the visual acuity that fighter pilots have achieved under ideal conditions, but 10 km might be sufficient for land combat in urban environments.

⁵⁵ For acoustic communications, issues hold with respect to the speed of sound in the support medium.

⁵⁶ In focussing on communications, the exploration here ignores the question of whether and how such information could be gathered, although this is an issue in its own right.

⁵⁷ A full hemisphere may be overgenerous compared to human-factors research on the field of view required to make the user “feel” present [86]. Similarly, 100 km (50 nautical miles) may be similarly generous when compared with the normal visual acuity goal of resolving a ~10 cm symbol at a distance of 6m (6/6 vision). For the purposes here, however, it is better to overestimate the demand on information transfer, with allowance for eyes traversing across the sky or vehicle closure rates.

⁵⁸ Animation can be achieved at rates as slow as 12 frames per second, using double or quadruple projection of frames to avoid flicker. The discussion here takes the 30 frames per second for digital video, with a factor 4 multiple to allow for resolution of crossing targets.

⁵⁹ $\left(\pi / (1 \text{ m} / 100 \text{ km}) \right)^2 (24 \text{ bit}) (240 \text{ s}^{-1}) \approx 568 \times 10^{12} \text{ bit/s}$.

Given, though, that the raw data flows for “human virtual presence” will not be easily attained in the near future – that is, no trivial solution⁶⁰ – near to mid term technology development may look at reducing the volume of data. Possible approaches include:

- *Compression.* The utility of compression, and of “lossy” compression⁶¹ in particular, is bound up in the question of the information needed for the human to gain a sufficient feeling of presence. By way of comparison, the success of the MP3 audio compression standard is due, in part, to its maintaining “CD quality” sound while achieving substantial compression efficiency, by discarding sound data not required for the feeling of “CD quality”.
- *Simulation.* Simulation runs in partnership with compression, in filling out information that is not being acquired or transmitted from the robot⁶². Again, the question is what the human needs in order to gain a sufficient feeling of presence. As an example, it is probably unnecessary and undesirable to convey the full g-loading of a robot aircraft pulling a 9g turn, but a sense of acceleration may be desirable to help the human operator stay physically oriented.

It is difficult to predict whether such approaches can, in the near to medium term, achieve “human virtual presence” across all application domains. However, in converse, it can be safely predicted that, at least within some applications, some autonomy will need to be engineered into the robot, to fill the gap from the human being physically remote.

To finish the analysis here, the final point considered is whether and how new approaches to computing will change the reprogrammability of robots, and hence the load on communications. Within this lies the implicit assumption of new computing hardware being introduced onto the robot, to enable new approaches to remote processing of data⁶³. The following technologies are then of interest:

- *Hardening and protection of electronics and optics.* The electronic and optical components used on robot systems have, to date, significantly lagged behind those used in home or office devices. A significant reason for this is the need to shield components against radiation, a problem that is particularly acute in aerospace. More powerful computing hardware will open options for robot system developers.
- *Maturation of software engineering.* The software industry is slowly, and somewhat painfully, learning how to develop complex yet reliable software. This capacity underpins robot system developers seeking to build software with high awareness.

⁶⁰ For comparison, during the 1980s, the data rates enabled by modems drove interest in compression for text files, at a time when 10 k was regarded as a large text file. By the late 1990s, e-mail attachments on the order of 1MB (some 100 times larger) were routine, an observation reflected in the Apple decision to omit the 3.5 inch floppy drive from the iMac and subsequent Macintosh computers.

⁶¹ “Lossless” compression, as the name implies, allows total reconstruction of the original data. “Lossy” compression deliberately discards “unimportant” data to achieve higher performance.

⁶² Viewed from a compression perspective thus: data compression yields a compressed form *together with* a dictionary for reconstruction. If the receiver already knows the dictionary, it does not need to be transmitted. Moreover, once the compressed data is received, the dictionary is the basis for building the simulation.

⁶³ Again the question of sufficiency, this time of computing and communications. If a new computer system was developed tomorrow, it could be networked to the robot ... if communications were sufficient.

D.2. Locating Robot in Environment

The bottleneck issue centres on providing the robot the capacity to locate itself within the physical environment, in both absolute and relative terms. Given that a robot is being built to inspect, move within or otherwise interact with the physical environment, the physical location of the robot may have both direct and collateral effects. This issue thus weighs into the severity of the situation against the awareness that can be achieved, and hence the autonomy that may be acceptable.

Overall, this issue is about the capacity to generate an *information model* of the robot's environment. The difficulty lies in articulating a modelling process that captures the salient information and discards the irrelevant, noting that the words "salient" and "irrelevant" are ill-defined and context-sensitive. Equally however, the process within the human brain is also ill-understood at this time.

To date, approaches in machine navigation have simplified the problem by simplifying the environment, robot system designers either identifying or providing the "salient information" in a form readily detected by machine (electronic) means. Important examples include:

- *Astronavigation*, with primary use in space systems. The position of bright stars can be plotted and predicted externally, and their spectra are unique and readily detected by electronic sensors. The robot's information model is that of an environment solely consisting of these bright star beacons, a model that is sufficient for space navigation.
- *Inertial*, with widespread use in air systems. Acceleration and rotation can be measured by a variety of means (weighted springs, mechanical or laser-ring gyroscopes), and integrated to record displacement from some home position. The robot's information model is that of an environment that is totally empty, with only the home position and displacements from it. This is sufficient for aerospace and maritime navigation, over timeframes limited by the drift in the navigation equipment.
- *Electromagnetic Referencing*, with notable use in maritime and land systems. The magnetic or electromagnetic phenomena may be natural (the Earth's magnetic field) or human-generated (radio waves for LORAN, radio or laser beam-riding). The robot's information model is that of an environment defined by the electromagnetic phenomena. This is sufficient for aerospace, maritime and some land navigation, to a precision limited by noise in the electromagnetic signal.
- *Comparison with Supplied Model*, with notable use in air systems. Under this approach, the environment is described numerically, typically as digital images (terrain comparison) or vector geometry (3D target modelling). This being the robot's information model, sensors can gather data from the physical environment for comparison and thus location. This approach has enabled the precision achieved in recent air strike systems (Tomahawk, JASSM), however, it requires an ISTAR system to generate the information model for the robot.

- *Global Positioning System*, with use in air, maritime and land systems. GPS may be regarded as a combination of astronavigation and electromagnetic referencing, in that a known satellite constellation transmits a readily processed radio signal. The robot's information model is thus that of the satellite constellation, the environment otherwise regarded as empty. The precision that can be achieved, together with low cost to GPS users, was one of the reasons for UAVs "coming of age" during 2000–2005 [89] [90]⁶⁴.
- *Electronic Markup*, with use in industrial and commercial settings. The basic idea is to simplify the physical environment (structuring or semi-structuring [91]) and provide readily detectable navigation markers (brightly colour panels [92]). This essentially constitutes an "electronic markup" of the physical environment, the "markup" becoming the robot's information model. This approach is sufficient for production, warehousing and similar industrial activities where the physical environment is static or otherwise controllable.

It can be seen that, by simplifying⁶⁵ the environment, robot designers have reduced the demand for data to populate the corresponding information model. The result is a reduced demand for awareness.

In the near to medium future, robot system designers can and will make use of these existing approaches, to the limits that the situation allows. As a particular example, the integration of UAVs into US Navy carrier operations is likely to depend heavily upon Shipboard Relative GPS [93], that is, simplifying the carrier environment into an information model based on GPS. The frontier for progress is that of environments that cannot be simplified without losing vital content, and/or where the pre-assembly of a model is prohibitively expensive; urban terrain being a key example⁶⁶.

For exploration purposes, the discussion here considers the problem of depth-mapping a sphere, originating from the robot, to a precision of 1 degree in angle⁶⁷ and with a depth error of 1 cm at 1 m range. A depth map of this size could enable a robot to be sufficiently aware of an urban setting to navigate it without damage. Notable issues include:

- *Memory Requirements*. To record depths of a sphere ($360^\circ \times 180^\circ$) to 1 degree of arc will require about⁶⁸ 254 kilobytes of memory. This is readily achieved with current solid-state memory technology⁶⁹. There may be intrinsic advantages to using holographic memory [94], by using 3D memory to store 3D data, in terms of access and understanding of relationships within the data⁷⁰.

⁶⁴ The three reasons identified in [90] were reliable sensors, via modern electronics; low-cost automated flight control and navigation, with GPS being a key enabler; and operator-UAV control, via datalinks.

⁶⁵ Arguably, *substituting* the environment.

⁶⁶ A robot operating within a city streetscape cannot operate solely off GPS or inertial guidance, since it will not have awareness of the buildings. A 3D map of the city could be supplied, however, this puts demands on the ISTAR system to keep the map up to date.

⁶⁷ Yielding an arc length error of about 1.75 cm at 1 m range.

⁶⁸ $(360) \times (180) = 64800 \text{ floats} = 259200 \text{ bytes} \approx 254 \text{ kilobytes}$.

⁶⁹ For comparison, an angular precision of 1 minute of arc requires 933120000 floats, or about 1 gigabyte of memory; again readily achievable using current solid state memory.

⁷⁰ For example, following surfaces and identifying corners.

- *Depth Mapping.* Present day laser rangefinding equipment⁷¹ approaches the depth range accuracy and angular precision goals. Present day ultrasound sensors also approach the depth range accuracy, though at lower angular precision⁷².
- *Refresh Speed.* Taking the laser rangefinding approach as indicative, the sphere could be imaged in under⁷³ 3 seconds.

This cursory analysis indicates that depth mapping a single sphere can be achieved with present-day systems. The challenge comes when the robot moves, in processing the next sphere of data. The question of deciding whether and how a given sphere of data overlaps, or otherwise has commonality with, another sphere is not trivial. In its base form, the problem involves the comparison of subsets⁷⁴ from the first sphere with those from the second. Although a given comparison may be executed quickly, the number of such subsets is combinatorially large, leading to an explosion in the number of tests to execute. For the near to medium future, this problem will *not* be solved *trivially*⁷⁵ – the combinatorial growth outpaces Moore’s Law progress in computational power. Intrinsically parallel approaches via DNA [97] or quantum [98] computing may yield solutions, over and above conceptual insights (heuristics).

D.3. Status of Robot

The bottleneck issue centres on providing the robot the capacity to monitor its physical status. Environmental drivers may compromise the robot or its subsystems, with follow through to mission capability, potentially requiring human intervention. The awareness that the robot has of its status, and its capacity to deal with it within the context of situation, will drive the autonomy that can be achieved.

At present, the status of robot subsystems (temperature, presence of moisture, structural stress, presence of corrosion) can be monitored by a variety of electromechanical and electronic means, assembled into an *information model* by conventional computing. It is noted that this information model of the robot is, by necessity, a limited representation of the full robot status, due to:

- *Sensor inaccuracies*, resulting in losses in fidelity from robot subsystem to sensor.
- *Limited sampling*, for sensors cannot be arbitrarily placed into the robot, but are themselves subsystems requiring integration.
- *Memory and computing limits*, in that the construction of the information model is a load on the robot’s computing capability.

⁷¹ DeltaSphere-3000 3D laser scanner product literature [95]: Range of 1 ft – 40 ft; Range Accuracy of 0.3 in at 40 ft; Angular Accuracy of 0.015 degrees; Scan Rate of 25000 samples/second; Field of View –60 degrees to +90 degrees in elevation, 360 degrees in azimuth.

⁷² Product data sheet for Honeywell 943 Series Ultrasonic Distance Sensors: Beam angle at 8 degrees [96].

⁷³ See footnotes 68 and 71. (64800 samples) / (25000 samples/second) $\approx$ 2.6 seconds.

⁷⁴ Comparisons may be based on imaging data (colour, brightness, ...).

⁷⁵ In the sense described earlier in Footnote 60.

In the near to medium term future, advances in microelectronics/optics and electronic computing will ease the engineering in all three aspects⁷⁶. Significant improvements could also arise from:

- *Smart materials*, and in particular, the embedding of microscale sensors directly into the materials making up the robot and its subsystems. Doing so reduces the inaccuracy and limited sampling issues – effectively, the sensor is the material. If, in addition, the embedded sensors can be rapidly and arbitrarily interrogated, then the memory and computational issues are eased – instead of having to assemble the sensor data into an information model, the smart material constitutes its own model within itself.
- *Grid computing*, and the thinking behind it [99]. The concepts and technology for assembling distributed and disparate systems into a seamless end-user whole, with recovery of individual component failure, applies to the problem of assembling a robot from subsystems. Moreover, the thinking also applies to the robot as a networked system, drawing upon and providing services.

D.4. Locating Other Entities in Environment

The bottleneck issue centres on providing the robot the capacity to locate other entities in the environment. This builds on the earlier discussion of the robot's location within the environment, but adds the dynamic element of entities moving in ways not easily predicted. The interaction the robot has with such entities will have direct and collateral effects. This issue thus weighs into the severity of the situation against the awareness that can be achieved, and hence the autonomy that may be acceptable.

The earlier discussion of physically locating the robot within its environment carries forward to the issue here. In particular, entities can be tagged by means readily detected by machine (electronic) means. In the warfighting context, this is somewhat easier to achieve with friendly units (IFF/GPS transponders) than with neutrals or hostiles, particularly as the latter may be actively seeking to *avoid* detection. The scope of the present issue thus encompasses the space of sensors and the exploitation of sensor data.

It is also necessary to note that the location of other entities is bound to the severity of consequences. A missile launched into hostile airspace may not need to detect entities (including anti-missile fire) if the costs of collision are acceptable. By contrast, the cost of a collision between a UAV and friendly traffic are serious⁷⁷.

In the near to medium term future, progress will come with the general push to improve battlespace management (notably airspace and friendly fire deconfliction). Conversely, robots will need to be able to report their location and intended movements, at tempos demanded by battlespace management systems. This may compromise robot freedom of movement^{78 79} and will add to communication loads.

⁷⁶ In particular, the earlier discussion of hardening and protection of electronics and optics also applies here.

⁷⁷ Commentary on the X-47A Pegasus, "... primary mission is to demonstrate technology for a catapult-launched, arrested-landing UCAV. This is a make-or-break issue: Traditionally, naval commanders regard any jet-powered, explosives-carrying unmanned vehicle approaching their vessel as a missile, and respond accordingly." [100]. Also compare with progress in clearing Global Hawk to operate in civil airspace [67].

⁷⁸ For comparison from World War I, British troops in Western Europe were expected to advance at a pace compatible with artillery barrages, and junior officers were ordered to stay within (telephone)

D.5. Target Modelling

The bottleneck issue centres on providing the robot the capacity to model targets. The word “target” can be applied to anything that the robot might interact with in its environment, however, the issues are particularly acute when the robot’s action generates a lethal effect. This issue thus weighs into the severity of the situation against the awareness that can be achieved, and hence the autonomy that may be acceptable.

To date, the basic approaches to target modelling have been:

- *Engage Unless Friendly*. Under this approach, the properties for being *friendly* are described, and the robot can engage anything else. Friendly identification can be achieved by physical deconfliction (minefields, free fire zones), or by information tagging (IFF). The key difficulty of this approach is in the risk of casualties among neutrals and/or non-combatants.
- *Comparison with Supplied Model*. Here, the target is described numerically, typically as digital images (digital scene comparison) or vector geometry (3D target modelling). This being the robot’s target model, sensors can gather data to seek a match. This approach has enabled the precision achieved in recent air strike systems (Tomahawk, JASSM), however, it requires an ISTAR system to generate the target model for the robot. Mobile or otherwise dynamic targets pose further challenges, in that the robot system designers have to seek and use numeric descriptions that are robust against target variation.

These approaches have been achieved to varying degrees of success with technology to date, and the near to medium term future should see refinements. For operations in cluttered or urban environments, in a world that is media aware and casualty averse, there is likely to be broad emphasis on the second approach. Conversely, the first approach will continue to be relegated to those rare, high-intensity, situations demanding rapid and decisive engagement on timelines too short to allow a human in the loop. This trend will be exacerbated by the increasing potential from communications technology, as explored above, and hence the capacity to put a human in the loop. It should be noted, however, that these are observations for conventional military forces – non-conventional forces may have neither the communications capability, nor the constraints from casualty aversion.

The third approach, and the ongoing goal of artificial intelligence research, is for *training by similarity*. Under this aspiration, the robot is supplied with examples of valid targets, and instructed to engage “those like them”. In contrast to the above approaches, the robot derives its own numeric description⁸⁰ of what constitutes a valid target, but it is targets that being modelled, not friendlies-with-exclusion.

communications range of senior officers [101]. By contrast, late World War I German tactics of insertion and bypass accepted a loss of control on troop movements to secure advances [101].

⁷⁹ This issue can also be seen in “free flight” concepts for air traffic management [102]. There are two extremes: Where aircraft are kept to within the corridors of airspace that the ground controllers are able to manage, versus aircraft being totally free to manoeuvre with no ground control. With robot aircraft, the “free flight” situation potential becomes more complex, but the central issue is about the granularity required in a central traffic management picture.

⁸⁰ It may be preferable for the descriptions to be comprehensible by humans, but it is not clear that technology can achieve this additional hurdle.

Key issues in progress towards training by similarity include:

- *Subset inclusion testing.* A key recurring problem is the capacity to identify whether a set of data includes a given signal as a subset. Examples including the testing of a pixel image for containing a near match to a supplied image, and the testing of a vector description of a scene for containing a near match to a supplied object. While testing a given match may not be difficult, the sheer number of possible comparisons is daunting, and grows combinatorially with the environment. . For the near to medium future, this problem will *not* be solved *trivially*⁸¹ – the combinatorial growth outpaces Moore’s Law progress in computational power. Intrinsically parallel approaches via DNA [97] or quantum [98] computing may yield solutions, over and above conceptual insights (heuristics).
- *Data extrapolation.* A second recurring problem is the capacity to decide whether and how a new piece of data fits to a given set, or equivalently, how a given set of data can be extrapolated. The problem is straightforward if an underlying data model is known, but training by similarity seeks to avoid the need to provide such a model⁸². While it is difficult to predict direct progress in the near to medium future, researchers in this area will be boosted by ongoing improvements in computing hardware and scientific software.

It is observed that these and similar problems involve the search for workable processes, and research has not progressed to the implementation and refinement of processes that are known to work. Hence, no prediction can be made as to progress in the near to medium future, other than that workable processes may be found at any time.

⁸¹ In the sense described earlier in Footnote 60.

⁸² A note in passing – for a set of data, *any* model can be fitted to it. However, the fit may be unusably bad or intrinsically wrong. As an example, it is mathematically possible to fit a raw sinusoid to stock market data, even though there is no compelling reason to do so; in contrast, fitting a parabola to data recorded from a free falling body is based on reasoning of underlying physics.

Acknowledgments

The author thanks Kevin Dean for identifying this as a topic of interest, and Mark Ellis for his extensive review and feedback. The work also benefited from the critical insight and review provided by Hugh Barkley, Tony Dekker, Michael Docking, Meredith Hue, Martin Oxenham and Jim Winkles.

References

- 1 C. J. Matheus, M. M. Kokar, and K. Baclawski. "A Core Ontology for Situation Awareness" In *Proceedings of the Sixth International Conference on Information Fusion*, 2003. pp545-552. <http://www1.coe.neu.edu/~kokar/publications/Fusion03.pdf> (at 2004-05-05).
- 2 R S Seymour, V Demczuk, A Filippidis, H T French, A-M Grisogono, W Johnson, D Reid, D Sands, and Y Yue. "Measuring the Benefits of Networked C4ISR Systems – a framework for experimental system development and evaluation". Land Warfare Conference (Sydney, November 2001). 2001.
- 3 Holly A. Yanco and Jill L. Drury. "A Taxonomy for Human-Robot Interaction". In *Proceedings of the AAAI Fall Symposium on Human-Robot Interaction*. AAAI Technical Report FS-02-03. (Falmouth, Massachusetts, November 2002). pp111-119. <http://www.cs.uml.edu/~holly/papers/yanco-drury-taxonomy-fss02.pdf> (at 2004-05-05).
- 4 Henry Hexmoor. "Case Studies of Autonomy". In *Proceedings of FLAIRS 2000*. <http://www.csce.uark.edu/~hexmoor/CV/PUBLICATIONS/CONFERENCES/FLAIRS-00/1HHexmoor00.doc> (at 2004-05-05).
- 5 Albert A Nofi. "Defining and Measuring Shared Situational Awareness". Center for Naval Analyses. CRM D0002895.A1/Final. November 2000. <http://www.thoughtlink.com/publications/DefiningSSA00Abstract.htm> (at 2004-05-05).
- 6 Dale A Lambert. "Situations for Situation Awareness". *FUSION 2001* – 4th International Conference on Information Fusion (Montreal, Canada, 7-10 August 2001). <http://www.inforfusion.org/fusion2001/FUSION2001.htm> (at 2005-09-05).
- 7 Kevin Bonsor. "How Landmines Work". *HowStuffWorks*. <http://people.howstuffworks.com/landmine.htm/printable> (at 2004-04-20).
- 8 "Landmines", Australian Department of Foreign Affairs and Trade. <http://www.dfat.gov.au/landmines/> (Accessed 2004-04-21).
- 9 "Convention on the Prohibition of the use, Stockpiling, Production and Transfer of Anti-Personnel Mines and their Destruction, (Ottawa Landmines Treaty)". 18 September 1997. http://www.un.org/Depts/mine/UNDocs/ban_trty.htm (Accessed 2004-04-21).

- 10 S A Moroz, P W Gower and R Pierson. "Sniper Detection and Counterfire Direction". NRL Review, 1997. <http://www.urf.com/madl/eo/viper/NRLRew.html> (at 2004-04-20).
- 11 "VIPER". Maryland Advanced Development Laboratory, University Research Foundation. <http://www.urf.com/madl/> → Viper (at 2004-04-20).
- 12 "SADS - Sniper Acoustic Detection System ". Rafael Armament Development Authority. <http://www.rafael.co.il/web/rafnew/products/land-sads.htm> (at 2004-05-12).
- 13 "PDCue Counter Sniper System". AAI Defense Systems. <http://www.aaicorp.com/defense/em/pdc.htm> (at 2004-11-09).
- 14 G L Duckworth, D C Gilbert, J E Barger. "Acoustic counter-sniper system". Presented at the SPIE International Symposium on Enabling Technologies for Law Enforcement and Security. Command, Control, Communications, and Intelligence Systems for Law Enforcement, SPIE Proceedings Vol 2938. 19-21 November 1996 (Boston, MA). http://www.bbn.com/sniper/docs/sniper_spie96.pdf (at 2004-04-20).
- 15 Michael Williams. "Remotely Operated Weapon Systems for More Effective Security". Presented in "Safeguards and Security: A New Era workshop", May 2003. <http://www.inmm.org/topics/contents/papers.htm> → ROWS (at 2004-04-20).
- 16 Precision Remotes Inc. <http://precisionremotes.com/> (at 2004-04-20).
- 17 G J de Somer. "Landmines: Weapons That Cause Unnecessary Suffering". *Australian Defence Force Journal*. No 127. Nov/Dec 1997. pp16-24. <http://www.defence.gov.au/publications/dfj/adfj127.pdf> (at 2005-04-04)
- 18 "MK 15 Phalanx Close-In Weapons System (CIWS)". Federation of American Scientists – Military Analysis Network. <http://www.fas.org/man/dod-101/sys/ship/weaps/mk-15.htm> (at 2004-04-20).
- 19 "The Gatling Gun – Phalanx CIWS (Close-In Weapon System)". US Army TACOM R1 – Aircraft Armament and Small Arms Commodity Business Unit. <http://tri.army.mil/LC/CS/csa/aagatlin.htm> (at 2004-04-20).
- 20 Tom Clancy. *Marine: A Guided Tour of a Marine Expeditionary Unit*. HarperCollins Publishers. 1997.
- 21 Jeffrey L Levinson, Randy L Edwards. *Missile Inbound – The Attack on the Stark in the Persian Gulf*. Airlife Publishing, England. 1997.
- 22 "ZSU-23-4 *Shilka* 23MM Antiaircraft Gun". Federation of American Scientists – Military Analysis Network. <http://www.fas.org/man/dod-101/sys/land/row/zsu-23-4.htm> (at 2004-04-21).
- 23 Images downloaded from <http://www.inetres.com/gp/sf/aliens/> (at 2004-04-21). All *Aliens* images Copyright 1986 Twentieth Century Fox Film Corporation.
- 24 "AEGIS Weapon System MK-7". Federation of American Scientists – Military Analysis Network. <http://www.fas.org/man/dod-101/sys/ship/weaps/aegis.htm> (at 2004-04-21).

-
- 25 "AN/SPY-1 Radar". Federation of American Scientists – Military Analysis Network. <http://www.fas.org/man/dod-101/sys/ship/weaps/an-spy-1.htm> (at 2004-04-21).
- 26 "AN/SPG-62 Fire Control Radar". Federation of American Scientists – Military Analysis Network. <http://www.fas.org/man/dod-101/sys/ship/weaps/an-spg-62.htm> (at 2004-04-21).
- 27 "RIM-66 / RIM-67 Standard Missile". Federation of American Scientists – Military Analysis Network. <http://www.fas.org/man/dod-101/sys/missile/sm-2.htm> (at 2004-04-21).
- 28 "Mk 7 Weapon System 'Aegis'; GWS Mk 34; FCS Mk 99". *Jane's Naval Weapon Systems* (Online edition). 02-Jun-2005.
- 29 LTC David Evans (Ret). "Vincennes – A Case Study". *US Naval Institute Proceedings*. Vol.119, Num.8. August 1993. pp49-58.
- 30 Tom Clancy, GEN Chuck Horner. *Every Man a Tiger*. Berkley Publishing Group, 2000.
- 31 Images of USS Hue City (CG66) and USS Vicksburg (CG69) taken by Thoralf Doehring, downloaded from <http://navysite.de/weapons/aegis.htm> (at 2004-04-21).
- 32 "FZG 76". Fiddler's Green. http://www.fiddlersgreen.net/AC/aircraft/V1/v1_info/vi_info.htm (at 2004-04-21).
- 33 "V1 flying bomb". *TheFreeDictionary.com*. <http://encyclopedia.thefreedictionary.com/V1%20Flying%20Bomb> (at 2004-04-21).
- 34 Bill Gunston. *Rockets & Missiles*. Salamander Books, 1979.
- 35 "AGM-84 Harpoon, SLAM [Stand-Off Land Attack Missile]". Federation of American Scientists – Military Analysis Network. <http://www.fas.org/man/dod-101/sys/smart/agm-84.htm> (at 2004-04-21).
- 36 "BGM-109 Tomahawk". Federation of American Scientists – Military Analysis Network. <http://www.fas.org/man/dod-101/sys/smart/bgm-109.htm> (at 2004-04-21).
- 37 Rick Atkinson. *Crusade: The Untold Story of the Persian Gulf War*. Mariner Books, 1994.
- 38 "JASSM". Lockheed Martin Missiles and Fire Control. http://www.missilesandfirecontrol.com/our_products/strikeweapons/JASSM/product-jassm.html (at 2004-04-22).
- 39 "AGM-158 Joint Air to Surface Standoff Missile (JASSM)". Federation of American Scientists – Military Analysis Network. <http://www.fas.org/man/dod-101/sys/smart/jassm.htm> (at 2004-04-22).
- 40 "AGM-158 JASSM – Joint Air-to-Surface Standoff Missile". *The Warfighter's Encyclopedia*. Warfighter Response Center, Naval Air Warfare Center - Weapons Division, United States Navy. http://wrc.chinalake.navy.mil/warfighter_enc/weapons/Develop/jassm/jassm.htm (at 2004-04-22).

-
- 41 Job Advert at Lockheed Martin Corporation, Ref ID 295324. Posted to *CareerBuilder.com* at http://science-biotech.careerbuilder.com/JobSeeker/Jobs/JobDetails.aspx?Job_DID=JZ3Q668R6NjXV0DGCR&cbRecursionCnt=1&cbsid=9fadfc37d1ca46578d81a26ddb10eb93-135917442-xm-2 (at 2004-04-22)
 - 42 2LT Brooke Davis. "B-1B releases first-ever inert JASSM cruise missile at China Lake". Air Force Flight Test Centre. November 21, 2003. http://www.edwards.af.mil/archive/2003/2003-archive-b1b_jassm.html (at 2004-07-26).
 - 43 "Boeing SLAM-ER Successful in Automatic Target Acquisition Launch". Boeing Press Release, 30 April 2001. <http://www.fas.org/man/dod-101/sys/smart/docs/man-sm-agm84-010430.htm> (at 2004-04-22).
 - 44 Christopher Bellamy. *Expert Witness*. Brassey's, 1993.
 - 45 Tom Harris. "How Machine Guns Work". *Howstuffworks*. <http://people.howstuffworks.com/machine-gun.htm/printable> (at 2004-04-23).
 - 46 Eric J Chaisson. *The Hubble Wars*. HarperCollins Publishers. 1994.
 - 47 "[SM3B] Servicing Mission 3B | Hubble Space Telescope". NASA. <http://sm3b.gsfc.nasa.gov> (at 2004-04-23).
 - 48 "[SM3A] Servicing Mission 3A | Hubble Space Telescope". NASA. <http://sm3a.gsfc.nasa.gov> (at 2004-04-23).
 - 49 "The Hubble Project - Servicing Missions". <http://hubble.nasa.gov/servicing-missions/> (at 2004-04-23).
 - 50 "A Reprieve For Hubble?". *CBSNews.com*. Jan.30 2004. <http://www.cbsnews.com/stories/2004/01/16/tech/main593826.shtml> (at 2004-04-23).
 - 51 Robert Roy Britt. "Space Junk: The Stuff Left Behind". *Space.com*. 19 October 2000. http://www.space.com/spacewatch/space_junk.html (at 2004-04-27).
 - 52 "Mars Microrover Telecommunications". NASA Jet Propulsion Laboratory, California Institute of Technology. January 21, 2000. <http://marsprogram.jpl.nasa.gov/MPF/rovercom/rovcom.html> (at 2004-04-27), and subsequent links.
 - 53 "Mars Exploration Rover Mission - Moving around Mars". NASA Jet Propulsion Laboratory, California Institute of Technology. 27 Apr 2004. http://marsrovers.jpl.nasa.gov/mission/tl_surface_nav.html (at 2004-04-27), and subsequent links.
 - 54 Robert A Hall. *Combat Battalion: The Eighth Battalion in Vietnam*. Allen & Unwin. 2000.
 - 55 "The DARPA Grand Challenge Commemorative Program : March 8 - 13 2004". Defense Advanced Research Projects Agency. 2004. <http://www.grandchallenge.com> (at 2004-04-27).

-
- 56 Marsha Walton. "Robots fail to complete Grand Challenge". *CNN.com*. March 14, 2004. <http://edition.cnn.com/2004/TECH/ptech/03/14/darpa.race/> (at 2004-04-27).
- 57 Image downloaded from the DARPA Grand Challenge 2004 Still Photos Gallery, <http://www.grandchallenge.com/gallery/index.html> (at 2004-04-27).
- 58 "Final Data from DARPA Grand Challenge". DARPA Press Release, March 13, 2004. http://www.darpa.mil/grandchallenge/media_news.htm (at 2004-04-27).
- 59 Greg Goebel. "Unmanned Aerial Vehicles". v1.2.0, 01 Jan 03. <http://www.vectorsite.net/twuav3.html> (at 2004-05-03).
- 60 Christopher A Jones. "Unmanned Aerial Vehicles (UAVs) An Assessment Of Historical Operations And Future Possibilities". Research Paper AU/ACSC/0230D/97-03, March 1997. <http://www.fas.org/irp/program/collect/docs/97-0230D.pdf> (at 2004-05-03).
- 61 Bruce Bailey. "Drones In Southeast Asia". http://www.55srwa.org/55_bruce.html (at 2004-05-03).
- 62 "RQ-4A Global Hawk (Tier II+ HAE UAV)". Federation of American Scientists – Military Analysis Network. http://www.fas.org/irp/program/collect/global_hawk.htm (at 2004-05-04).
- 63 "Fact Sheet – Global Hawk". *Air Force Link*. April 03. <http://www.af.mil/factsheets/factsheet.asp?fsID=175> (at 2004-05-04).
- 64 "Global Hawk Completes First Flight". Teledyne Ryan Press Release March 2, 1998. <http://www.fas.org/irp/program/collect/docs/980302-News-Release.htm> (at 2004-05-04).
- 65 Karl Schwarz. "Global Hawk for ELINT Tests at Nordholz Naval Air Base". *Flug Revue*. 1/2004. <http://www.flug-revue.rotor.com/FRheft/FRHeft04/FRH0401/FR0401a.htm> (at 2004-05-04).
- 66 Frank Colucci. "Air Force Refines Training Programs for UAV Operators". *National Defense*. May 2004. <http://www.nationaldefensemagazine.org/article.cfm?Id=1427> (at 2004-05-04).
- 67 David Mulholland. "Global Hawk authorisation lifts UAV industry". *Jane's Defense Weekly*. Vol 40, No 10. 10 September 2003. p40.
- 68 Carlo Kopp. "The UCAV Ascendancy: What are the Problem Issues?". Invited Paper, UAV Australia. ADSC-ADFA/Flight International, February 8-9 2001. <http://www.csse.monash.edu.au/~carlo/mpubs.html> (at 2004-11-09).
- 69 Robert Valdes. "How the Predator UAV Works". *Howstuffworks*. <http://science.howstuffworks.com/predator.htm/printable> (at 2004-05-04).
- 70 Frank Bokulich "Technology update – JSF voice recognition". *Aerospace Engineering*. May 2000. http://www.sae.org/aeromag/techupdate_5-00/03.htm (at 2004-05-04).
- 71 David H Freedman. "Inside the Joint Strike Fighter". *Business 2.0*. February 2002. <http://www.business2.com/b2/web/articles/0,17863,514388,00.html> (at 2004-05-04).

- 72 Nick Cook. "Predator closes sensor-to-shooter gap for USAF". *Jane's Defence Weekly*. Vol.37, Num.7. 13 February 2002. pp28-29.
- 73 "Convention (VIII) relative to the Laying of Automatic Submarine Contact Mines. The Hague, 18 October 1907.". International Humanitarian Law - Treaties & Documents. <http://www.icrc.ch/IHL.nsf/WebList?ReadForm&id=215&t=art> (at 2006-06-13).
- 74 Leon Lederman. *The God Particle*. Bantam Press, London. 1993.
- 75 Phillip F Schewe and Ben Stein. "Physics News Update". *The American Institute of Physics Bulletin of Physics News*. Number 350. December 10 1997. <http://www.aip.org/enews/physnews/1997/physnews.350.htm> (at 2004-05-17).
- 76 Malcolm Tarlton. "Quantum Teleportation (Graphic)". American Institute of Physics. <http://www.aip.org/physnews/graphics/html/teleport.htm> (at 2004-05-17).
- 77 "Space Launch Report: Delta IV Data Sheet". *Space Launch Report*. Last Update: March 6, 2004. <http://www76.pair.com/tjohnson/delta4.html> (at 2004-05-17).
- 78 "U.S. Air Force Procures Boeing Delta IV Launches for EELV Program". Boeing Press Release 98-078, October 1998. http://www.boeing.com/news/releases/1998/news_release_981016a.html (at 2004-05-17).
- 79 "Small Satellite classification". Surrey Satellite Technology Limited. October 14th 1996. <http://www.ee.surrey.ac.uk/SSC/CSER/UOSAT/classification.html> (at 2004-05-17).
- 80 "Australian Participation on the Joint Strike Fighter Program". Australian Government – Department of Defence. <http://www.defence.gov.au/jsf/> (at 2004-05-17).
- 81 Joint Standing Committee on Foreign Affairs, Defence and Trade. "Funding Australia's Defence". Report No 83. Presented out-of-session to the Presiding Officers on 8 May 1998. Tabled in the House of Representatives on 12 May 1998 and in the Senate on 13 May 1998. http://www.aph.gov.au/house/committee/jfadt/def_funding/reportind.htm (at 2004-05-17).
- 82 Barbara Opall-Rome. "From Foot Soldier to Network Node: Israel Seeks Quick, Cheap Application of Personal Technology". *DefenseNews.com*. 21 October 2003.
- 83 Efi Landau. "Cellcom, Mountain Rose don't conflict says former IDF Communications Officer". *Wireless News Update*. http://www.wireless.co.il/Full_print.asp?689 (at 2004-05-18).
- 84 Brendan Nicholson. "Seismic and mobile phone records got Bali bombers". *The Age*. May 12, 2005. <http://www.theage.com.au/news/War-on-Terror/Seismic-and-mobile-phone-records-got-Bali-bombers/2005/05/11/1115585028263.html> (at 2005-08-31).
- 85 "Robotics: at your command: control your own remotely operated vehicle". The Tech Museum of Innovation. 2000. <http://www.thetech.org/robotics/atyourcommand/> (at 2004-05-17).

-
- 86 J D Prothero and H G Hoffman. "Immersive Virtual Environments". Technical Report TR-95-2. Human Interface Technology Laboratory, University of Washington. <http://www.hitl.washington.edu/publications/r-95-5/> (at 2004-05-17).
- 87 Keith Sink. "DirectX 8 and VB Development: Animation Techniques". *Sams*. Jun 21, 2002. <http://www.sampublishing.com/articles/article.asp?p=27309&seqNum=3> (at 2004-05-17).
- 88 Brent D Johnson. "Boeing, Ball Aerospace Collaborate on Satellite Communications System". *Photonics BusinessWorld*. July 2003 Edition. <http://www.photonics.com/spectra/business/XQ/ASP/businessid.889/QX/read.htm> (at 2004-05-18).
- 89 Nick Cook. "Out in front". *Jane's Defence Weekly*. Vol.37, No.3. 16 January 2002. pp24-27.
- 90 Bill Sweetman. "Endurance above all for UAVs". *International Defense Review*. June 01, 2003. p58.
- 91 "Mobile Detection Assessment and Response System (1989-)". SPAWAR Systems Center San Diego. 22 October 2003. <http://www.nosc.mil/robots/land/mdars/mdars.html> (at 2004-05-18).
- 92 Illah R Nourbakhsh. "The failures of a self-reliant robot with no planner: or, What might a planner buy you?". The Robotics Institute, Carnegie Mellon University. <http://www-2.cs.cmu.edu/~illah/SAGE/sagefailures.pdf> (at 2004-05-18).
- 93 Bill Sweetman. "UCAVs spread their wings". *International Defense Review*. May 01, 2001.
- 94 Kevin Bonsor. "How Holographic Memory Will Work". *Howstuffworks*. <http://computer.howstuffworks.com/holographic-memory.htm/printable> (at 2004-05-19).
- 95 "DeltaSphere-3000 3D Scene Digitizer". Product Data Sheet. 3rdTech. v060202. <http://www.3rdtech.com/images/deltadatasheet02v5forweb2.PDF> (at 2004-05-19).
- 96 "943 Series Ultrasonic Distance Sensors". Product Data Sheet. Honeywell. 100464-EN Issue 1 02/04. <http://content.honeywell.com/sensing/prodinfo/ultrasonic/product%20sheets/UDS943dts.pdf> (at 2004-05-19).
- 97 Vincent Kiernan. "DNA-Based Computers Could Race Past Supercomputers, Researchers Predict". *The Chronicle of Higher Education*. Information Technology, Page A23. November 28, 1997. <http://chronicle.com/data/articles.dir/art-44.dir/issue-14.dir/14a02301.htm> (at 2004-05-19).
- 98 Andrew Steane, "Quantum Computing". *Reports on Progress in Physics*. Vol 61, 1998. pp117-173. <http://www.qubit.org/library/intros/compSteane/qcintro.html> (at 2004-05-19).
- 99 "IBM Grid Computing". IBM. <http://www-1.ibm.com/grid/> (at 2004-05-21).
- 100 Bill Sweetman. "Naval UAV concepts in flux". *Jane's Navy International*. April 01, 2002.

101 Martin Van Creveld. *Command in War*. Harvard University Press. 1987.

102 Mica R Endsley. "Situation Awareness, Automation & Free Flight". Air Traffic Management Research & Development (Saclay, France, June 17-20, 1997). <http://atm-seminar-97.eurocontrol.fr/> (at 2005-09-05).

DISTRIBUTION LIST

The Generation of Situational Awareness within Autonomous Systems –
A Near to Mid term Study – Analysis

Patrick Chisan Hew

AUSTRALIA

DEFENCE ORGANISATION	No. of copies
Task Sponsor	
Director Future Warfighting	1 Printed
S&T Program	
Chief Defence Scientist	1
Deputy Chief Defence Scientist Policy	1
AS Science Corporate Management	1
Director General Science Policy Development	1
Counsellor Defence Science, London	Doc Data Sheet
Counsellor Defence Science, Washington	Doc Data Sheet
Scientific Adviser to MRDC, Thailand	Doc Data Sheet
Scientific Adviser Joint	1
Navy Scientific Adviser	Doc Data Sheet & Dist List
Scientific Adviser – Army	Doc Data Sheet & Dist List
full copy for an Air Force task or a topic relevant to Air Force. Otherwise Doc Data Sht & Dist List only	Doc Data Sheet & Dist List
Scientific Adviser to the DMO	Doc Data Sheet & Dist List
Chief of Defence Systems Analysis Division	Doc Data Sheet & Dist List
Research Leader Strategy and Concepts	1
Research Leader Maritime Operations Research	1
Research Leader Operations Research (Land Operations Division)	1
Research Leader Air Operations Research	1

Hugh Barkley	1 Printed
Kevin Dean	1 Printed
Roland Keir	1 Printed
Mark Ellis	1 Printed
Patrick Hew	1 Printed
DSTO Library and Archives	
Library Edinburgh	1 Printed
Defence Archives	1 Printed
Library Canberra	Doc Data Sheet
Capability Development Group	
Director General Capability and Plans	Doc Data Sheet
Assistant Secretary Investment Analysis	Doc Data Sheet
Director Capability Plans and Programming	Doc Data Sheet
Chief Information Officer Group	
Director General Australian Defence Simulation Office	Doc Data Sheet
AS Information Strategy and Futures	Doc Data Sheet
Director General Information Services	Doc Data Sheet
Strategy Group	
Assistant Secretary Strategic Planning	Doc Data Sheet
Assistant Secretary Governance and Counter-Proliferation	Doc Data Sheet
Navy	
Maritime Operational Analysis Centre, Building 89/90 Garden Island Sydney NSW	Doc Data Sht & Dist List
Deputy Director (Operations)	
Deputy Director (Analysis)	
Director General Navy Capability, Performance and Plans, Navy Headquarters	Doc Data Sheet
Director General Navy Strategic Policy and Futures, Navy Headquarters	Doc Data Sheet
Air Force	
SO (Science) - Headquarters Air Combat Group, RAAF Base, Williamstown NSW 2314	Doc Data Sht & Exec Summary
Army	
ABCA National Standardisation Officer	e-mailed Doc Data Sheet
J86 (TCS GROUP), DJFHQ	Doc Data Sheet
SO (Science) - Land Headquarters (LHQ), Victoria Barracks NSW	Doc Data Sht &

SO (Science), Deployable Joint Force Headquarters (DJFHQ) (L), Enoggera QLD	Exec Summary Doc Data Sheet
---	--------------------------------

Joint Operations Command

Director General Joint Operations	Doc Data Sheet
Chief of Staff Headquarters Joint Operations Command	Doc Data Sheet
Commandant ADF Warfare Centre	Doc Data Sheet
Director General Strategic Logistics	Doc Data Sheet
COS Australian Defence College	Doc Data Sheet

Intelligence and Security Group

AS Concepts, Capability and Resources	1
DGSTA , Defence Intelligence Organisation	1
Manager, Information Centre, Defence Intelligence Organisation	1
Director Advanced Capabilities	Doc Data Sheet

Defence Materiel Organisation

Deputy CEO	Doc Data Sheet
Head Aerospace Systems Division	Doc Data Sheet
Head Maritime Systems Division	Doc Data Sheet
Program Manager Air Warfare Destroyer	Doc Data Sheet
Guided Weapon & Explosive Ordnance Branch (GWEO)	Doc Data Sheet
CDR Joint Logistics Command	Doc Data Sheet

OTHER ORGANISATIONS

National Library of Australia	1
NASA (Canberra)	1
Library of New South Wales	1

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy

Library	1
Head of Aerospace and Mechanical Engineering	1
Hargrave Library, Monash University	Doc Data Sheet

OUTSIDE AUSTRALIA

INTERNATIONAL DEFENCE INFORMATION CENTRES

US Defense Technical Information Center	1
UK Dstl Knowledge Services	1

Canada Defence Research Directorate R&D Knowledge & Information Management (DRDKIM)	1
NZ Defence Information Centre	1

ABSTRACTING AND INFORMATION ORGANISATIONS

Library, Chemical Abstracts Reference Service	1
Engineering Societies Library, US	1
Materials Information, Cambridge Scientific Abstracts, US	1
Documents Librarian, The Center for Research Libraries, US	1

INFORMATION EXCHANGE AGREEMENT PARTNERS

National Aerospace Laboratory, Japan	1
National Aerospace Laboratory, Netherlands	1

SPARES	5 Printed
--------	-----------

Total number of copies: 40 Printed: 13 PDF: 27

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA					
				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)	
2. TITLE The Generation of Situational Awareness within Autonomous Systems – A Near to Mid term Study – Analysis			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)"		
4. AUTHOR(S) Patrick Chisan Hew			5. CORPORATE AUTHOR DSTO Information Sciences Laboratory PO Box 1500 Edinburgh South Australia 5111 Australia		
6a. DSTO NUMBER DSTO-TR-1896		6b. AR NUMBER AR-013-702		7. DOCUMENT DATE July 2006	
8. FILE NUMBER 2005/1105607/1		9. TASK NUMBER STR 02/211		10. TASK SPONSOR DFW	
				11. NO. OF PAGES 6+80	
				12. NO. OF REFERENCES 102	
13. URL on the World Wide Web http://www.dsto.defence.gov.au/corporate/reports/DSTO-TR-1896.pdf				14. RELEASE AUTHORITY Chief, Defence Systems Analysis Division	
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i>					
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111					
16. DELIBERATE ANNOUNCEMENT No Limitations					
17. CITATION IN OTHER DOCUMENTS Yes					
18. DEFTEST DESCRIPTORS automation, futures analysis, robotics, situational awareness.					
19. ABSTRACT This study aims to clarify and capture the nature of electronic situational awareness and its interface with electro/mechanical systems. It argues that "autonomous situation awareness" is about the sufficiency of awareness for autonomy in the situation at hand. The approach is calibrated through historical case studies, and the study then considers the potential from near to mid term technology.					