



INTELLIGENT SYSTEMS REFERENCE LIBRARY
Volume 3

Anthony Finn
Steve Scheduling

Developments and Challenges for Autonomous Unmanned Vehicles

A Compendium

 Springer

Anthony Finn and Steve Scheduling

Developments and Challenges for Autonomous Unmanned Vehicles

Intelligent Systems Reference Library, Volume 3

Editors-in-Chief

Prof. Janusz Kacprzyk
Systems Research Institute
Polish Academy of Sciences
ul. Newelska 6
01-447 Warsaw
Poland
E-mail: kacprzyk@ibspan.waw.pl

Prof. Lakhmi C. Jain
University of South Australia
Adelaide
Mawson Lakes Campus
South Australia
Australia
E-mail: Lakhmi.jain@unisa.edu.au

Further volumes of this series can be found on our homepage: springer.com

Vol. 1. Christine L. Mumford and Lakhmi C. Jain (Eds.)
Computational Intelligence, 2009
ISBN 978-3-642-01798-8

Vol. 2. Yuehui Chen and Ajith Abraham
Tree-Structure Based Hybrid Computational Intelligence, 2009
ISBN 978-3-642-04738-1

Vol. 3. Anthony Finn and Steve Scheduling
Developments and Challenges for Autonomous Unmanned Vehicles, 2010
ISBN 978-3-642-10703-0

Anthony Finn and Steve Scheduling

Developments and Challenges for Autonomous Unmanned Vehicles

A Compendium

Anthony Finn

Defence Science & Technology Organisation
Department of Defence
Edinburgh, SA 5111
Australia
E-mail: anthony.finn@dsto.defence.gov.au

Steve Scheduling

Australian Centre for Field Robotics
University of Sydney
Sydney, NSW 2006
Australia
E-mail: scheduling@acfr.usyd.edu.au

ISBN 978-3-642-10703-0

e-ISBN 978-3-642-10704-7

DOI 10.1007/978-3-642-10704-7

Intelligent Systems Reference Library

ISSN 1868-4394

Library of Congress Control Number: 2010920022

© 2010 Springer-Verlag Berlin Heidelberg

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typeset & Cover Design: Scientific Publishing Services Pvt. Ltd., Chennai, India.

Printed in acid-free paper

9 8 7 6 5 4 3 2 1

springer.com

FOR LISA

Foreword

The past decade has seen a sea change in both the vision and practical deployment of robotics and autonomous systems in defence applications. This has been driven both by increasing technical capabilities and by the imagining of what can be done by many researchers and application experts. The visible successes of unmanned air vehicles (UAVs), as brought to you daily on your evening TV news, of bomb disposal and sentry robots in regular use in urban conflicts, and of underwater mine-hunting robot vehicles deployed as standard features on ships, are all changing the way we think about defence and tactical conflict. Our future thoughts on what robots and autonomous systems might be capable of in the future are also expanding and maturing with many active projects developing unmanned ground vehicles (UGVs), unmanned combat aircraft (UCAVs) and conceiving entirely new “systems of autonomous systems”. These future thoughts increasingly rely on an extended idea of autonomy, in which increasing amounts of sensor data processing and decision making are undertaken by the robot itself, reducing human intervention and replacing the many roles of current unmanned system operators.

These rapid changes in technology and thinking are, as yet, not backed up by corresponding changes in our concept of operations, by our understanding of the limitations, as well as the advantages, conferred by this technology, or by our study of the legal and ethical consequences of deploying autonomous systems in a human conflict. Indeed, we are still as children in a toy shop.

This book makes a serious attempt to address these challenging issues in a manner which is at once accessible, technically sound, and thought provoking for practitioners and decision makers in the defence realm. The book provides an excellent overview of technologies, their history, capabilities, limitations and future development; in a manner which aims to assist in understanding of what robotics can do for defence. The book also explores some of the issues surrounding the likely changes in operational concepts, the effects this may have on areas such as policy and training, and the manner in which the performance of unmanned systems might be evaluated. Uniquely, and bravely, this book discusses some of the legal issues surrounding the use of robots and autonomous systems in defence. These are complex issues which will likely come to dominate the manner in which robots are deployed and operated as their technical competency matures. This book provides a serious analysis of these legal and ethical issues.

For the defence policy and decision maker, this is a “must-read” book which brings together an important technology summary with a considered analysis of

future doctrinal, legal and ethical issues in unmanned and autonomous systems. For research engineers and developers of robotics, this book provides a unique perspective on the implications and consequences of our craft; connecting what we do to the deployment and use of the technology in current and future defence systems.

Professor Hugh Durrant-Whyte
ARC Federation Fellow
Director, Australian Centre for Field Robotics
University of Sydney, Sydney NSW, Australia

Preface

It is widely anticipated that autonomous vehicles will have a transformational impact on military forces and will play a key role in many future force structures. As a result, many tasks have already been identified that unmanned systems could undertake more readily than humans. However, for this to occur, such systems will need to be agile, versatile, persistent, reliable, survivable and lethal. This will require many of the vehicles 'cognitive' or higher order functions to be more fully developed, whereas to date only the 'component' or physical functions have been successfully automated and deployed.

Such intelligent automation clearly presents a range of challenges for researchers, users, developers, planners, and procurers of such systems, many not widely appreciated or shared outside their own disciplines. Having worked together for a number of years the authors noted a mutual misunderstanding between those who develop and innovate and those who plan and use: scientists like to innovate in a vacuum and provide abstract solutions, engineers like to build to specification and deliver useful applications, planners like to know what is going to happen next and what the best decision is right now, and users – like the rest of us – just want their current jobs made easier.

Some of the misunderstanding undoubtedly stems from technology aspirations that are unrealistic because they are unbounded by the extraordinarily tough requirements of persistent autonomy in dynamic military environments; some because the potential of the technology is misunderstood either relative to opportunity or to its application domain. As a result, we undertook to write a book that would communicate across the cultural boundaries by articulating what we see as some of the key, enduring, and unexpressed challenges in the area. We do this because, like many others, we recognise the potential of the technology and the rate at which this disruptive innovation is evolving. However, having both been involved in the development of real systems for many years we also appreciate the complexity of the environment and the many challenges this presents.

The book draws upon a broad range of others' work with a view to providing a product that (we hope) is greater than the sum of its parts. The discussion is intentionally approached from the perspective of improving understanding rather than providing solutions or drawing firm conclusions. Consequently, researchers reading this book with the hope of uncovering some novel theory or approach to automating an unmanned vehicle will be as disappointed as the capability planner who anticipates a catalogue of technical risks and feasibility options against his

favoured list of component technologies and potential applications. Nevertheless, it is hoped that both will at least learn something of the other's world and that progress will ensue as a result.

In addition to their military applications, however, robotic and unmanned systems are also applied to an array of commercial tasks. As a result, it is recognised that many of the challenges discussed in this book are not unique to military systems, but are generally applicable to civilian applications and systems also. While the areas of commonality are not explicitly highlighted, it is hoped that they are readily identifiable so that benefit may be brought to the widest possible audience.

Acknowledgements

We are very grateful to a number of colleagues: Brig (Ret.) Steve Quinn, Chief of Land Operations and Dr Warren Harch, Deputy Chief Defence Scientist (Information and Weapons Systems), both from the Defence Science & Technology Organisation; Professor Hugh Durrant-Whyte, Director of the Australian Centre for Field Robotics at the University of Sydney; and, Professor Lakhmi Jain, Director KES at the University of South Australia. All have kindly contributed through discussion or assisted in the publication of this work. We are also grateful to our many colleagues at DSTO and the University of Sydney who have provided valuable advice over many years.

Contents

1	Introduction.....	1
2	Background.....	5
	2.1 Early UVS.....	5
	2.2 Modern UVS.....	15
	2.3 Looking Forward	25
3	Autonomous UVS.....	35
	3.1 UVS Components	35
	3.2 UVS Architectures.....	39
	3.3 Human-UVS Interaction.....	41
	3.3.1 Human-Machine Interfaces.....	49
	3.3.2 Supervisory Workloads.....	50
4	UVS Technology Issues.....	55
	4.1 Technology Challenges	55
	4.2 Contextual Decision-Making.....	59
	4.2.1 Planning in Dynamic Environments	65
	4.3 Verification and Validation	68
	4.3.1 Trust and Reliability in UVS	69
	4.3.2 Systems V&V for Autonomous UVS	72
	4.3.3 Simulation-Based V&V	78
	4.3.4 Health and Usage Monitoring.....	80
	4.4 Multi-vehicle Systems	82
	4.4.1 Multi-UAV ISTAR Example.....	87
	4.4.2 Multi-UVS Coordination	91
	4.4.3 Autonomous Multi-UVS Task Allocation	94
	4.4.4 Multi-UVS Navigation, Localisation and Mapping.....	99
	4.4.5 Capability and Systems Integration	102

5	Force-Integration of UVS.....	105
5.1	Capability Challenges.....	105
5.1.1	Policy Considerations.....	108
5.1.2	Training and Maintenance.....	109
5.2	A Systems Approach to UVS Analysis.....	112
5.3	Simulating UVS.....	118
5.3.1	Interactive Simulation.....	124
5.4	Measuring the Effectiveness of UVS.....	128
5.4.1	Measuring the Force Effectiveness of UVS.....	129
5.4.2	Measuring the Systems Performance of UVS.....	138
5.4.3	Measuring UVS Performance.....	145
6	Legal Issues for UVS.....	155
6.1	Legal Issues for UVS Platforms.....	156
6.1.1	Maritime Vehicles.....	157
6.1.2	The Legal Status of UVMs.....	157
6.1.3	Operations in Ocean Areas.....	159
6.1.4	Navigation Environment.....	159
6.1.5	Air Vehicles.....	161
6.1.6	Ground Vehicles.....	164
6.2	UVS and Tort Law.....	165
6.3	UVS and the Law of Armed Conflict.....	168
6.3.1	The Law of Armed Conflict.....	171
6.3.2	UVS and Jus ad Bellum.....	176
6.3.3	Autonomous Weapons.....	178
6.4	Accountability and Liability.....	180
6.4.1	Legal Personality for UVS.....	184
6.4.2	UVS as Tools.....	185
6.4.3	Human Intervention.....	186
6.4.4	Technical Equivalence.....	186
6.4.5	UVS as Agents.....	187
6.4.6	Separate Legal Personality.....	188
6.4.7	Impediments to Artificial Agency.....	189
6.4.8	Technical Impediments to Agency.....	191
6.4.9	Is the UVS the Real Decision-Maker?.....	192
6.5	Ethical Control of UVS.....	194
6.5.1	Could UVS Ethically Out-Perform Human.....	196
6.5.2	Implementation Issues.....	198
6.5.3	Architectural Considerations.....	202
	Concluding Remark.....	209
	References.....	213
	Acronyms.....	229

Table of Figures

Fig. 2.1	When will TRL 6 be reached for the key functional areas of UVS?	27
Fig. 3.1	The key functional relationships for an autonomous UVS.....	36
Fig. 3.2	Levels of Autonomy (after [202])	42
Fig. 3.3	Levels of Autonomy (after Sheridan & Verplank, 1978 [250]).....	43
Fig. 3.4	Intra-Vehicle Levels of Autonomy (after Cummings, 2004 [74]).....	44
Fig. 3.5	The Duality of Autonomy	45
Fig. 3.6	Relationship between degree of autonomy and mission complexity Adapted from [202]	47
Fig. 3.7	Relationship between system complexity & mission autonomy Adapted from [202].....	48
Fig. 4.1	Approach of linking perception to action through cognition (Adapted from [58])	62
Fig. 4.2	Alternative decomposition of the perception-actuation problem (Adapted from [58]).....	63
Fig. 5.1	Systems Engineering Approach to UVS-Force Integration [32]	116
Fig. 5.2	System model for determining UVS effectiveness.....	117
Fig. 5.3	Overview of Simulation Methods (adapted from [225])	119
Fig. 5.4	Value of Simulation Techniques	122
Fig. 5.5	Tendencies of characteristics of measures (adapted from [203])	131
Fig. 5.6	Validity criteria of measures	131
Fig. 5.7	Reliability criteria of measures.....	132
Fig. 5.8	Scenario Framework (adapted from [203])	135
Fig. 5.9	Examples of MOFE and associated MOE's	137
Fig. 5.10	Typical Measures of Effectiveness for Human-UVS Teaming	141

Chapter 1

Introduction

“The whole art of war consists in getting at what is on the other side of the hill”
The Duke of Wellington, 1769-1852

The intrinsic demand for automated, robotic and unmanned systems is largely driven by applications that are inherently repetitive, unpleasant, or dangerous. At present, these tasks typically include agricultural, container handling, intelligent transportation (repetitive), scientific exploration, mining, waste management, (unpleasant), search and rescue, fire-fighting, and military applications (dangerous). Additionally, as many of these applications necessitate the employment of vehicles that are relatively expensive, the additional cost of the automation components and integration is often modest relative to the gains made by better use of the platform. This tends to auger well when establishing robust use cases for such systems, particularly in the commercial arena. For military applications there are a number of other drivers:

- Autonomous UVS have the capacity to extend the reach and access of operations, thereby reducing risk to warfighting personnel;
- If appropriately networked to each other and to higher value manned platforms, they can significantly increase coherency of effort and operational tempo, thereby potentially offering much higher operational effectiveness and increased capability across the battlespace;
- In the longer term, they could reduce the cost of acquisition and operations; and
- They are one of the few areas of current technology that legitimately qualifies as having a potentially revolutionary impact on military operations

In 2001 the US Congress was sufficiently persuaded by the military potential of these systems that it directed its Department of Defense (DoD) that one third of all operational deep strike force aircraft must be unmanned by 2010 and one third of its

operational ground combat vehicles must be unmanned by 2015.¹ While this may not be achieved in terms of fully autonomous UVS, the deadlines have applied considerable pressure on the US military to introduce large numbers of tele-operated or semi-autonomous UVS into capability, and many nations will almost certainly follow their lead. As technology improves the level of autonomy will increase.

As a result, autonomous and Unmanned Vehicle Systems (UVS) now play a key role in the modern battlespace and would have allowed the Iron Duke to spend much more of his career pondering other matters. Furthermore, their contribution in recent conflicts has led to public recognition of their utility by the most senior leaders of several nations.

Broadly, UVS fall into classes defined by their operating environment (air, land and sea), autonomy and size, which in turn defines weight, payload configuration, endurance, mission, etc. They operate in all battlespace environments and are usually referred to as: Unmanned (or Uninhabited) Aerial Vehicles (UAV), Unmanned Maritime Vehicles (UMV)² or Unmanned Ground Vehicles (UGV). Their combat variants usually have a “C” inserted between the Unmanned and the environmental descriptor (i.e. “UCAV” for Unmanned Combat Aerial Vehicle).

The military benefits of these systems are proven repeatedly on an almost daily basis with several thousand of them currently in service around the world. The fielded systems provide critical support to operations and comprise a diverse mix of Commercial off the Shelf (COTS) and fielded prototypes that vary in size from a few pounds to tens of tons. They are employed in a range of roles, including rapid environmental assessment (REA), improvised explosive device (IED) detection and defeat, explosive ordnance disposal (EOD), countermining, force protection, obstacle clearance, battle damage assessment (BDA), electronic warfare (EW) and intelligence, surveillance, and reconnaissance (ISR).

There are also a number of key areas of capability where UVS can supplement warfighter activity in ways that they currently do not: logistics, medicine, engineering, security, and maintenance. Similarly, for the foreseeable future most western defence forces will continue to be designed around highly trained, well equipped personnel, selected for their resourcefulness and ability to improvise [13]; but that are dependent upon sophisticated high value assets that act as force multipliers. For survivability, these high-value platforms have historically depended upon sensors and data links to maintain situational awareness, with Electronic Warfare (EW) self-protection often used as a last resort. However, increasingly capable air, land, and sea combat systems now use a combination of sensors networked together to provide an adversary with the capacity to precisely track and target these high value assets at long ranges. As a result, the projection

¹ Section 220(a) (2) of the Floyd D. Spence National Defense Authorization Act for Fiscal Year 2001 (as enacted into law by Public Law 106–398; 114 Statute 1654A–38), mandates that “It shall be a goal of the Armed Forces to achieve the fielding of unmanned, remotely controlled technology such that

- By 2010, one third of the operational deep-strike force aircraft fleet are unmanned, and
- By 2015, one-third of the operational ground combat vehicles are unmanned.”

² Comprising Unmanned Underwater Vehicles (UUVs) and Unmanned Surface Vehicles (USVs).

of spectral and physical force in such a way as to have an effect beyond the immediate self-protection of the platforms will almost certainly become vital to the survival of these high value (and likely manned) assets. This is also a role ideally suited to autonomous UVS.³ Irrespective of their domain or application, the general requirements of UVS are [205]:

- Persistence, low cost, stealth, and ready deploy/retrieve-ability;
- The capacity to detect, locate, track, identify and engage targets autonomously;
- The ability to gather, disseminate and act on several types of information;
- That they are networked together and to the higher-value, manned assets;
- That the individual platform and sensor elements can self-organise; and,
- That they do not impose significant risk or burden upon the operators.

Clearly, application of UVS will not result in ‘bloodless battlefields’ or circumstances in which we ‘press the button and fight the war’ any time soon. Nor will they provide solutions to every capability challenge faced by today’s defence force planners. They will, however, be exploitable in terms of their capacity to carry out dangerous, repetitive and mechanically-oriented tasks currently undertaken routinely by warfighters thereby freeing them for other missions. Furthermore, UVS are particularly well-suited to well-structured and uncluttered environments and tasks and applications where system or mission failure has little or no impact on humans.

In the last few decades while much progress has been made towards these aspirational goals, realisation of the somewhat utopian vision of UVS working persistently and seamlessly together and with manned vehicles in adversarial environments still requires significant scientific advances to be made in a range of areas. As a result we have organised this book into a number of sections:

- The **Background** chapter represents a tapestry of UVS capability from its earliest, faltering steps through to a projection of what is likely within the planning cycles of most defence forces. The objective is simply to provide context for the reader.
- The next chapter, **Autonomous UVS**, describes the functional elements of a UVS, some architectural considerations and human-UVS interaction. The chapter acts as ‘grounding’ for the discussion on the technology, force integration and legal challenges that follow.
- In the chapter on **Technology Challenges** the requirements of next-generation, autonomous UVS operating in dynamic and complex military

³ Prior to the Global Financial Crisis of 2008-2009, spending on UVS in the US alone was predicted to be in excess of US\$33B over the period 2008-2016 (i.e. around US\$3.5B annually), with the expenditure profile divided roughly 50%-40%-10% between the procurement, research, and operations and maintenance budgets, respectively [112]. The expenditure on UAVs is expected to dominate with UGVs second (~ US\$500m annually) and UUVs rising from about US\$20M to US\$100M annually over the period. The 2007 Unmanned Systems Roadmap [278], 2009 Unmanned Systems Integrated Roadmap [279] and USAF Unmanned Aircraft Systems Flight Plan 2009-2047 [282], which more than any other recent documents provide an insight into how this money will be spent and what steps will be taken to realise the future potential of UVS, go further than the equivalent documents of other nations by anticipating that fully autonomous swarms of UVS will be achieved by 2015.

environments are considered. The challenges examined pertain more to higher-order tasks such as contextual decision-making, planning and re-planning in dynamic environments, verification and validation, and trust and reliability in human-UVS relationships than to the technological or systems integration challenges that exist for the functional elements of a UVS.

- In **Force Integration** attention is drawn to the fact that the value proposition for many military UVS has yet to be formally quantified. As a result, this chapter discusses some of the challenges for inserting autonomous UVS into existing force structures and presents a methodology that allows them to be assessed in terms of their mission, systems and technological performance based on a range of simulation and assessment techniques.
- In the final chapter on **Legal Issues** the symbiotic relationship between advanced technology, its capability exploitation, and the legal framework in which it must exist is explored. The chapter discusses some of the conundrums that result from UVS having the potential to exercise their own ‘judgement’ in regard to the lethal prosecution of targets and having to operate in environments shared by people, property and other vehicles.

Many of these challenges are equally applicable to commercial or non-military unmanned vehicle systems. For example, irrespective of application – and regardless of whether this is commercial or military – the payload-mission combination defines the overall system requirements. This does not usually impact the key functional elements of the system, which inevitably include mobility, localisation, navigation, and planning, irrespective of the mission. Furthermore, a degree of ‘cognitive’ or higher level planning capability relative to task complexity often translates to mission flexibility. In this book, the challenges are interpreted solely in terms of military missions, and their interpretation in the context of commercial application is left as an exercise for the reader.

Similarly, it is also recognised that both commercial and military vehicles will ultimately occupy environments shared by other users, which will in turn require them to obey the established ‘rules of the road.’ However, even though such technology and legal challenges are equally relevant to both military and commercial users and applications, they are not explicitly discussed in this book in the context of the many potential commercial uses. Additionally, even though a number of advances in military automation and unmanned systems are likely to extend beyond the commercial applications into the area of manned transport,⁴ these ‘dual use’ technology issues are identified and discussed solely in terms of their suitability to military unmanned systems.

⁴ For example, technologies that allow UAVs to detect, see, and avoid other aircraft are likely to benefit pilots of manned aircraft who are otherwise occupied or have heavy workloads. Similarly, while it is technically feasible for manned aircraft to be flown for short periods in nuclear, biological or chemical contaminated environments and to protect pilots from the blinding effects of high powered lasers, technologies that allow such aircraft to operate in such adverse environments and land safely after such exposures are inherently advantageous (lasers that have the potential to blind pilots are considered illegal, but so long as there are adversaries that do not respect the Laws of Armed Conflict such weapons remain a potential threat to manned aviation).

Chapter 2

Background

What we refer to as automation changes over time; when functions previously carried out by humans are carried out reliably by technology they become ‘machine operation’ rather than automation. On the other hand, certain functions preserve their automation label as they continue to be carried out by both humans and machines. The aim of this chapter is therefore to provide the reader with a historical perspective on the development of UVS and a view of what might become. This frames likely technological developments within the context of autonomous military vehicles so that the reader may appreciate the degree of disruptive innovation and capability promise becoming available within the financial planning cycles of most defence forces.

2.1 Early UVS

The desire to extend one’s reach while simultaneously reducing the risk to one’s own forces has existed as an aspiration for warfighters around world and throughout history. Moreover, like many other inventions, the earliest known descriptions of UVS appear to pre-date their development by many centuries. For instance, around 400 BC the Greek mathematician, Archytas of Tarentum, speculated about the possibility of robot birds propelled by steam and in 1275 a Syrian engineer, Hassan al-Rammah, showed a rocket-propelled device skimming across the sea surface, pre-dating the self-propelled torpedo by about six hundred years. Moreover, the notion of an autonomously controlled aircraft relying on television to ‘see’ was first described by the science fiction writer Hugo Gernsback in 1924.⁵

A detailed history of UVS would occupy a volume in its own right. Nevertheless, in their most simple form, and depending upon your point of view, the first ones date back either to around 400 BC when the Chinese invented the kite and attached bronze mirrors to them; to the Romans, who were the first to use fire ships⁶; or to August 22nd 1849 when the Austrians used unmanned balloons to attack the Italian city of Venice. Having loaded the balloons with explosives, they were floated over the city and set on fire using an electrical current discharged through long copper

⁵ Prophetically, Gernsback also predicted that most of those reading the article would live to see the reality.

⁶ Fire ships were filled with flammable (and often explosive) material and deliberately set on fire. They were then either steered or allowed to drift into an enemy fleet. Records indicate that control of the ships once underway, however, was somewhat haphazard as the wind was prone to change direction after the ships had been set alight.

wires trailing back to the Austrian forces. The attempt was less than successful with several balloons seemingly blown back over the Austrian lines.

In 1861 during the American Civil War, and not long after the Austrian UAV program presumably came to an end, Professor Thaddeus Lowe placed cameras onboard tethered balloons to gather intelligence of Confederate forces and their positions. It is not known how these cameras were controlled, but it seems to have been more militarily useful to risk placing men onboard the balloons as this was still common practice some fifty years later during World War I. However, both the Union and Confederate forces appear to have attempted to float and detonate explosive-laden balloons over their adversaries.

A more recognisable form of UVS, and the earliest vehicle with self-propulsion and onboard guidance and control was named after the Torpedo fish, which is an electric ray capable of delivering a stunning shock to its prey⁷. The torpedo was primarily developed by a British engineer, Robert Whitehead. In the 1860's, Whitehead teamed up with his 12-year old son, a technician named Annibale Ploech and Giovanni Luppis, an engineer who had recently retired from the Austrian navy. Together they demonstrated the first self-propelled torpedo to the Austrian Navy on 21st December 1866. By 1870 they had demonstrated a torpedo that could travel at 7knots and hit a target at 700yards.




Robert Whitehead with one of his early torpedos (circa 1870)

⁷ The earliest torpedos (the phrase was coined in about 1800) were in fact explosive devices such as mines and booby traps that were hidden, allowed to drift or steered using early submarines, the first of which was used in the Battle of Connecticut in 1775.

Unfortunately, the technical success and capability promise did not stop the company that produced the torpedo from going bankrupt. Nevertheless, in 1875 Whitehead transformed the remnants of that company into a private venture and made a success of it, eventually selling it on to Vickers Ltd. By 1881 more than 1500 torpedos were in the military inventories of a dozen or so countries. The key innovations were the self-regulating device that kept the torpedo at a pre-set depth and the propulsion mechanism which was provided by a small reciprocating engine run by compressed air. In the late 1890's Whitehead purchased a gyroscope from an Austrian inventor, Ludwig Obry, which enabled the device to be directionally stabilised.

If we distinguish weapons from vehicles, the first remote control vehicle appears to have been demonstrated publicly by Nikola Tesla in 1898 at an exhibition in Madison Square Garden (and covered in US patent number 613,809). He demonstrated a small remote control boat by transmitting specific frequencies to tuned circuits that controlled motors in the boat (although Tesla pretended the boat was capable of receiving verbal commands from the audience). A few years later in 1904 the Englishman, Jack Kitchen, demonstrated a remote control Windermere steam launch and in 1906 a Spanish mathematician, Leonardo Torres y Quevedo, demonstrated the remote control of a small ship before the King of Spain and a large crowd, having demonstrated "telekino" (remote control) some three years earlier.



Nikola Tesla's remote control tele-automaton (1898)


The first pilot-less aircraft was probably developed by Charles Kettering⁸ and was known as the Kettering Aerial Torpedo (or the Kettering “Bug”). It first flew in about 1915, although rather like the early torpedos its targeting, guidance and control systems were very primitive and had to be set prior to launch. A small gyroscopic autopilot designed by Elmer Sperry guided the aircraft to its designated target using a pneumatic/vacuum and electric control system; a barometric altimeter maintained altitude. A mechanical device that counted the number of engine revolutions determined distance travelled, wind speed and direction being accounted for manually prior to take-off. When the target was reached the engine shut down and bolts attaching the wings to the fuselage retracted. As the wings fell off, the “Bug” entered a ballistic trajectory and the impact of the aircraft detonated the 80kg of explosive onboard. The UAV was intended to hit a target such as a ship at a range of about 50miles, but was considered insufficiently accurate for practical military use and was never used in anger.



The Kettering Aerial Torpedo or Kettering ‘Bug’ (1915)

⁸ Assisted by Elmer Sperry, Orville Wright and Robert Millikan.

The first remotely controlled aircraft is believed to be the Ruston Proctor AT⁹ developed by an Englishman, “Professor” Archibald Low around the time of the Kettering Bug. The Ruston Proctor AT was launched from the back of a lorry using compressed air. The first full military demonstration of the system, before 30-40 allied generals, took place on 21st March 1917 at Upavon Central Flying School, near Salisbury Plain, England. Although the aircraft crashed due to engine failure, the team were able to demonstrate controlled flight and follow-on development was authorised. Following the end of the war, however, the promise of RC aircraft appears not to have been pursued very vigorously by the British.¹⁰ Nevertheless, by 1927 they had developed LARYNX (Long Range Gun with Lynx Engine), which was an early anti-ship cruise missile and in 1931 they developed the Fairey Queen an RC target version of the Fairey III F flying boat, which is usually considered to be the first reliable remote control aircraft.



LARYNX, Long Range Gun with Lynx Engine (1927)

During the period between the two World Wars, the US appears to have been rather keener on developing remote control aircraft than their British counterparts, putting obsolete manned aircraft to work as target drones. However, the concept was not fully developed, suffering the fate of many R&D programs: cancellation due to lack of funds. As a consequence, the Kettering Bug’s potential – even after it had been provided with a 200 mile range, remote control and improved guidance – never seems to have been exploited practically, although it was briefly considered as a serious candidate for long range bombing missions by the allies during WWII.


In 1944, however, the US Eighth Air Force did convert around 25 B-17 bombers that had outlived their operational usefulness to enable them to be flown

⁹ AT stands for Aerial Target, a misnomer intended to fool the Germans as to the real intention of the device – a flying bomb or “Aerial Torpedo” to be used against ships or zeppelins.

¹⁰ Interestingly, the Germans appear to have recognised potential and tried twice to assassinate Low!

jointly by pilots and by remote control. The program was code-named Operation Aphrodite (the Navy's equivalent was called Operation Anvil) and the aircraft were nicknamed "Weary Willy" after the war-weary aircraft that were being modified. The aircraft, which were renamed BQ-7 Flying Bombs, were stripped of their usual military equipment (including the cockpit canopy) and packed full of about nine tons of Torpex high explosive. They were then used to target bomb-resistant German fortifications such as submarine pens and V-1 missile sites. The drones were flown part-way to their destination by pilots who then baled out, leaving another pilot in a nearby manned aircraft to control the Weary Willy by viewing its cockpit instruments and the external environment remotely via a television camera placed in the nose of the drone.


The first Weary Willy mission took place on August 4 1944 with the target a V-1 site in Pas-de-Calais. Unfortunately, one of the drones went out of control shortly after the first crewman bailed out and crashed near the French coastal village of Orford, creating an enormous crater and destroying two acres of trees. The second drone was successfully flown toward the Pas-de-Calais, although clouds obscured the television view from its nose and it missed its target by 500 feet. Subsequent missions were no more successful as they tended to fly predictable trajectories and were easily shot down by German flak or were simply inaccurate. In one rather alarming episode, the remote control developed a 'mind of its own' forcing the explosive-laden drone into a circle over the industrial heart of the city of Ipswich in the UK before eventually crashing harmlessly in the North Sea. The program was eventually cancelled after only about a dozen missions, with the only drone that did manage to hit its target failing to explode, thereby providing the Germans with an almost intact B-17 bomber and a complete set of remote controls. In the final analysis, the program was probably more dangerous to the pilots flying the BQ-17 drones as it was to the Germans.



Air-launched variant of the V-1, launched from a Heinkel HE-111 (1945)

Germany, on the other hand, seems to have exploited the capability potential of UVS more successfully, having focused its considerable attention on the development of its V-1 flying bomb (also known as the “Doodlebug”). The key technical innovation appears to have been the reliability of the German guidance systems and engines. The guidance system used a simple autopilot to regulate the UAV’s height and speed and a countdown timer driven by an anemometer on the nose to calculate range relative to distance from the intended target. The countdown timer was used to arm the bomb and set in motion the terminal dive phase of the V-1’s mission. The Argus-Schmidt pulse (or resonant) jet engine initially required appreciable air flow to operate, which in turn required steam-powered launch ramps. These were conspicuous and hence vulnerable to air attacks. Later versions were able to operate at zero airspeed owing to the nature of the intake and acoustically tuned resonant combustion chamber. This allowed it to be rail launched from more discrete sites and hence deployed covertly at ‘mobile’ facilities away from air bases. Late in the war several air-launched variants of the V-1 were also developed.

Development of remote control and autonomous ground vehicles seems to have lagged behind that in other environments, presumably for the same reasons that it still does today: the complexity of the terrain. The first military unmanned ground vehicles appear to have been patented in 1915 during WWI. The invention of an American, Victor Villar, and an Englishman, Stafford Talbot, the Land Torpedo was designed to clear a channel through obstacles such as barbed wire. In their patent application, Villar and Talbot proposed transporting a torpedo across no-man’s land using a two-cylinder steam engine. The engine had no reverse gears and only a most basic control cable. It is not known if any were ever manufactured.



Villar & Talbot’s Land Torpedo (1917)

The first RC ground vehicles appear to have been the TT-26 tanks developed by the Soviet Red Army during the 1930's and used during WWII and the Winter War of 1939-40 against Finland. Two tele-tank battalions were built, each tele-tank being controlled from another tank some 500-1500m away. The tele-tanks were armed with machine guns, flame throwers, smoke canisters and/or 200-700kg bombs that could be dropped near enemy fortifications and detonated remotely. The tele-tanks could also deploy chemical weapons, but these were never used. When not in combat, the tele-tanks were driven manually.


During WWII the Russians also investigated what could conceivably pass as 'unmanned' anti-tank vehicles, although it is doubtful that the project received RSPCA or PETA approval. Significant amounts of explosive were strapped to dogs, which had previously been trained to associate the underside of tanks with food. Unfortunately, the dogs had not been trained against German tanks and they only associated the underside of Russian tanks with food. Fortunately for all, the project was quickly discontinued.



Russian TT-26 Tele-Tank (circa 1941)

The Germans also employed tele-operation on ground vehicles during WWII in the form of the Beetle Tank (also called the Goliath Tracked Mine). This invention was originally developed by a French engineer, Adolphe Kegresse, but in 1940, after France had been over-run, it started to be manufactured by the Carl Borgward Corporation in Bremen. The Beetle Tank was a tracked vehicle


approximately 1m long that carried around 75kg of high explosive. Used sacrificially as an anti-tank, anti-infrastructure, or bridge demolition device it was initially powered by unreliable electric motors, and later by a diesel engine. Although used extensively during the Warsaw uprising of 1944, it was not considered a great success as it travelled very slowly (9km/hr), had thin armour and poor ground clearance (11cm). Furthermore, the main unit was connected to a small control box via a triple-strand of telephone wires which were vulnerable to entanglement and damage. Around 7,500 Beetle Tanks were manufactured; several are now in museums.



Bundesarchiv, Bild 146-1080-053-53
Foto: o. Ang. | 1943/1944 ca.


German WWII Beetle Tank or Goliath Tracked Mine (1940)

The first fully autonomous UVS (i.e. they responded to their environment and were not externally controlled by humans) were UGVs developed by William Walter at the Burden Neurological Institute in Bristol, England around 1948-49 [287]. The vehicles responded to contact with objects (and each other) and had phototube ‘eyes’ that could sense light. They relied upon these to navigate using two vacuum tube amplifiers that drove relays to control steering and drive motors. Resembling a tortoise and called “Elsie” they exhibited simple intelligence, dancing near lighted re-charging huts until their batteries ran flat. Walter actually developed them to study the emergent and complex behaviour of networks of simple organisms.



Left: William Walter's 'Elsie' in its original shell (1949), Right: Johns Hopkins 'Beast' (circa 1959)

The next, and significantly more sophisticated autonomous vehicle, appeared in the late 1950's or early 1960's. Johns Hopkins "Beast" could apparently wander the University's hallways, centring itself using sonar until its batteries ran low. It would then locate black wall sockets using optics, plug itself in and recharge; whereupon it would resume its patrolling duties. Later in this decade Stanford University and SRI developed the first autonomous UGVs to be controlled by a computer: "Cart" and "Shakey" [98]. Although the computers filled a room and communicated with their vehicles via a radio link, both systems 'saw' with TV cameras (Shakey later got a laser range finder). Shakey would take pictures of its surroundings and then plan a path to the next room that avoided obstacles; move a little, take another picture, re-plan and so on. While Shakey is generally accepted to have been rather unreliable (each move forward was only a few metres and took about an hour for the processing to complete), the Cart was more successful surviving as a research platform into the next decade.



Left: The Stanford Research Institute's 'Shakey' (circa 1966), Right: The Stanford 'Cart' (circa 1963)

2.2 Modern UVS

Progress on military UVS from these earliest incarnations appears to have been fairly steady, with most major military powers taking an interest and progressively working through the levels of autonomy described in the next section. For instance, in the latter part of WWII the British successfully used “drone ships” for laying smoke during the allied invasion of Normandy in 1944. Such UVS were also used by the US to collect radioactive samples following the atomic bomb blasts at Bikini Atoll in the South Pacific in 1946; and by 1954 remotely operated minesweepers had been developed. By the 1960’s the technology had progressed to the point where it was regularly employed for missile firing practice and gunnery training.

Similarly, although the technological promise of UAVs was not realised in time to be fully exploited by the allied forces during WWII, they were seriously considered as an alternative to their manned counterparts in an attempt to avoid the carnage that resulted in the loss of 40,000 aircraft and 160,000 crewmen. Moreover, their promise does appear to have been recognised soon after the cessation of hostilities with the USAF awarding a contract to the Ryan Aeronautical Company in 1948 for the XQ-2, a sub-sonic, jet-propelled UAV [147]. During the 1950’s and early 1960’s their usefulness was also explored much more fully by the US during the Korean and Vietnam wars in their now traditional surveillance and reconnaissance roles.



Ryan Aeronautical Company’s XQ-2 UAV (1948)

Since the 1960’s there have been significant advances in miniaturisation, computer processing, sensor, signal and image processing, communications techniques and materials science. As a result, UVS are now able to sense their world using electro-optic and infra-red cameras, microphones, pressure sensors, and electronic olfactory sensors. Furthermore, localisation for most UVS has become routine since the Global

Positioning System (GPS)¹¹ was declassified in 1983 following the shooting down of the Korean airliner KAL007 over the Kamchatka Peninsular.

Likewise, advances in computer processing techniques have resulted in an explosion of sophisticated estimation, optimisation, feature extraction, tracking, and adaptive and machine learning techniques, which have all combined to enable rapid progress to be made in the science of autonomy and military UVS more generally. As a result, autonomous UVS can now capture, represent and interpret relevant environmental cues (location, geometry, spectral content) and then autonomously combine and manipulate this information such that the result is a series of control actions representing system priorities relevant to a mission and allocated by humans or another UVS. Additionally, the price of the technologies has dropped significantly such that UVS are approximately 80% cheaper now than they were in 1990. As a result, most countries around the world are investing in military UVS technology, most notably the United States. Although Japan probably has a technical lead in the area of humanoid robots (which are not discussed here) and the human-machine interface, the US probably has a lead in almost all other areas. Even non-state actors have started to develop their own UVS¹².

From the perspective of understanding the current and short-term military future of UVS one can do a great deal worse than studying the United States DoD Unmanned Systems Roadmap 2007-2032 [278], the Integrated Roadmap for 2009-2034 [282], or any of the equivalent documents such as the Navy's Unmanned Undersea and Surface Vehicle Master Plans [204] [205] [206] or the Joint Robotics Plans [141]. These documents essentially articulate six key goals, which also emerge consistently in the capability planning documents of other nations such as Australia, the UK and Canada¹³. Furthermore, these roadmaps also map

¹¹ UUVs are unable to localise using GPS unless they surface. Consequently, they must rely on dead-reckoning, relative navigation techniques such as SLAM (Simultaneous Localisation & Mapping) derived from their own organic sensors. There are a number of alternative 'underwater GPS' systems that use acoustic beacons [286], but these are typically able to communicate with UUVs only at rates of 2-3kbps over ranges of around 10km.

¹² For instance, Hezbollah have flown several UAVs into Israel and the Israelis have found several model aircraft fitted with explosives. There is also evidence to suggest that insurgents in the current theatres in the Middle East and Afghanistan re-cycle crashed allied UAVs. Similarly, the FARC in Columbia also has a number of UAVs, while other criminal elements now regularly use unmanned systems for surveillance and drug transportation, particularly in the maritime environment.

¹³ These key goals, which are discussed in detail in the section on Force Integration are: improve the effectiveness of combatant commanders and coalition unmanned systems through improved integration and joint services collaboration; emphasize commonality to achieve greater interoperability among system controls, communications, data products, and data links on unmanned systems; foster the development of policies, standards, and procedures that enable safe and timely operations and the effective integration of manned and unmanned systems; implement standardized and protected positive control measures for unmanned systems and their associated armament; support rapid demonstration and integration of validated combat capabilities in fielded systems through a flexible prototyping, test and logistical support process; and aggressively control cost by utilizing competition, refining and prioritizing requirements and increasing interdependencies (networking) among DoD systems.

out a aspirations for UVS autonomy through reference to increasingly sophisticated levels of autonomous control: (i) remote guidance, (ii) real time health monitoring & diagnosis, (iii) adaptation to failure & local conditions, (iv) onboard route re-planning, (v) group coordination, (vi) group tactical re-planning, (vii) group tactical goals, (viii) distributed control, (ix) group strategic goals, and (x) fully autonomous swarms.

For a complete reference to the huge variety of vehicles that exist (there are thought to be in excess of 10,000 in the inventory of the US alone) the reader is referred to the many works that catalogue some of the better known ones (e.g. [49] [141] [202] [230] [234] [237] [245] [244] [268] [276] [277] [278] [279] [280] [281] [282]). Many of these works also detail technical developments in the international arena, with some providing descriptions and projected goals on a country-by-country basis. To provide a contemporary cross-section of those systems actually flying, floating or driving in operations today, however, a small miscellany of UVS is described here.

UAVs, UMVs and UGVs currently in operation, range in size (and cost) from hand-held systems that cost less than a few hundred dollars to UAVs the size of manned aircraft and costing in excess of \$100M or UGVs weighing in excess of 40tons. The performance characteristics of these vehicles are very different and intimately tied to their operational requirements. Most operationally deployed systems, however, are tele-operated or at most semi-autonomous.



Northrop Grumman RQ-4 A/B Global Hawk UAV (2006)

The largest and most capable operationally deployed UAVs are the Medium Altitude Long Endurance (MALE) MQ-1 Predator and MQ-9 Reaper (armed with AGM 114 Hellfire missiles) and High Altitude Long Endurance (HALE) RQ-4

Global Hawk, which has endurance in excess of 30 hours, a payload capacity of around 1,360kg and can operate at altitudes in excess of 65,000 feet. These UAVs typically carry electro-optic and infra-red sensors and synthetic aperture radar (SAR) with a moving target indicator (MTI) capability, providing continuous (i.e. day/night) all-weather surveillance to ground control stations located beyond the horizon and usually thousands of miles away.




General Atomics MQ-9 'Reaper' Predator-B UAV armed with Paveway Laser Guided Bombs (LGB) Missiles (2007)

Shorter range tactical UAVs like the Hermes 450, the Sperwer, the Hunter and the Shadow typically have much shorter ranges (100-300km) and operate at altitudes closer to 15,000feet. Smaller again are the 'minis', which have wingspans around 2-3m and much smaller payload capacities. These smaller UAVs usually have significantly lower endurance, typically of the order of 1-2hours¹⁴ and operate at altitudes between 100m and 1km and are frequently electrically powered.

A number of UCAVs have been developed, the most prominent of which are the Boeing X-45 and Northrop Grumman X-47, known as Joint Unmanned Combat Air Systems (J-UCAS). They were developed for strike missions such as the Suppression of Enemy Air Defence (SEAD), electronic warfare and associated operations. The X-45 J-UCAS has gone through a number of iterations and is currently a swept wing, stealthy design of composite construction using foam matrix core and a composite fibre-reinforced epoxy skin, with a wingspan of 10m

¹⁴ Although some (e.g. Aerosonde and Scan Eagle) have significantly greater endurance. For example the Aerosonde has a demonstrated endurance approaching 40 hours.

and overall length 8m. The fuselage carries two internally housed weapons bays and an internally mounted Honeywell non-afterburning turbofan engine. It has a payload capacity of around 2,000kg, a cruising altitude and speed of about 35,000feet and 0.75 mach, respectively; and a range of about 600km.



Boeing X-45-C Joint Unmanned Combat Air System (2007)

In fact, UAVs now come in almost all shapes and sizes and with various strategies for lift and mobility, including standard helicopter rotary wing and ducted fans for Vertical Take-Off and Landing (VTOL), flexible wings such as parafoils and airships that hold considerable promise for long endurance communications relay missions. In particular, they now come in ‘micro’ – Micro Air Vehicles (or MAVs) typically having a wingspan less than 30cm and weigh around 300-400gm. Like their much larger counterparts, they come in fixed wing and rotary wing versions, but only have an endurance of around 30mins. The science associated with designing MAVs and flying them autonomously and stably outdoors in moderate weather conditions (i.e. light precipitation and <20knot winds) has been demonstrated repeatedly. However, these systems still require nocturnal capability, quick launch mechanisms, autonomous landing capability, or the capacity to fly in and out of buildings/caves autonomously. Improved portability and more physically robust MAV designs are also needed. Nevertheless, several MAVs are now available from a number of online electronic stores for <\$100.



Left: Aerovironment WASP II UAV (1998), Right: The ‘Mosquito’ MAV – available for <US\$50 (2008)

Like UAVs, UMVs also occur in a spectrum of shapes and sizes. In the case of UUV’s they vary from the smaller, ‘man-portable’ systems typically around 10-20cm diameter and 0.3-1m long with less than 50kg displacement and that carry sensors with limited range and resolution, through the light-weight and heavy-weight systems that have around 250kg and 1,500kg displacements, respectively and greater endurance and payload capacity, up to the much large displacement UUVs, which have around 10tonne displacement and significantly increased payload and endurance capabilities. The lighter weight UUVs are usually inexpensive and are designed to survive only a few missions or not be retrieved at all; some extending their endurance by perching on local terrain.

Large numbers of UUVs have diameters tailored to torpedo launch tubes (i.e. 12.75 or 21 inches) and have already seen operational service. Typical examples include man-portable systems such as the Semi-Autonomous Hydrographic Reconnaissance Vehicle (SAHRV), the Remote Environmental Measuring Unit (REMUS) family of UUVs and the multi-reconfigurable unmanned undersea vehicle (MRUUV).¹⁵



Left: Hydroid’s REMUS-100 UUV & Right: The Slocum Glider

The MRUUV has dimensions similar to those of a heavyweight torpedo and is used onboard Los Angeles and Virginia class attack submarines. It is 6m in length,

¹⁵ The Phase I RDT&E funding for MRUUV was cancelled in December 2008. However, the 2010 budget plan contains for the future of this program.

has a 21in diameter, and weighs about 2,000kg. Like many UUVs it can be reconfigured for a range of different missions by switching modules, such as maritime reconnaissance, undersea search and survey, communications and navigation aid, submarine trail and track and mine identification. The MRUUV has an open architecture for technology spirals that enable less expensive upgrades to its system over the course of its service life.

Most UUVs carry acoustic payloads such as sonar or Synthetic Aperture Sonar (SAS) for identifying mine-like objects. SAS, however, are power-hungry devices and severely limit endurance. Similarly, communications between the operators and the UUVs is often severely bandwidth limited due to the nature of the medium in which they operate. As a result, many UUVs often take the form of Remotely Operate Vehicles (ROVs) and communicate with their operators over electrical and/or fibre-optic cables. ROVs with endurance of up to 70 hours are reasonably common, whereas stand-alone UUVs have endurances an order of magnitude less.

UUV gliders, such as Slocum, are also now common. These UVS harness the ocean's natural energy to 'fly' or 'porpoise' through the water by alternating their buoyancy to either rise or fall. As these UUV's zigzag vertically, their fins allow the vehicles to traverse the oceans horizontally, and although they travel relatively slowly (~2km/hr) it is common for the vehicles to have endurances of several months and therefore the capacity to travel great distances (i.e. in excess of several thousand kilometres).



Left: Spartan Scout USV (2003) & Right: Rafael Protector USV (2005)

USVs such as the Spartan Scout, Protector, Sentry, Stingray and SWIMS (Shallow Water Influence Mine Sweeper) and Remote Mine sweeper (RMS) have also been deployed operationally. The RMS and SWIMS missions are to detect, classify, localize, and identify bottom and moored mine threats in shallow and deep water. RMS is an air-breathing, diesel-powered semi-submersible that autonomously follows a preplanned mission plan. The vehicle deploys a variable depth sensor, comprising acoustic and EO sensors to positively identify objects of interest. The data link sends information collected by the USV back to the host ship via line-of-sight or over-the-horizon transmissions.

The Spartan Scout, which is fairly typical in terms of performance, is a 7-11m Rigid Hull Inflatable Boat (RHIB). The 7m version has an 8hour endurance and 150mile range, and can be operated by remote control or as a modular,

reconfigurable, multi-mission, high-speed, semi-autonomous or fully-autonomous USV. It is capable of transit speeds of up to 28knots and carrying payloads up to 1,500kg (7m RHIB) and has been demonstrated in a range of roles including mine warfare, force protection, and precision strike scenarios. The larger version (2,500kg, 11m RHIB) has commensurately greater endurance and speed. Weaponised versions that carry stabilised machine guns and Hellfire or Javelin anti-armour missiles have also been developed.

The autonomous search and hydrographic (ASH) vehicle and the Roboski were initially developed in the 1990s as jet-ski type target drones for ship self defense training. They are now also used as reconnaissance vehicles and operate as remotely controlled vehicles. The Owl USV is a commercially available modification of ASH, with a low-profile hull for increased stealth and payload capacity, although it too operates in a remote control mode. Several variants for stealthy USV sensor platforms are under development.

UGVs are currently used for a range of surveillance, EOD, mine clearance, obstacle breach fire fighting duties. They range in size from automated and remotely controlled versions of the Abrams tank (over 40 tons) to an explosion in the operational deployment of small and remotely controlled UGVs, which are used for a range of EOD, IED and sub-tactical ISR duties (e.g. Talon SWORDS, Andros, PackBot, ToughBot, iBot and Dragon Runner), frequently having been lobbed into buildings or caves to search for hostile personnel or devices. Other operational systems include the Automated Ordnance Excavator AOE (34tons), D7G, a combat engineering vehicle (28tons), the deployable universal earthmover DEUCE (18tons), and a range of smaller vehicles of the order of 1-15tons such as the All-purpose Remote Transport System (ARTS) or the experimental unmanned vehicle (XUV), which provide a variety of combat-support roles.



Left: Foster-Miller TALON SWORDS (2003), Top Right: Packbot (2003) & Bottom Right: Dragon Runner (2002)

Until recently, most of the heavier UGVs were automated or tele-operated versions of manned vehicles. In the last few years, however, more autonomous, dedicated design tactical UGVs (e.g. Gladiator) have also been deployed, particularly in urban terrain, to provide ISR and firepower ahead of dismounted soldiers, although these have relatively simple autonomous navigation capabilities. There are also larger, more rugged systems designed for a wider variety of terrain to provide perimeter surveillance (e.g. Guardium and TAGS), some of which carry remote control weapons systems (e.g. See-Shoot, which is deployed in Gaza, or SGR-A1, which is deployed in the Demilitarised Zone between North and South Korea).

A range of much larger, more autonomous UGVs with greater mobility such as MDARS, Crusher and Multi-function Utility Logistics Equipment (MULE) vehicle are also being developed as a component of the US Future Combat System. Crusher and MULE are both autonomous vehicles whose technologies are, among other things, aimed at automating logistic support. These systems provide relatively high degrees of automation in their basic functions, but have limited autonomy in more complex tasks.



US Marine Corps' Gladiator UGV (2004)

UVS with the potential to become available within the next five-ten years range from nanobots¹⁶ that are devices that range in scale from about 0.1-10 μ m and are

¹⁶ Based largely on Micro-Electronic Machines (MEMs) technology, many nanobot programs are currently hypothetical or at best projected capability, although some have been tested in regard to their sensing and control capabilities albeit largely in the medical domain. These UVS are not realistically expected to deliver significant kinetic effects, but may be useful for indoor ISR, communications, or cyber attack.

constructed from nano-scale or molecular structures to UAVs that rely on solar (or hydrogen) power and aspire to have the capacity to carry 500kg payloads and stay aloft for up to five years.



Crusher UGV (2007)

Other UVS perhaps closer to deployment include a walking robotic mule¹⁷, submersible UAVs¹⁸, soldier-portable weaponised UAVs, boomerang UAVs¹⁹,

¹⁷ DARPA robotic mule program, known as Big Dog, is currently a prototype quadruped walking system capable of carrying around 200lbs and following a dismounted soldier across 8miles of unstructured terrain, obstacles, etc over a period of 3hours (and surviving significant disruptive side impact). The version to be released within the next year will be significantly larger (1,250lbs, payload capacity 400lbs), responsive to hand and voice commands and have a 24hour (20mile) endurance.

¹⁸ The performance requirements for the proposed submersible UAV are that it can cover 1,850km by air or 185km by sea-surface, or 22km underwater in eight hours or less; and carry a 1,000kg payload [202].

¹⁹ The “boomerang” UAV is an intriguing, almost comical device, comprising a three-pronged fuselage that is thrown by hand, rather like a boomerang. As the unit spins through the air the EO camera at its centre takes images that are automatically corrected for the rotation of the UAV and transmitted to the users at a ground station as a mosaic. By tilting the UAV as it is thrown, it can be made to fly around corners [169].

micro UAVs with payloads weighing less than 2g,²⁰ UMVs capable of operating in the surge-dominated surf zones, mechanical insects, UGVs that disintegrate²¹ and UGVs that consume environmental matter to extend their endurance. Furthermore, there are at least as many programs pertaining to the functional elements of UVS that are likely to become available within a timescale of a few years that could have a dramatic impact on UVS and their autonomy. These include such things as 1GPixel cameras, artificial gecko feet to allow micro and nano UAVs to land on ceilings and walls, brain-based interfaces with real time contextual memory interpretation, management and data analysis techniques for Petabytes of imagery, self-repairing materials and materials that allow the structure of the platform to act as the sensor (i.e. the UVS *is* the sensor). The list is as endless and as diverse as human imagination.

2.3 Looking Forward

“In future military conflicts, having a strong bladder and a big butt may turn out to be more useful physical attributes than being able to do a hundred push-ups”

P.W. Singer (2009)

Forecasts on autonomy, robots, artificial intelligence have been made over a period of more than 50 years without bearing much fruit. Furthermore, many level-headed people involved in the use and capability planning of UVS will be justifiably sceptical about such predictions. Nevertheless, in 2008 Hans Moravec, co-founder of Carnegie Mellon University’s Robotic Institute, said “I see a strong parallel between the evolution of robot intelligence and the biological intelligence that preceded it. The largest nervous systems doubled in size about every fifteen million years since the Cambrian explosion 550 million years ago. Robot controllers double in complexity [processing power] every year or two. They are now barely at the lower range of vertebrate complexity, but should catch up with us within a half century.” Here are his educated predictions [241]:

²⁰ Typically, autonomous MAVs weigh about 400g. However there are programs aimed at developing Nano Air Vehicles (NAVs) that have the capacity to carry out military missions and yet weigh less than 10g (and <7.5cm) [135]. Some experimental systems already demonstrated, such as the Delfly Micro, weigh around 3g, (the payload and battery weigh about 1g each). It is common for navigation systems of NAVs to be modelled around insects such as dragonflies. Even smaller UAVs have been demonstrated, but to date have been linked to an external energy source.

²¹ The disintegrating UGVs comprise two larger, two-wheeled tubular vehicles joined together by two smaller, similarly shaped tubular UGVs. The single, aggregated UGV is robust and portable and can be thrown into buildings etc. As an integrated system it can carry larger payloads and has better mobility characteristics than its component UGVs. However, on command, the system can break apart into its four constituent smaller UGVs, each with cameras and each independently controllable. Alternatively, the UGV can be carried as four independent UGVs.

2010: A first generation of broadly-capable ‘universal robots’ will emerge. These ‘servant’ robots will be able to run application programs for many simple chores. These machines will have mental power and inflexible behaviour analogous to small reptiles.

2015: Utility robots host programs for several tasks. Larger ‘utility robots’ with manipulator arms able to run several different programs to perform different tasks may follow single-purpose home robots. Their $\sim 10^{10}$ ops computers would support narrow inflexible competencies, perhaps comparable to the skills of an amphibian, like a frog.

2020: Universal robots host programs for most simple chores. Larger machines with manipulator arms and the ability to perform several different tasks may follow, culminating eventually in human-scale ‘universal’ robots that can run application programs for most simple chores. Their 10^9 lizard-scale minds would execute application programs with reptilian inflexibility.

2030: Robot competence will become comparable to larger mammals. In the decades following the first universal robots, a second generation with mammal-like brainpower and cognitive ability will emerge. They will have a conditioned learning mechanism, and steer among alternative paths in their application programs on the basis of past experience, gradually adapting to their special circumstances. A third generation will think like small primates and maintain physical, cultural and psychological models of their world to mentally rehearse and optimize tasks before physically performing them. A fourth, humanlike, generation will abstract and reason from the world model.

As such systems are not likely to be with us anytime soon in military terms, let us set aside the longer-term, thought-provoking future scenarios (and others that have humans interfaced with robots so that our brains can be transferred into more capable computers or that have robots disposing of humans once they have become intellectually superior to us and *we* inhabit *their* planet [159]). Nevertheless, let us remember that many of these scenarios are not delivered by science fiction writers, but by highly intelligent, respected and experienced researchers in the field. Let us also note that by these calculations by about mid-way through the 21st Century these predictions indicate that all physical and intellectual tasks currently carried out by humans should be within the grasp of robots.

In our world, military UVS will not be expected to assume all of the intellectual tasks of a human, at least for some time to come. Furthermore, if we base our predictions on an assessment of the US DoD Master Plans and Integrated Roadmaps (e.g. [44] [205] [206] [237] [278] [279] [281] [282]), we can also establish estimates of when Technology Readiness Level (TRL) 6 [218] will be achieved in some of the key functional areas of UVS technology (Figure 2.1). [279] in particular provides a breakdown of the component technology enablers by year and TRL, which indicates that fully autonomous UVS will likely become a reality around 2015.

Technology Areas	Remote Ctrl/ Tele-operated	Semi- Autonomous	Platform- Autonomous	Network- Autonomous
Autonomous Navigation				
Planning & Dynamic C2				
Health & Usage Monitoring				
Power & Propulsion				
Human-Machine Interaction				
Machine Intelligence				
Sensors & Data/ Info Fusion				
Comms & Networks ²²				

COTS or MOTS	
Avail. 2010 – 2015	
Avail. 2016 – 2025	

Fig. 2.1 When will TRL 6 be reached for the key functional areas of UVS?

Consequently, even though three or four years of mainly software engineering will elapse between the availability of any systems at TRL 6 and systems that are operationally deployable (i.e. TRL 9), if these predictions are even only approximately correct, then many militaries will have access to deployable, fully autonomous UVS within time scales commensurate with the soft end of their current financial planning cycles. The challenge is then for defence forces to adapt their practices to accommodate this technology.

Instantiation of Intelligent Decision-making Techniques (IDT) within a UVS that allows functional replacement of a manned asset obviously requires development of behaviours approaching those of a human and there are currently many limitations that make this aspiration very unrealistic. For example, portable computer processing cannot currently mimic the processing capacity of the brain, the relevant architecture is not yet sufficiently well-understood and therefore optimised to allow (for example) complex perception or high level reasoning, and the software and algorithms cannot yet imitate the contextual decision-making or visual perception capabilities of human behaviour.

In 1997 two influential analyses [196] and [47] looked at the computational capacity of the brain and extrapolated the processing power of micro-processors to determine how long it would be before the intelligence onboard robots [196] or super-intelligence [47] achieved human equivalence. Let us briefly re-visit these analyses with updated data to understand the relevant timeframes, estimate the degree of sophistication that we might expect from UVS over the next twenty years, and understand how much room is still needed for improvement.

The human brain occupies a volume of around $1,500\text{cm}^3$, draws around 25 W, contains about 10^{11} neurons and operates in a similar fashion to a micro-circuit; each neuron has about 5,000 synapses and signals are transmitted along these synapses at an average of about 100Hz. Each signal contains around 5 bits. This limits brain performance to an upper limit of 10^{11} millions of instructions per second (MIPS). Functional replication of elementary auditory processing and observations of the signal processing requirements of the retina, however, indicate that the brain is in fact capable of 10^8 MIPS [47] (the discrepancy is thought to imply redundancy).

Brains also require memory, which estimates place at around 10^8 Mbytes [6] [47]. The ratio between semi-conductor based computer memory and speed has remained relatively constant at about 1byte/ops over the years. As a result, in terms of achieving human equivalence, building and accessing enough memory is likely to be no harder – and quite possibly considerably easier – than developing a fast enough micro-processor. At present, a high-end computer that draws $\sim 1\text{kW}$ and is considered reasonably portable is capable of $\sim 10^5$ MIPS [137] at present or perhaps between 10^3 and 10^4 MIPS for a less expensive laptop, which would draw around 100W.

The future performance capabilities of such systems are traditionally estimated using Moore's law, which states that the number of transistors on a chip doubles almost every 18 months (two orders of magnitude every 10 years).²³ Although based on a somewhat brute-force approach, these crude extrapolations provide estimates of 10^8 MIPS for high-end portable computers somewhere around 2025. Interestingly, these estimates accord with those derived by applying the historical exponential reduction in the cost of computer power. For example, super-computers that cost around \$100,000,000 exceeded the 10^8 MIPS requirement in 2005 [178] and have more recently surpassed 10^9 MIPS [114]. Using the historical data, we can anticipate that such processing power will be available for \$1,000 around 2025. As a result, it can reasonably be argued that this largely establishes a timescale for the processor requirements for tackling the most demanding or high-level 'thinking' tasks, such as ethical governance or contextual decision-making (see later). Simultaneously, the current three-to-four orders of magnitude shortfall also indicates why current military UVS exhibit such primitive levels of autonomy and intelligence.

UVS and brains are, of course, required to do many things aside from high-level reasoning and the brain's major conscious task, visual perception is aided by the fact that 90% of the processing is carried out in the retina, which joins the 10^8 pixels of the eye to the brain via 10^6 nerve fibres that operate using pulse code modulation at about 50Hz [196]. In this regard, several decades of vision research have shown that basic image processing tasks for UVS (real-time detection and tracking of gray-scale patches in simple images) require between 10-100 MIPS, whereas more complex image processing tasks (tracking complex 3-D objects in clutter) require around 10^4 MIPS [196]. Furthermore, [6] and [119] have estimated that around 10^5 MIPS will be required for retinal-equivalent image processing (i.e. taking account of both image processing and data fusion requirements) and 10^6 MIPS for good human-driving performance.

A high end portable CMOS or CCD camera (10^7 pixel array operating at a readout rate of 10Hz with a 14-bit ADC) can deliver information at a rate of about 1GHz; two orders of magnitude short of the human visual perception system. Other systems, with lesser pixel density cameras can deliver much higher readout rates (10^6 pixels at 1,000Hz), closing the gap by an order of magnitude in terms of

²³ It is recognised, however, that Moore's law is not really a law but an observed reality (or self-fulfilling prophecy, depending on your point of view) used as a guideline by the semi-conductor industry to ensure that the technology advances uniformly across a complex front.

data rate. Regardless, in June 2008 NVIDIA released its Desktop 1.2 Teraflop Tesla GPU, which retails for around \$2,000 and has a bandwidth of the order of 100GBps, indicating that the raw processing power for the graphic manipulation may already be here. As a result, we can anticipate parity for perception processing capabilities around 2018, even though “driving” is not all that we will expect of an autonomous military UVS.

This indicates that raw computational power does not appear to be the major roadblock to achieving perception-based autonomy. The main problem would appear to be one of systems engineering and the need to optimise the perception architecture so that the sensors, embedded processors, algorithms and communications are tightly integrated. The use of embedded simulation, however, can aid in resolving this (see Simulation later).

It is appreciated that the above rather crude approach of using MIPS as the basis for establishing the functional performance of an automated system without taking account of the architecture is flawed. For example, a fruit fly with a brain containing only about 250,000 neurons and a fraction of the processing power of a modern PC, can effortlessly perform three-dimensional navigation manoeuvres well beyond the most sophisticated UVS. Furthermore, the human brain is thought to be massively parallel and the modern computer only moderately so²⁴ (and historically massively sequential). Nevertheless, this analysis is only meant as an approximate guide. Additionally, as [196] and [47] both point out, brains – like UVS – also need ‘software’.

Furthermore, based on the above projections, speeding up IDT is probably comparatively unimportant compared to the creation of better IDT. Additionally, based on current experience and the variety of contexts in which military UVS will be placed, it is likely that a range of research techniques will need to be applied and integrated if UVS are to be capable of undertaking the ‘general’ duties of humans and manned vehicles. This is because no individual technique is robust to the broader problem space; as the problem space changes, so must the technique. As a result, IDT will probably require software techniques that interpret context and select a process or solution technique accordingly. In this regard, [142] notes that the three most significant algorithmic impediments are probably the development of pattern recognition techniques, the encapsulation of common sense, and the capacity to generate synthetic emotion, which is discussed briefly in the chapter on UVS Architectures.

In determining the algorithmic requirements, however, it should be noted that we know the problem can be solved in principle as humans have achieved the requisite levels, albeit in a form that routinely makes mistakes. However, it is only recently that, as processing capabilities have increased to levels commensurate with lower order animal cognitive function, we have begun to gain real insights into some of the more sophisticated artificial intelligence techniques (with many describing progress to date as ‘glacial’). There are a number of potential research avenues that currently show promise in this area including: general computer

²⁴ [93] reports that by 2015 we can expect massively parallel computing architectures to evolve, with systems on a chip being produced on wafer scale. As a result, these systems may be embedded directly into sensors or other ‘mechanical’ devices.

science and mathematics (in particular adaptive machine learning and evolutionary techniques), cognitive science, brain imaging, biotechnology and brain imaging. For the prospect of intelligent autonomy to become a reality only one of the above, some other technique, or some combination needs to deliver the break-through. As research in autonomy and intelligence is a highly interdisciplinary field, however, each will benefit from developments in the other.

Some other techniques showing promise at the moment involve reverse engineering the brain and getting it to ‘bootstrap’ itself into intelligent thought through a process akin to evolution. In this regard, it should be noted that the Blue Brain project has recently been able to accurately simulate a section of the neocortical column²⁵ of a rat [178]. In theory, as the column now works, it only needs to be scaled up to replicate the entire brain. Furthermore, the scientists at Ecole Polytechnique Federale de Lausanne (EPFL) have indicated that they are confident that they will be able to start simulating the entire brain (including sentience) within the next few years.²⁶

A major impediment in determining when IDT software is likely to achieve human equivalence is that IDT is the study not of a single field, but rather the aggregate of many: some techniques are ‘top down’, while others ‘bottom up’. As with many such matters, it is likely to be a combination of both approaches that ultimately resolves the ‘general’ problem. That is, the required software is likely to learn from a combination of low-level environmental interaction (by “bumping into reality” as [142] puts it) and the higher-level ‘cognitive’ processing of derived or pre-programmed instructions.

It is also hard to measure human intelligence and hence where on the curve the performance of current algorithms lie. In essence, we are attempting to assess the state of science. Nevertheless as perception systems must abstract information from sensed data to plan and execute tasks, we may measure the ability of one or more algorithms to perform the following [208]:

- Sense the environment and the internal state of the UVS
- Perceive and recognize objects, events and situations
- Remember, understand, and reason about what is perceived
- Attend to what is important and ignore what is irrelevant
- Predict what will probably happen in the future under a variety of assumptions
- Evaluate what is perceived and predicted
- Make decisions, plan, and act so as to achieve goals
- Learn from experience and from instructions

We can then record the historical, state-of-the-art, and human levels of performance in terms of accuracy, speed, efficiency and cost-benefit and then

²⁵ The neo-cortical column is a tiny slice of brain containing approximately 10,000 neurons, with about 30 million synaptic connections between them.

²⁶ It is important to note, however, that it is not possible to simulate a brain that is not perfectly understood. The best we can achieve is to simulate our current understanding of the brain and test that understanding against reality.

extrapolate as others have done previously for hardware to arrive at some equivalent predictions.

However, we should note that until recently research communities have typically focused on narrow areas of IDT algorithmic development. For example, in the field of image processing, they have concentrated on the development of edge detectors, image segmentation techniques, texture analysis, colour classification and stereo vision, all based on single-sensor modalities and with much of the processing done off-line. Any predictions should therefore be made on the basis of multi-sensor modalities carried out on real-time processors and in aggregate with all of the tasks highlighted above.

One consideration not often taken account of in predictions of this nature is software stability. An algorithm capable of performing (say) human-equivalent perception will probably comprise around 1,500,000 lines of code (LOC) (i.e. about 10,000 function points), which may not have been produced by the same organisation, let alone programmer.²⁷ Estimates indicate that it will take a further 3-6 months to document such code, 3-6 months to appropriately architect such code, 18-24 months to re-implement this on suitable processors, and a further 6 months for systems and performance testing [237]. Furthermore, the effects of software quality on perception performance and reliability are unknown. Additionally, software capable of contextual decision-making and other higher order or creative “thinking” will be significantly larger and more complex. In this regard, it is perhaps instructive to note that the first facial recognition programs were developed in the 1970’s, and only now are becoming stable and effective enough for regular use.

Clearly, it is preferable if we understand how the system works. However, it is possible that we may replicate some of the human sensory perceptions (which essentially perform Fourier transforms) and connect them to a copy of a ‘bootstrapped’ intelligence without fully comprehending how this ‘brain’ works. Regardless of the strategy we eventually use, it does not seem to make sense to make predictions beyond those already made by others more qualified to do so.

Nevertheless, let us note some considerations for the future [47] [195] [208]:

- Primitive levels of autonomy are likely to advance very rapidly once established; as soon as any relevant techniques are developed and stable, they may be copied and run on smaller, cheaper processors (i.e. unlike human evolution, the lessons are instantly transportable from entity to entity and do not need to be re-taught).
- Military UVS are only one of the many fields that will be accelerated by progress in the area. There will consequently be a synergistic relationship between these fields and the development of military UVS.

²⁷ Using the ‘rule of thumb’ (number of function points)^{1.22} = total number of defects to be eliminated and (number of function points)^{1.24} = total number of test cases required [93], 10,000 function points translates to around 76,000 defects, of which we can expect 20% to have high severity impact, with 95%-99% of these found prior to delivery (depending on the originating organisation’s maturity).

- As IDT approach human parity, they may be incestuously applied to the design and development of key techniques and functional components. However, it is likely that a law of diminishing returns will apply rather than an exponential acceleration as predicted by some.
- On reaching parity, the processors will rapidly exceed the capabilities of the humans. Moreover, using Moore's law as a guide we can assume that within ten years of this date there will be an abundance of machines capable of thinking 100 times more rapidly than humans. Militaries will then have to adapt.
- Such sophisticated machines would be independent agents; capable of initiating and making their own plans. That is, these entities are no longer really just tools.

Clearly, it is likely to be the development of stable software techniques that are likely to prove the greatest obstacle to the achievement of autonomous vehicles that reproduce or approach human equivalence. Before moving on, however, let us briefly consider some other limitations to IDT achieving such sophisticated levels of functionality observed by [142], [168] and others:

- In regard to micro-processor technology development, a number of estimates terminate Moore's law around 2020 when the limit of shrinking the current transistor is reached.²⁸ This is still two orders of magnitude short of human-equivalent performance on a portable scale.
- Given the magnitude of the global semi-conductor market any such limits are likely to be "hard" (i.e. additional research funding by any single corporation or nation is unlikely to result in progress). By the same token, such trends have been maintained for over 50 years through a number of other 'roadblocks' and a number of alternative strategies have already been proposed. However, such techniques and technologies are currently still in their infancy.
- The global financial crisis (or some other economic event) might curtail the rate at which technological advances are made. However, a study of the historical data indicates that such events do not actually affect the speed of research.
- Emotion appears to be a key ingredient for value-based decision-making.²⁹ As a result, if we are to enable IDT to observe their environment and then on this basis set and prioritise goals, process the related decisions and then communicate the outcomes of this decision-making efficiently to others, we may need to model or synthesise (and hence measure) emotions.

²⁸ At around this time the layer of atoms comprising the chip will become so thin that electricity will 'leak' out of the chip and the computer will short circuit.

²⁹ The study of patients that have specific brain-injuries allowing them to retain their reasoning ability but preventing them from feeling their emotions has shown that such people struggle to decide between the range of options [168].

- IDT will likely require sophisticated software that is based on the notion of mathematical logic. However, as Kurt Godel proved in 1931, there are true statements in arithmetic that can never be proven within the axioms of arithmetic. The result is that, as a means for expressing or describing our world, mathematics may prove to be incomplete. Furthermore, even if mathematics suffices as an approximation, for all practical purposes the computation must still be carried out 'quickly enough' or the problem is effectively incomputable.
- According to [224], the physiological processes underlying thought may involve the super-position of a number of quantum states, each of which performs a key calculation, before the differences and distributions in mass and energy cause the states to collapse to a single (measurable) state. This process cannot be replicated by any computer yet conceived.
- As the autonomous UVS become more complex, system certification costs will increase exponentially due to a projected increase in the testing requirements (the UVS can do more) and required testing resources (the systems are more complex and therefore require hardware in the loop and testing labour).
- As the results start to show real promise, the concept of sentience in a machine may trigger a religious or ethical debate which results in restrictions placed on the research. While this is possible, it is likely to be temporary, and Western nations are likely to care more about such things than other nations.
- If research results start to dry up the number of graduates studying in the area may decline. Alternatively, if multi-disciplinary approaches show little progress those studying in the area may choose instead to focus on narrower endeavours.

Chapter 3

Autonomous UVS

Cost drivers and the requirement for UVS to be mobile and adaptive means that a number of systems – either component technologies or multiple UVS – will need to be controlled by a relatively small number of people. This implies both a high degree of autonomy, and that humans will be retained within the decision-making loop at some level. In this chapter, therefore, the functional components and architecture of an autonomous UVS are described, together with the nature of human-UVS interaction. This is intended to provide an appreciation of the complexity and scale of such systems, and an understanding of the degrees of freedom that are often coupled when the performance of autonomous UVS are measured. The chapter lays the ground-work for discussing the technological challenges, the issues pertaining to technology-force insertion, and the legal conundrums.

3.1 UVS Components

The architectural representation of a modern autonomous UVS is complex, an example of which can be seen in Figure 3.1. It has multiple interacting and independent components such that the sum behaviour is a function of the individual components, the interactions between them and communication with any human supervisors. Many of the component technologies for autonomous UVS exist independently in largely mature forms, although due to their complex interactions significant systems engineering is still required to bring the sub-systems together such that the end result is a system that can operate autonomously in a complex military environment [91]. The development of robust integration schemes probably represents the most significant technical challenge to the deployment of practical military UVS.

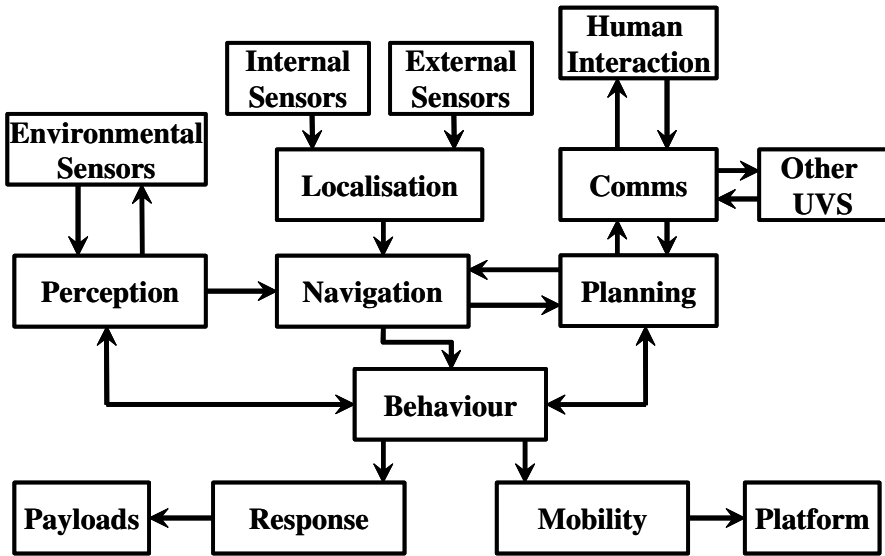


Fig. 3.1 The key functional relationships for an autonomous UVS

As the degree of autonomy increases, so it becomes increasingly difficult to predict the sum state of the system. Moreover, the system is actually a function of a number of linked hardware and software processing elements and humans (programmers, engineers and users) which becomes significantly more complex as these systems are networked either to each other or other technologies. Typically there are multiplicities of architectures, data formats, operating systems, programming languages, compilers and communications protocols, not to mention an infinite variety of hardware combinations.

The key UVS components, as given in Figure 3.1, can be described as follows:

- **Internal sensors** may measure wheel velocity (odometers), steering angle, ground (radar or laser) or sea floor (acoustic) Doppler or depth/altitude (pressure sensors). Wheel, thruster and steer encoders are used both for low-level closed loop control and dead-reckoning navigation. In UGVs, when used in conjunction with external sensors, wheel and steer encoders have a secondary role in making traction, slip, and other terrain characteristics observable. These sensors may also be classed as ‘proprioceptive’ – i.e. they perceive internal factors that are effected by the environment and the UVS’ own behaviour.
- **External sensors** (i.e. Inertial Measurement Unit (IMU), inclinometers, magnetic compass, and GPS) provide data that are usually fused and provide observations of the position and orientation of the UVS in some absolute frame of reference. These sensors are also usually considered to be proprioceptive.
- **Environmental sensors** (i.e. radar, LADAR, EO, IR and acoustic) are required by the UVS for its perception algorithms to observe and develop

a map of its environment within the particular segment of the spectrum. The choice of which sensor to rely on for any given function or action is often made by the perception algorithms. These sensors are often described as ‘exteroceptive’ – i.e. they perceive external factors that are not under the control of the UVS.

- **Localisation** typically provides estimates of the position, velocity, attitude, altitude rate and acceleration (or some combination thereof) of the vehicle. Localisation is usually an output-only function when viewed by the rest of the system [88].
- **Perception** is defined here as capturing, representing, and interpreting relevant environmental cues (e.g. location, geometry, spectral content, etc) observed by the sensors and relating these to features in the real world for the vehicle’s moment-to-moment control, mission and task planning, payload control, etc.
- **Navigation** is concerned with the generation of a map of the vehicle’s immediate surroundings, how to navigate around the immediate surroundings and to the next waypoint or final destination, and the detection of hazards relative to the vehicle’s mobility. It takes input from the localisation and perception functions and uses this information in conjunction with the behaviour function to execute a mission.
- **Planning** is the process of generating a trajectory or sequence of actions from a specified starting point to a goal position or activity while avoiding obstacles and impediments (i.e. mission planning). This function has no direct links with either sensory input or supervisor output, but must use an understanding of these, in conjunction with world maps and defined mission objectives, to produce appropriate commands.
- **Behaviour** (or cooperative tactics) combines the outputs of navigation, planning, and perception (often through a world map) and translates them into actuator commands for platform mobility and payload response. The specific behaviours depend upon the mission and the payload onboard the UVS, which will vary with the operational requirements. Behaviour modules are often associated with behaviour-based (or ‘reactive’) architectures, but may not be included in deliberative architectures.
- **Communication** provides the link between the vehicle and any other elements of the system, including operators and other (manned or unmanned) vehicles.
- **Human Interaction** provides the interface between the user/supervisor and the UVS and defines how the two systems work together. It covers the functional Human Machine Interfaces (HMI) that assist users to understand what the system is doing in relation to its mission or tasking, how the tasking is allocated between the human and the UVS, and how the impact of uncertainty and cognitive overload can be mitigated.
- **Mobility** is the ability of the vehicle to traverse its environment and for a UGV is often expressed in terms of the magnitude of an obstacle that the vehicle can negotiate. For UMVs and UAVs mobility usually includes the hydro and aero-dynamic properties of the platform. Mobility is often directly addressed through the design of the UVS platform.

- **Response** is the tasking of payloads to create or enhance a situational awareness picture for the user, or to respond to and engage with objects within this picture. Typical activities might include aiming, arming, and firing a weapon or panning, tilting & zooming an ISR payload to assist in the assessment of objects by a user.
- **Payloads** comprise an EO/IR, radar, LADAR, EW or acoustic sensor (or some combination thereof) employed for the purposes of (e.g.) ISR or REA, but may also include manipulator arms, EOD, MCM, weapons, etc.
- **Platforms** are the integrating frames for the total system, which must be tailored to the missions they are to accomplish. Relevant technologies include mechanical design, structural mechanics, materials, launch/recovery techniques, etc.

The following functional components are not shown in Figure 3.1, but are considered key elements of any autonomous UVS platform:

- **Energy storage** and the rate at which it can be used is key to any battlefield UVS. For small UVS the energy systems are typically electrical; larger UVS typically use fuel or hybrid-electrical systems.
- **Propulsion** systems are also key functional components of UVS and are typically designed around the tasks and missions likely to be undertaken. Designers must also take into account their signature, which can be acoustic, electromagnetic, infrared and visual.
- **Health & Usage Monitoring Systems (HUMS)** and built in test equipment are often ignored in smaller UVS but are frequently used for self-monitoring, diagnosis and remediation of systems or functional components in the larger ones.

Of the functional elements identified, [91] [245] and others classify the areas of platform, energy storage, propulsion, HUMS, payload, communications (and its management) and mobility as largely mature and existing in a deployable form. Localisation is also mature for most domains. Response is similarly well-established, but due to the uncertainty surrounding the systems integration and information assurance – to say nothing of the legal and policy considerations – a ‘safe’ response unit for autonomous weapons has yet to be formally certified for operational use. Navigation and mission planning (including dynamic re-planning) techniques are also largely mature, but their application to military missions still needs to be demonstrated in complex and unstructured environments. Perception of the environment and the instantiation of tactical behaviours are the least well demonstrated functions. There are also issues in regard to the fusion of data across the sensor domains to the level of integrity required for the safe operation of UVS in the presence of people, property and other vehicles.

Almost all of the functions in the IDT (Figure 3.1) are software modules that are likely to be distributed across a number of programs and processors such that no one processor, program, programmer or user knows the full context of certain decisions. Moreover, if the UVS forms a component of a Network Centric

Warfare³⁰ (NCW) environment, it may be that software modules originating at another node in the network are executed within the UVS (or vice versa), such that any users are unaware either when or where the agents are executing or from where they originated (the agents simply cooperate with one another across platforms and operating systems autonomously).

Some added complexity comes from the fact that the IDT of an autonomous UVS is polymorphic, interpreting data from a number of different perspectives and manipulating information in accordance with properties of the UVS and the nature of the mission or problem at hand. Additionally, some defence IDT may be designed to be unpredictable so as to inject a degree of flexibility or creativity into the UVS, as predictable systems are not necessarily optimal for military operations. From a human operator's perspective, consistency is desirable, but in an adversarial context the capacity to predict exactly what the teams will do may be unhelpful. Consequently, a balance must be struck between consistency, unpredictability, and explanation that still allow the operators to understand and trust the action's of the UVS at any time.

3.2 UVS Architectures

The hardware and software architecture of UVS underpins the technology as well-defined architectures allow efficient engineering development and deployment of complete systems. Furthermore, together with the software development environment and tool sets, their selection is usually a key programmatic decision. As a result, the architecture needs to be defined at several levels of abstraction [237]:

- The basic computing infrastructure (processors and operating systems)
- The inter-application communications infra-structure and services (middleware)
- The ancillary support infrastructure (HUMS, recovery and software loading)

Open-systems architecture is almost mandatory if UVS are to be constructed rapidly from existing components. Such architectures also leverage industry-standard application programming interfaces and draw on the best civilian and defence standards for software engineering of autonomous systems and are essential if long term engineering evolution of the system is desirable. Additionally, open-systems architectures reduce non-recurring engineering (NRE) costs by allowing COTS and open-source hardware to be used for major infrastructure components such as processors, operating systems, communications stacks, and middleware [237]. Maintenance and evolutionary development costs

³⁰ Network Centric Warfare exploits concepts such as information superiority to provide a competitive edge in warfare. It encompasses “the ability to collect, process, and disseminate an uninterrupted flow of information while exploiting and/or denying an adversary's ability to do the same” [2].

are also usually reduced. However, while the use of such freely available architectures allows researchers the freedom to experiment inexpensively, it also allows potential adversaries to exploit weaknesses and vulnerabilities.

The use of COTS and/or open source hardware, operating systems and middleware also ensures a degree of hardware and software independence. Of these, the middleware layer is perhaps the most critical as it permits the definition of hardware-independent services and inter-application interfaces that enforces modular, platform and device independent, communication and computing, the integration of low-level enabling capabilities such as navigation, trajectory planning and collision avoidance, and the integration of higher-level capabilities such as mission planning, data fusion, task allocation, and coordinated control. Middleware also allows location of cooperating applications as needed to meet systems requirements such as responsiveness or redundancy and enables applications to be developed and deployed without concern for processor or network topology and technology. Additionally, changes in processing hardware or network configuration may impact the open source or middleware implementations, but the infrastructure would only need to be done once for a large system and software applications may only need to be re-compiled [237]. All of which makes the integration of real and simulated entities simpler, particularly those simulation tools that use embedded component models, abstraction and intelligent reasoning methods.

Selection of an appropriate architecture allows the implementation of formal approaches for building reliability into autonomy through the integration of prognostic methods for designing and predicting mission performance of UVS in terms of component models, platform operation, sensor effectiveness, and metrics for mission performance. It also allows verification and certification of autonomous UVS operations by allowing the implementation of mathematical, structural and algorithmic methods for modelling reliability and safety. This in turn allows principled approaches to reliability and safety in autonomous UVS to be developed, which will lead to the development of formal certification for mixed human-autonomy systems.

Taking an open-system approach to UVS architectures means that developers have the capacity to develop and test applications and systems on general purpose computers and workstations, which simplifies transition of development software from hardware-in-the-loop test and simulation environments to prototypes and technology demonstrators (see UVS Simulation and Verification & Validation). Appropriately engineered (i.e. if the test environment accurately replicates the embedded hardware), this can be as simple as copying the binary code, which allows fast spiral development software development and integration without the need for specialised embedded hardware development.

Similarly, the use of industry standard middleware allows an object-oriented component-based application development. However, as object-oriented technology is often incompatible with safety-critical systems, this is not without its problems unless specific restrictions are rigidly enforced. In this regard, such systems could benefit from the development of more formal methods that can prove software integrity through formal modelling languages and model checkers. Coupled to this problem is the need to develop a better and more fundamental

understanding of how human operators and UVS can best interact and cooperate to achieve a given task, and hence how autonomous function should be mapped to critical software architecture requirements.

All of the above addresses the key tenets of basic systems engineering philosophy; that the architecture should be designed with a purpose in mind and that the implementation of the architectural vision should be staged. Three other principles articulated by [184] also lead to hardware and software UVS infrastructure that allows efficient engineering development and deployment of complete systems:

- “The best is the enemy of the good” – UVS developers should plan on making compromises rather than being disappointed that you cannot reach Nirvana via your perfect architecture
- “Don’t throw in the kitchen sink” – the best architectures are simple and do not look as nasty as the problems they were designed to solve, and
- “A good architecture limits the number of ways in which software components may interface with one another” – because software components alone do not prevent their own inappropriate use!

Finally, a persistent UVS must be interfaced to a large number of sensing and actuation modalities so that the IDT may develop both self and situational awareness. In this regard, a successful architecture may require a significant amount of information to be input. For comparison, humans have more than 10^8 sensory inputs in each eye, 10^4 in each ear and 10^4 for taste; all of which are tolerant to sensor noise and error.

3.3 Human-UVS Interaction

Autonomous UVS that execute tasks for their military supervisors without any human intervention are fascinating, but more akin to a biological species. Furthermore, for the foreseeable future engineers will continue to develop “collaborative” UVS; that is, UVS that work with humans as a team, doing what they are asked to do. This means that we are only interested in systems that leverage human attention and ability rather than in UVS that do whatever they want whenever they want. In this regard, the requirement for UVS (or a network or ‘cooperative’ thereof) to be mobile and adaptive means that a number of systems (either component technologies or UVS) must be controlled by a relatively small number of people. This implies both a high degree of autonomy, and that humans will be retained within the decision-making cycle at some level, perhaps to initiate the use of the UVS or to assist them to extricate themselves from difficult or ambiguous situations.

For military UVS the categories of autonomy, which essentially represent increasingly sophisticated levels of intelligence (or commensurately decreasing levels of operator burden), are usually defined as [237]:

Remote control & tele-operated - A human operator controls a robotic vehicle from a distance. The human performs all of the cognitive processes. The onboard

sensors and communications enable the operator to visualise the location and movement of the platform within its environment and its onboard effectors enable the human to act on the information it provides.

Semi-autonomous - These systems have advanced navigation, obstacle avoidance, and data fusion capabilities that minimise the need for operator interaction (e.g. to achieve point-to-point mobility or target engagement). They also have sufficient on-board processing to adapt to simple changes in objective designated by an operator.

Platform-centric autonomous - A fully autonomous UVS can undertake complex tasks/missions, acquiring information from other sources as required. Alternatively, it can respond to additional commands from a controller without the need for further guidance.

Network-centric autonomous - These systems have sufficient autonomy to operate as independent nodes in a Network Centric Warfare (NCW) engagement. They should be capable of receiving information from the network, incorporating it in their mission planning and execution, and responding to other information requests, including the resolution of conflicting commands.

There are several other scales for levels of autonomy: the scale for the US Army Future Combat System (FCS) is shown below (Figure 3.2) and Sheridan & Verplank [250] describe human-computer interaction more comprehensively using ten levels (Figure 3.3). The US Air Force has also defined one based around the well-known OODA loop [52].

Level 1 (Manual Operation)
➤ The human operator directs and controls all mission functions
➤ The vehicle still maneuvers autonomously
Level 2 (Management by Consent)
➤ The system automatically recommends actions for selected functions
➤ The system prompts the operator at key points for information or decisions
➤ Today's autonomous vehicles operate at this level
Level 3 (Management by Exception)
➤ The system automatically executes mission-related functions when response times are too short for operator intervention
➤ The operator is alerted to function progress
➤ The operator may override or alter parameters and cancel or redirect actions within defined time lines
➤ Exceptions are brought to the operator's attention for decisions
Level 4 (Fully Autonomous)
➤ The system automatically executes mission-related functions when response times are too short for operator intervention
➤ The operator is alerted to function progress

Fig. 3.2 Levels of Autonomy (after [202])

Automation Level	Description of Automation Function
1	Computer offers no assistance: human takes all decisions and actions
2	Computer offers a complete set of decision/action alternatives
3	Computer narrows the selection down to a few alternatives
4	Computer suggests one alternative
5	Computer executes that suggestion if the human approves
6	Computer allows human time to veto before automatic execution
7	Computer executes automatically, then necessarily informs humans
8	Computer informs the human only if asked
9	Computer informs the human only if it (the computer) decides to
10	Computer decides everything and acts autonomously, ignoring human

Fig. 3.3 Levels of Autonomy (after Sheridan & Verplank, 1978 [250])

Unfortunately, the different focus of the groups means that the categories of one scale often correspond poorly to those in others. This is mainly because the scales attempt to categorise the problem as one-dimensional when in reality the application of autonomy is more complex and has several degrees of freedom.

A single UVS comprises two control loops: the supervisory control loop and the sensor-actuator control loop of the UVS.³¹ With the first of these, the UVS interacts with the human supervisor via the HMI. The human receives information on such things as UVS status and surroundings via the HMI, synthesises the information and provides feedback to the UVS via the control element of the HMI. If fully independent of its supervisor (i.e. Level 10 on the Sheridan & Verplank scale), the UVS uses its lower-level control loop to gather data about its environment based on its sensors and then acts solely on this information using its actuators. In general, however, the UVS combines both the supervisor's input and that from its own sensors before acting on the world, providing a natural decomposition for measuring the performance of UVS: that relating to the human-operator interaction and that relating solely to the UVS.

³¹ In reality there are more than two control loops (see Contextual Decision-Making). For example, from inner to outer a typical system might comprise: actuator control, rate control, trajectory control, mission control, and supervisory control.

Vehicles are in full collaborative ³² communication, and individual vehicle-tasking changes according to autonomously re-assigned, re-prioritised, cooperative goals. There is no human intervention.	<p style="text-align: center;">Maximum Network Autonomy</p> <p style="text-align: center;">↑</p> <p style="text-align: center;">↓</p> <p style="text-align: center;">Minimum Network Autonomy</p>
Vehicles collaborate with one another, but the human interacts by dynamically re-assigning or re-prioritising the cooperative's goals	
Vehicles communicate with one another for separation and threat deconfliction but still depend on human for new tasking	
Vehicles do not communicate with one another and follow original tasking unless human identifies a new task	

Fig. 3.4 Intra-Vehicle Levels of Autonomy (after Cummings, 2004 [74])

Cummings [74] describes five levels of intra-vehicle cooperation (Figure 3.4). As a result, when supervisors interact with multiple UVS, the nature of each human-UVS interaction is similar to the single UVS case, except in one important case: when interaction with one UVS is via another and not directly between the HMI and the UVS. Consequently, when a UVS network is (say) tasked with locating a target and the individual platforms do not have any capacity to collaborate, the levels of supervisor-UVS interaction can vary from one to ten against the Sheridan-Verplank scale. Alternatively, in the case where there is full intra-UVS collaboration, the human-UVS interactions must exist only at the higher levels of the Sheridan-Verplank scale and the network would determine the best candidate to engage the target.

As one of the main issues for the interaction of the vehicles and the human supervisor is the impact of the human decision-making process on the system performance, this duality in the levels of automation (Figure 3.5) presents a problem for the UVS designer. In single vehicle problems, there are ten discrete levels of autonomy which can allow direct comparisons of the system's overall performance to be made against one another. However, when there are networks of vehicles the problem space becomes significantly larger and more complex.³³

³² In this work coordination, cooperation, and collaboration are used rather loosely (as is often the case in robotics more generally). However, it is recognised that more formal definitions exist (e.g. [122]), where coordination implies a level of data exchange (e.g. coupling through a passive mechanism) and cooperation and collaboration imply differing levels of mutual agreement through data and action (or predicted action) exchange as well as coordination.

³³ Interaction between multiple vehicles and a human supervisor is made more complex by most existing distributed multi-robot coordination algorithms not being particularly well-suited to human interaction; most eliciting emergent behaviour such that the individual UVS follow simple coordination rules rather than teamwork models or goals. As a consequence, in complex environments the techniques usually break down because the UVS cannot explain their actions or role to other members of their team or their human supervisors.

Consequently, when designing a support system that allows the humans to interact with multiple vehicles it is necessary to assess the impact of the levels of human-UVS automation, the effects of various levels of collaboration between the UVS and the indirect influences of interaction between the automation schemes. Predominantly, this is because if the UVS mission is complex or the automation is not highly reliable, the cooperative (or even individual UVS) may perform poorer than one with no automated assistance [86].

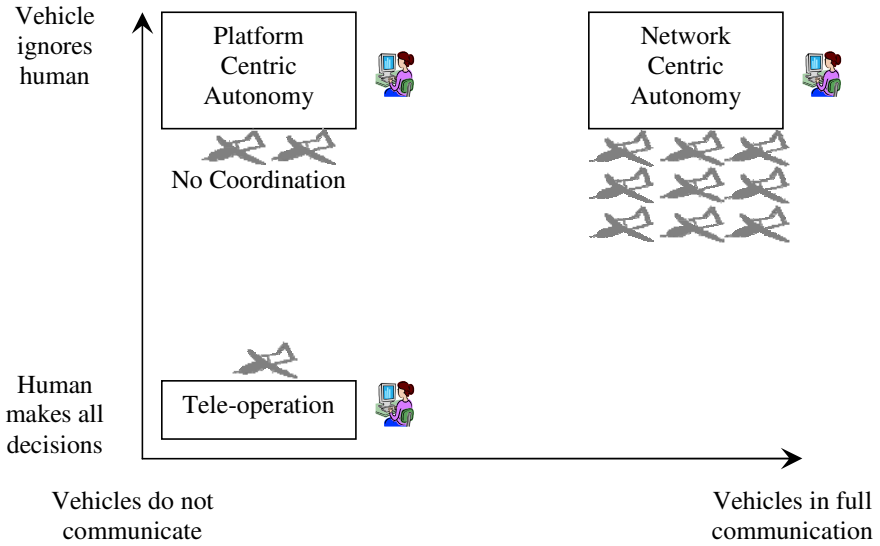


Fig. 3.5 The Duality of Autonomy

As a result, in addition to the degree of individual and cooperative autonomy, there is at least one other axis against which we need to measure autonomy: scenario or mission complexity (Figure 3.6). Scenario complexity is essentially represented by the number of mission-level activities that can be undertaken by the system, regardless of whether they are undertaken by UVS, humans, or some combination thereof. It recognises that the human-UVS enterprise is the system and measures its functional capability holistically. This axis is considerably less well-defined than those pertaining to single and multiple human-computer interaction as it is itself a function of many complex and interdependent variables, such as the degree of environmental difficulty and the complexity of the mission within this environment.

We may characterise the degree of environmental difficulty against metrics that include static elements such as “terrain” (traversability, soil, sea state, etc) and dynamic elements such as the number, density and type of objects in the environment and the frequency/rate at which they change or move. The framework might also characterise the environment in terms of luminescence/visibility and the electromagnetic spectrum, as well as operational considerations such as the presence of threats or decoys and whether the environment is rural, urban, blue water, surf, port, in-land waterway, etc. It should also take account of weather effects (e.g. [134]).

In order for a UVS to successfully achieve its mission objectives it must perform the correct actions in the right order. Consequently, the degree of mission complexity is effectively the ratio of the number of incorrect ways to perform a task relative to the number of correct ways to tackle it, where the more likely the wrong choice the higher the mission or task complexity. As a rule, in order to guarantee success, this means that the number of feasible actions a UVS can identify, select between and carry out must be at least as numerous as the number of wrong ways that the task can be undertaken. This is often referred to in complexity theory as the “rule of requisite variety”. Furthermore, while UVS are sometimes designed with unnecessary complexity, only high complexity UVS can carry out highly complex missions successfully. The complexity of a mission should not be confused with its scale, which [35] defines as the number of actions that need to be undertaken for successful completion. The scale of a UVS (or cooperative) may be thought of as the number of systems or sub-components that need to interact in a coordinated fashion. In terms of Shannon’s theorem, these concepts may be thought of as the amount of information required to describe a mission (complexity) and the resolution or level of detail in the description (scale). The notions of mission complexity and scale are revisited in the section on Multi-Vehicle Systems.

At a more simplified level, however, UVS mission complexity may be thought of as the number and nature of sub-tasks and decisions that must be made, where the sub-tasks and decisions include such things as organisational and command and control considerations, the degree of collaboration with other manned assets that may be required (which is separate to the UVS collaboration and human-independence axes), and the amount of a priori and ‘external’ situational awareness or knowledge available. Thereafter there will be high level phases, such as launch, ingress to area of operations, conduct of specific operational tasks, egress and recovery; and lower level phases, such as (for an ISR mission) follow terrain, find a target or area and search, prosecute a target, pursuit-evasion, respond to threats, etc. From the descriptive length of these phases/tasks we may then categorise the scenario complexity in notional terms such as high, medium and low.

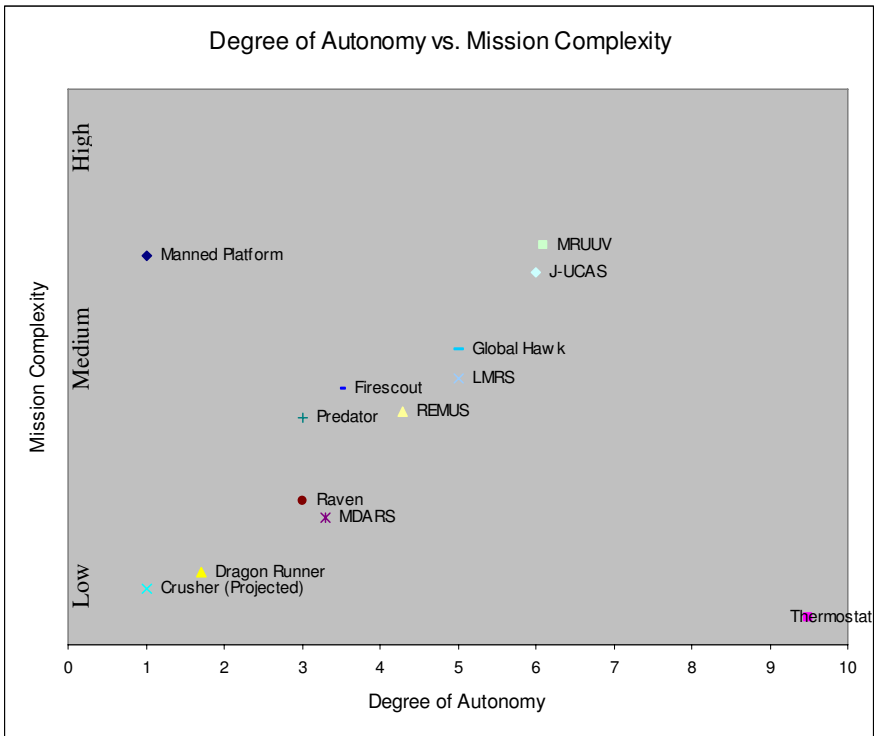


Fig. 3.6 Relationship between degree of autonomy and mission complexity Adapted from [202]

Interestingly, as [202] indicates, if we plot the degree of autonomy against mission complexity for current UVS³⁴ we find an almost 100% correspondence between the two variables (Figure 3.6). That is the more autonomous the UVS the more challenging the task that it can undertake. Furthermore, if we now define “mission autonomy” as the product of these two parameters, and plot these against mission complexity, we again find two interesting trends:³⁵ that for current UVS, higher mission autonomy corresponds to higher system complexity; and, that the degree of functional capability is broadly defined by the environmental domains in which the UVS operate (Figure 3.7).

³⁴ As there are no multi-UVS cooperatives in operations, degree of autonomy corresponds directly to single UVS autonomy; in other words, the Sheridan & Verplank scale shown in Figure 4.

³⁵ The scale of 0-10 is a notional one, with 10 being high mission complexity or high mission autonomy.

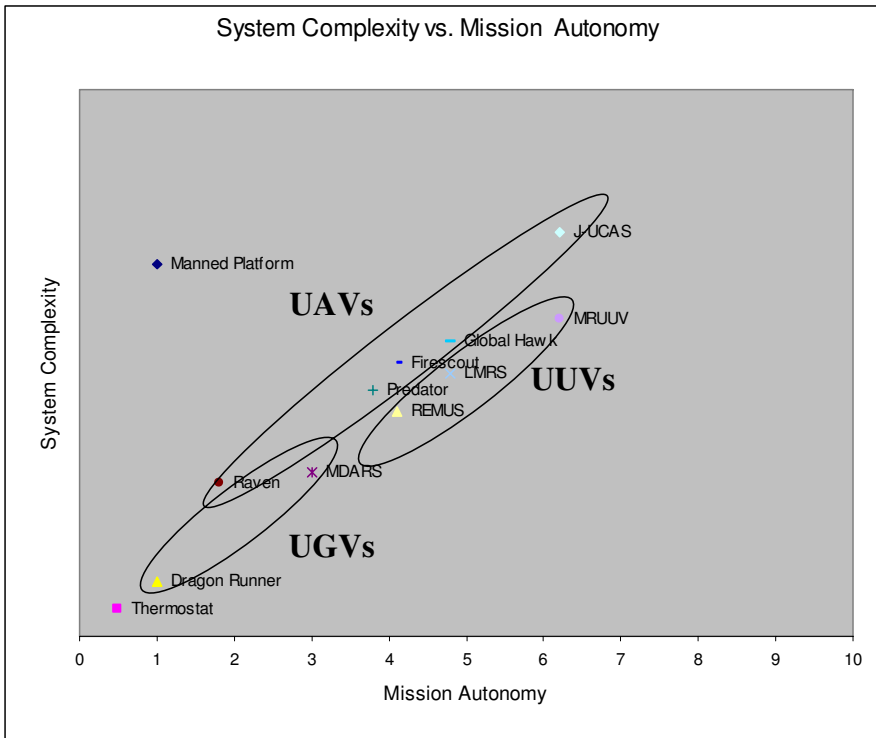


Fig. 3.7 Relationship between system complexity & mission autonomy Adapted from [202]

There are at least three reasons for this [202] [219] [220]: that the system's functional capability is distributed between the IDT onboard the UVS and the human supervisor; that this reflects the broad complexity of navigation and other operations in each of the environmental domains; and, that we anticipate more complex missions will result in less UVS reliability, which in turn requires increased human oversight and intervention.

As UVS mature and our confidence in their capabilities grow, more and more capability will be placed onboard them. To achieve this, however, we need to introduce automation primarily where it replaces the 'correct' mission or task responsibilities and presents the residual cognitive or physical tasks to humans appropriately, be they on or removed from the battlespace. The problem, of course, is identifying the key high priority tasks for complex and dynamic decision-making scenarios.

Furthermore, selection of these responsibilities is dependent upon a number of factors that include the nature of the task, operational tempo, levels of operator training, and experience [13]. Modelling and simulation tools – in particular hardware in the loop simulations – are a useful aid in the evaluation of the cognitive saturation points of the humans and the overall systems performance of the user-cooperative combination. However, application of these techniques also

require the development and definition of a framework and a set of metrics that enables the research results to be evaluated in the form of experiments with data that is quantitative or fiducially referenced. These are discussed in later sections of this work.

3.3.1 *Human-Machine Interfaces*

Interaction with UVS is enabled through their Ground Control Stations (GCS) that are usually specialised human-machine interfaces (HMI) for the particular needs of a given UVS. Most are not synonymous with human-centric computing (i.e. augmentation of human ability), although such principles should be applied when appropriate. Furthermore, over the last decade human-UVS interaction has emerged as a discipline in its own right, due mainly to the shift in the thinking on artificial intelligence away from independent action and more toward complementary action [146]. Furthermore, it addresses how users interact with single and/or networks of UVS, how any responsibilities are dynamically and contextually allocated, and how the impact of uncertainty in self and situational awareness and information overload is managed. Correctly implemented, human-UVS interfaces reduce training time and allow users to transfer their skill sets across UVS suites and environmental domains (i.e. a common interface allows users to control (say) different types of UAVs or UAVs and UGVs) [53].

At present, there are two key deficiencies in most ground station interfaces: the interfaces themselves and the control schemes they enact [76]. The current interfaces need to be more naturalistic (e.g. speech or gesture based), however this is complicated by the fact that a human will not normally share both the self and situational awareness of the UVS or be within line of sight of it. On the other hand, the challenge for the control schemes, which provide the underlying organisation and allocation of responsibilities, is that they are coupled to the interface. In other words, a badly designed interface can make it hard to use a well-designed control scheme. The complexity of the problem is increased when the cooperative of UVS is heterogeneous as the cooperative then comprises UVS with different hardware and software properties with each system providing a different local self and situation awareness perspective. It then becomes much harder for any users to monitor or predict task evolution and the ramifications of asset variability.

Teamwork also appears particularly relevant to the control of military UVS and HMI, with users reacting much more favourably to naturalistic and social interfaces. Furthermore, while many researchers are now looking at multi-UVS architectures, task and mission allocation, and the organisation of teams, significantly fewer are considering HMI that allow a single UVS within a team to work effectively with more than one user [237]. For instance, how it can schedule its tasks between them on the basis of its own sensor optimality while simultaneously avoiding user information overload?

Task allocation across an NCW environment is also important because each member of the team or NCW node can be expected to have skills that others lack. Furthermore, each may have a perspective that the other does not and it is not

clear what we need to know about environments, tasks, humans, and UVS in order to optimise task execution even if we knew what the capabilities of humans and UVS were [200]. At present, the human-UVS relationship (and hence the HMI) is usually based around some form of turn-taking behaviour, which can introduce delays and inefficiencies or even cause frustration. There is consequently a need to design HMI and UVS that work more fluently with their human partners. That is, regardless of the level of UVS automation or latency, the two are able to work together collaboratively at a higher level of performance, particularly when they are familiar with the both task and each other.

To achieve this we need to develop mathematical models, not just of the technological performance of the UVS and the human teams, but also of factors affecting manned-unmanned team performance. These factors are driven by the uneven responses exhibited by UVS in unstructured or complex environments and are typically expressed in terms of ‘trust’ or ‘confidence’ and account for a range of UVS performance variables such as risk, reliability, dependability, response accuracy, etc and human performance variables such as training, experience, workload, stress, etc (e.g. [124] [162] [259]).

Based on neurological and psychological evidence team performance is currently thought to be achieved best through the use of anticipatory perceptual simulation where principles of embodied cognition and top-down perceptual simulation are used as the basis for the development of a cognitive architecture for the human user [132]. This approach essentially works on the basis that collaborative joint action relies upon the human and the UVS having the capacity to anticipate one another’s actions. In other words, based on a model of past events, the resulting anticipatory expectation is modelled as perceptual simulation affecting a top-down notion of perception processes.

The challenges here appear to be that evaluating the relative contributions of the human and the UVS and defining or establishing similar learning curves are non-trivial. This is particularly true for evaluation techniques that rely upon self-reporting as there appear to be significant differences in the perception of the human and the UVS contributions. Also, whereas two UVS of the same origin might be expected to behave identically under similar conditions, the same cannot be said for humans – or perhaps military UVS that are intentionally ‘creative’ to avoid an adversary predicting their functional behaviour. Humans exhibit behavioural differences, both subtle and significant, implying a need to model each potential user separately. We return to this topic again in the section on Measuring Systems Performance.

3.3.2 Supervisory Workloads

In simplistic terms, the limit to the working memory of a human is seven, plus or minus two items. For dynamic working memory this number drops to two or three items [45]. As a result, we might expect that users can actively control a maximum of one or two UVS due to the sizeable number of tasks that a user typically might need to undertake:

- Navigate and control the UVS using potentially limited FOV sensors;
- Appraise the relevant situational awareness pictures (UVS and human);
- Acquire, discriminate, identify, track, locate, and designate potential targets;
- Coordinate the activities of multiple UVS dispersed across the battlespace;
- Resolve/fuse conflicting or potentially contradictory information;
- Control and aim any payloads/weapon(s) at the designated target;
- Make suitable response decisions within the Rules of Engagement (ROE);
- Arm and then fire any weapon(s), as appropriate.

The impact of the combination of the finite mental capacity of the operator and the need for the system to carry out a number of complex tasks simultaneously means that, unless multiple operators are to be used for each system and they are able to coordinate their actions, very high levels of automation are required for a range of UVS functions. For instance, a single operator might direct vehicles in a manner akin to an air traffic controller, or a single operator might act as a member of a team (perhaps the captain) in a collaborative environment, where most of the team are UVS. Irrespective of the strategy, as more UVS are added there is clearly a potential to overload the human operator.

In any system resources must be balanced against time constraints and the relative importance of the tasks that must be undertaken. In a human-UVS team the limited cognitive processing capabilities of the human operators is one such resource and there are two task intervention issues that must be addressed [87]:

- Information expected from the processors onboard the UVS has not arrived and a decision must be made (i.e. the user must act on incomplete information).
- The user has multiple dynamic tasks that vary in priority and must be balanced in priority and urgency against one another; and

In relation to the first of these, one of the main benefits of carrying out tasks using an autonomous UVS is its ability to process large amounts of information in a relatively short period of time, thereby reducing the operator's high mental workload and more optimally achieving high-level goals such as surveillance, target location, identification, engagement, etc. As a consequence, one might expect that, where such options were available, operators might choose automation over manual operation when their mental workload was high, particularly when human operators often cite excessive workload as a factor in their choice of automation. Surprisingly, there does not appear to be much evidence to support this and the subjective perceptions and objective measures are often dissociated [73]. Furthermore, other results indicate that when the users are supplied with decision aids that accurately predict future periods of high workload they fixate on attempts to globally optimise an uncertain future schedule to the detriment of solving specific, local problems. Additionally, if the UVS is highly

autonomous and only presents a filtered summary of what it considers relevant to the human supervisor, it may in fact filter out the relevant information.

In regard to the second issue, a number of studies involving target confirmation requests [95] [68] [185] [250] have concluded that the cognitive load on the operator is likely to become saturated when performing multiple tasks. These studies all used UAV teams and systems that had limited collaborative control and/or decision-aiding for the operators. However, those studies carried out on systems that have team performance monitoring and prompts to the user to intervene when trouble is detected³⁶ [165] have indicated that a user can effectively supervise much larger teams of UAVs³⁷. Some forms of interaction, such as those that specify goals as waypoints or tactical areas of interest appear largely immune to scaling problems if the HMI is appropriately designed. Other tasks, such as monitoring team performance and intervening when trouble is detected, are expected to increase in difficulty with team-size and estimates for UGVs indicate operator-to-vehicle ratios ranging between 2:1 and 1:5, depending upon mission complexity and operational tempo [202].

To avoid this, military UVS will need a scheme that identifies situations where human input is required and announces this to the operators to bring their attention to the problems at hand. Simultaneously, they must also transfer decision-making responsibility to the human. Typically, these decisions require projections into the future or the application of global judgements that are unlikely to be considered by the reactive and/or localised nature of a decentralised process. There are likely to be three problems [95]:

- Tasks that have not been unallocated. Under these circumstances, the autonomous tasking can be suspended (temporarily or permanently) or cancelled and the process handed to humans for a period.
- Team members that have not been tasked. When a team member does not have a task for an extended period the UVS can be forced to join a team, do nothing, moved to some convenient location pending further instructions, or the human take over control.
- Unusual performance relative to a task plan. Most tasks have logical conditions and metrics that indicate when the trajectory or plan has become unachievable or irrelevant. The mission plans need to identify such conditions a priori and the attention of the supervisor drawn to any performance criteria that are not satisfied (e.g. a waypoint may be missed or the time taken exceeds some threshold). A decision whether to cancel the plan or allow it to continue must then be made.

Moreover, because a human may not be able to respond to some prompts in a timely manner, and the value of the decision may be lessened, mathematical models of the transfer of control that capture the increasing appropriateness of (say) terminating a long-running plan may need to be employed [242]. In this

³⁶ The scheme simply identified situations where human input might be needed and then explicitly transfers responsibility for decision-making to the human.

³⁷ These experiments involved simulations with teams of eighty UAVs.

regard, studies commonly find large individual differences in the selection patterns of automation – based on a wide range of factors such as fatigue, training, vocation, age, etc.

The limitations in the working and long-term memory of UVS supervisors, however, means that they tend to rely upon mental models, schemas, scripts and heuristics, with heuristics based on prior experience providing a particularly high success rate [110]. When confronted with problems not framed within their experiential context, humans have the capacity to generalise and transfer these mental models to the new contexts. As a result it will continue to be necessary to retain humans within the decision-making cycle to assist the UVS to extricate itself from the most difficult situations.

Additionally, while higher levels of autonomy usually provide the best solutions for rigid tasks that require limited flexibility in decision-making and with a low probability of system failure [94], higher levels of automation are not usually so successful for systems that must deal with dynamic environments that might have many external and changing constraints [268]. This is because of their inability to reliably expedite the decision-making process in the face of uncertainty or unforeseen problems and it is predominantly this – and the human capacity to re-frame the context – that drives the need to include a human in the supervisory loop.

As a result trust takes on added significance within multi-UVS cooperatives as it declines rapidly with increasing numbers of vehicles if the reliability of the autonomous control decreases [87]. This lack of trust causes an increase in the cognitive workload of the users and as a consequence, the reliability of the automation is one of the key drivers that determine the capacity of human-to-vehicle ratios. That said when operators are involved in the planning and execution of decisions trust increases.

Ultimately, in order to collaborate fully in an ongoing sense the UVS and the humans must share a common view of each other's goals and sub-goals. This requires them to share a common world view, but not necessarily all of the information that each observes. For example, it is generally acknowledged that even a small amount of information-transfer provides great benefit for certain tasks. Consequently, we need to measure the effect of different communications strategies on the systems performance of UVS, such that we focus on the development of dedicated languages or protocols and the grounding of these representations in the physical world.

A number of groups are studying how fault-tolerance within networks of UVS can be improved, particularly in relation to maintaining situational awareness when connections between UVS change dynamically or unexpectedly. Similarly, other groups are focusing on improving the systems performance of multi-UVS teams that, due to their size, rely on low capacity communications links. For instance, how each team member could subscribe to the total pool of information available such that each is presented with only sufficient data to achieve its own goals; a notion popularised by Milgram's experiment [191]. In this experiment Milgram required an incompletely addressed letter to be passed from acquaintance to acquaintance within a network the size of the United States until it arrived at its

correct destination. The letter, addressed to “A Pittsburgh botanist who plays the flute”, started in Los Angeles and progressively approached the correct recipient on average in six steps.

In the UVS context, this would mean that the individual (and geographically distributed) components of the system must each have sufficient – but not complete – knowledge of the role of their neighbours,³⁸ which they must determine (potentially in real time) on the basis their global knowledge of the mission and the overarching goals of the cooperative. In colloquial terms, while the individuals may each understand that they are only stonemasons, carpenters, plumbers, electricians, and so on, they must also all understand that they are building a cathedral and have adequate corporate knowledge of and trust in their fellow workers. In essence this is placing trust in ‘the system’ (i.e. the architecture of the cooperative and a sub-set of its individual components) rather than in the mutual, cooperative and global planning capabilities of the UVS. In this regard, [18] indicates that in order for such control to become a reality we may need the ability to model or synthesise (and hence measure) emotions,³⁹ as neurobiology has shown that emotions communicate simplified but high impact information (using the vocabulary of neuro-modulation) between key areas of the interacting subjects’ brains.

Superficially, such techniques are attractive as reduction in communication requirements leads to reduction in systems latency. However, a reduction in the amount of information that needs to be passed between two entities is also a measure of their mutual levels of automation and hence complexity. Furthermore, it is not always either straightforward or effective to simply reduce the amount of information transferred between two entities and still maintain a high degree of synchronism as many tasks are difficult to automate reliably. For example, the decisions may involve the application of contextual reasoning or judgement, which in a military context may mean that they are linked to the execution of higher order command and control obligations. Furthermore, for autonomous military UVS, it is not clear whether such systems would be more or less vulnerable to adversarial exploitation. More research is required.

³⁸ In this case, the term “neighbour” refers to the next appropriate interactive component (i.e. architectural neighbour), which may not necessarily be geographically proximate. Furthermore, these neighbours may change discretely or continuously over time.

³⁹ This is an area known as ‘affective computing’ – computing that relates to, arises from or influences human emotion [227] – which is beyond the scope of this book.

Chapter 4

UVS Technology Issues

The aspirational capability objectives for UVS in the air, land and maritime environments articulated by several national roadmap documents (e.g. [202] [205] [206] [267] [268] [278] [281]) may effectively be expressed as: 2008-10 conduct of ISR missions; 2015-2020 autonomous patrol; and, 2025-2030 strike capability and combat missions. This implies a need for persistent UVS autonomy⁴⁰ in complex dynamic military environments and the automation of a range of higher order or ‘intelligent’ functions. As a result, the challenges discussed in this chapter pertain mainly to next-generation UVS and pre-empt any ‘validated’ current military requirements. The chapter focuses on the complexities of contextual decision-making, planning in dynamic environments, verification and validation, sensory deprivation and trust and reliability for autonomous military UVS. It is recognised that there are a large number of other technological challenges⁴¹ that may result in improved UVS, however, these have been covered by a range of other teams in studies cataloguing the state-of-the-art and projected requirements in each of these fields against various likely missions and applications (e.g. [44] [141] [202] [204] [205] [206] [230] [234] [237] [244] [245] [268] [276] [277] [278] [279] [280] [281] [282]).

4.1 Technology Challenges

UVS are subject to the laws of physics: they have mass and inertia, their moving parts wear, their electrical components emit heat, their sensors are corrupted by noise, no two systems are exactly alike, they fail and the environments into which militaries place them are complex, dynamic and unstructured. As a result, we cannot accurately predict their behaviour in advance and there are several factors that currently limit them from achieving their full potential. For example:

⁴⁰ It should be noted that requirements such as ‘persistent surveillance’ specified in several of the same planning documents differs from ‘persistent autonomy’ as the former can be carried out by UVS in combination with humans, whereas the latter requires complete independence.

⁴¹ Including communications, sensing, signal processing, data and information fusion, systems engineering and integration, launch and recovery, human factors, platform, aero/hydro-dynamics, mobility, collision avoidance, mission planning/re-planning, propulsion, size, and energy storage.

The projected affordability of future autonomous and unmanned systems is higher than it needs to be. UVS do not require humans to be onboard and consequently do not need life support systems, space for the humans, special armour or protection, etc. As a result, UVS can theoretically be made smaller and lighter than their manned counterparts. As the procurement cost of vehicles is roughly proportional their mass (about US\$3,300/kg) [141] [268] [278] a reduction in mass can be expected to translate into cost savings and a commensurate drop in the support required for the vehicle. However, current experience indicates that UVS and automation do yet exhibit a level of savings (say) enjoyed by computers. There is, of course, every reason to believe that as the technology matures, the costs will start to fall in line with the trends shown by manned vehicles, although further research into the hidden costs of operating UVS is also needed. For instance, at present, multiple operators are required for even single systems: a PackBot UGV may only require one operator, but it requires two people to transport it into the field and several more to protect the user. Similarly, Predator, Global Hawk and many other operational UAVs require a small number of people to operate the sensors and fly the vehicle, but significantly larger numbers to support it (planning, maintenance, image analysis, and so on).

Related to cost are issues of UVS survivability. That is, to survive, a UVS must detect, identify, classify, plan and respond to a threat. If a high degree of automation is employed and the sensing modality mixed, the cost of designing and producing such a UVS will likely be high (and the UVS may also be physically large). It is then more difficult to justify the sacrificial use of such systems. Alternatively, if the system is inexpensive a human will likely be required to interpret some or all of the sensed data. As a result, response options will be delayed and the likelihood that the vehicle will be lost increases. The probability of losing UVS is also linked to the reliability of sub-systems, which tends to decrease with an increase in system complexity. A second order effect is that as system complexity increases so does the level of mission complexity in which the UVS is employed, which increases the likelihood of losing the asset. Another impact of increased autonomy is a reduction in the communications signature, which allows more covert operations to be undertaken, although this is often offset by an increase in the sensing modality required, which may lead to an increase in other signatures such as radar cross section or an emission signature in another band.

Linked to cost and survivability are issues of affordability: the more survivable the UVS the fewer that need to be acquired. Although this needs to be balanced against the likely attrition rate that will come with the increased mission complexity and threat exposure that they will experience. Similarly, although an increase in the level of UVS autonomy theoretically leads to a reduction in the number of human supervisors required to operate it, there may be an increase in the maintenance and training requirements.

Another key technology issue at present is that many EO sensors are able to detect almost all of the light entering the camera aperture, with sensor noise near the lower limits set by the laws of physics. Thus, the challenges for these cameras lie not in improving the sensitivity of the sensors to light, but in increasing the size

of their imaging array and hence the capacity of sensors to have sufficiently broad fields of view and resolution to allow detection of entities at long ranges so that sophisticated image interpretation techniques can perceive and ‘understand’ the key elements in their environments.

These techniques and sensors then need to be combined with all-source data fusion and advanced machine learning or adaptation techniques to make the perception more robust and insensitive to environmental variations. This would allow mission and path planning beyond a platform’s organic sensor range and greater persistence in the battlespace through the provision of continuous, all-weather, 3-D terrain and target classification, mapping and localisation. Such improvements would also allow detection, recognition and interpretation of human, vehicle and other threat activity such that the UVS could distinguish friend from foe and anticipated versus unanticipated movement, thereby improving survivability and their capacity to operate in shared environments.

UVS also need IDT capable of making plans relative to a leader, manned vehicle or environmental changes so that they can adjust their resource usage or properties, join or leave teams (for example relative to communications, sensor scheduling, surveillance points, target kill, etc). Similarly, we need IDT that allow UVS to independently identify and make intelligent, complex operational and tactical decisions (e.g. self-concealment, lethal or non-lethal self-protection, avoidance of threats, and mimic leader action).

Linked to this is the need for communications and image compression technologies to be developed that allow beyond line of sight (BLOS) transfer of high resolution imagery and sonar data between UVS and their manned counterparts. For example, the most capable underwater systems currently achieve around 10kbps, but have a high signature for detection. Alternatively, other, LPI systems are typically able to communicate at rates of only 2-3kbps over ranges of around 10km, although using larger arrays or techniques that predict propagation conditions longer ranges and larger bandwidths are possible. Similarly, while RF propagates freely in the earth’s atmosphere up to about 100GHz and it is possible for small directional antennas (~ 20cm) to be combined with low power (1W) amplifiers and then used to exchange data at rates approaching 10Gbits/sec between a UAV⁴² and its GCS more than 100km away, these systems are still far too heavy and large to be of use to UAVs in the small-medium sized class.

Launch and recovery of UAVs and UUVs from ships is also a major issue. Fortunately, many longer endurance UAVs fly slowly so they can take off and land at speeds similar to those of ships at sea. As a result, it is necessary only to contribute to or absorb a small amount of the UAV’s energy if the vehicles are appropriately aligned. The same is true for UUVs, although many do not travel fast enough to keep up with operations at sea. Furthermore, most ships do not want to wait for sea state zero (or stop) before launching or recovering a UVS, and recovery of any UVS at sea is a hazardous undertaking – even if it were only damage to the UVS that were being considered.

⁴² This theoretical data rate will be reduced by a factor of up to 100 if anti-jam protection is afforded.

There is also a need to establish a clear product certification process for UVS that includes safety cases and regulatory regimes that address the very real dangers and the issues of public perception. Similarly, the application of autonomy to weaponisation and automatic target detection and recognition also needs to be addressed (this includes the related architectural designs). Related to both these issues are the use and safe manoeuvre of such UVS in the presence of people and other vehicles; the use and application of UVS within a human command and control network that changes; the level and modality of interoperability between different UVS and their control stations; developing flexibility in the levels of automation and adaptive interfaces; optimisation of vehicle-to-operator ratio for manned-unmanned collaboration; and, development of adaptive knowledge management systems for UVS.

All these and many other deficiencies relating to component technologies⁴³ of UVS are largely responsible for UVS not yet providing a persistent presence on our battlefields. However, as a number of highly qualified teams have published studies cataloguing the state-of-the-art and likely requirements in each of these fields against various capability projections, likely missions, and potential applications, and a full catalogue of the spectrum of technological challenges currently faced by UVS developers and programmers is simply beyond the scope of this book, this chapter focuses on the higher-order functions required to instantiate persistence rather than the physical ones.

- **Human-UVS Interaction:** UVS currently lack the ability to interact with humans and other UVS in an efficient and naturalistic manner that enables the human-vehicle system to perform a full range of complex tasks in unstructured environments. This is largely covered in the previous chapter, but discussed throughout this section.
- **Contextual Decision-Making:** Metrics for good decision-making, particularly for a context unspecified at mission commencement, are usually poorly defined. Understanding the basic patterns of stability and predictability for the decision-making paradigms is a pre-requisite for robust autonomy.
- **Verification & Validation:** The integrated and polymorphic nature of the sub-systems that make up a UVS combined with the requirements for stand-alone operation in a broad spectrum of unpredictable environments, which may be critical to mission success means that V&V, poses a significant challenge.
- **Trust & Reliability:** Trust and reliability are key issues that drive the levels of confidence and autonomy that we place in UVS. Currently UVS lack the capacity to understand their state such that they can predict their performance or detect functional or component failures autonomously, which affects our trust in them.

⁴³ For example, communications, sensing, signal processing, data/information fusion, systems integration, launch and recovery, human factors, platform, aero/hydro-dynamics, mobility, collision avoidance, mission planning/re-planning, propulsion, size, and energy storage.

- **Persistence:** UVS need to achieve improved performance over time, particularly in regard to repeated operations in the same environment, while learning from their experiences.
- **Dynamic Environments:** UVS frequently lack the ability to detect, locate and track moving objects while simultaneously accounting for longer-term changes in their environments.
- **Sensory Deprivation:** UVS perceive their environment through a limited sensory perspective, which may “blind” their supervisors; or force them to attend to the demands of a particularly burdensome task.
- **Robustness:** UVS lack robustness in the systems integration of their functional components and in the reliability of the system in dynamic environments. As a result, at some level UVS will malfunction and we are unlikely to be able to predict the specific nature or timing of these failures. Furthermore, UVS frequently fail, not through a manufacturing or design flaw but through routine dynamic loading, collision with an obstacle, or operators using them beyond their design limits. The operational need for cannibalisation of parts and specialist support has potential force implications.

4.2 Contextual Decision-Making

Reflecting on a widely used definition of intelligence [4], “the ability of a system to act appropriately in an uncertain environment, where appropriate action is that which increases the probability of success, and success is the achievement of behavioural sub-goals that supports the system’s ultimate aim,” we can see that intelligent autonomy is conceived within the context of a UVS within its environment rather than independent of it. As a result there are three aspects associated with testing such autonomous behaviour [6]: novelty in the environment or in the problem to be solved, uncertainty regarding what is to happen, and dealing with difficult situations.

In this regard, the fundamental building block of good decision-making for automation is a high degree of Situational Awareness (SA),⁴⁴ where SA is defined as having three levels [96]: perception of elements in the environment; comprehension of the current situation; and projection of the future status. Issues for each component of SA include:

Perception - Humans rely upon their five senses (or combinations thereof) to perceive their environment across the application domains. Their degree of success is often linked to their capacity to “notice things” while other events are unfolding. Many UVS ignore certain events as they are programmed to detect or interpret only particular ones.

⁴⁴ Here, for convenience, we include self awareness within the definition of situational awareness, although we shall return to discuss self-awareness in more detail in a later section (where we separate the two).

Comprehension - Humans comprehend a situation by fusing their environmental perceptions with relevant contextual information and mission goals. Most UVS rely upon their supervisors to prioritise the importance and meaning of information, except possibly in regard to aspects of their navigational aspirations. For instance, some autonomous UGV navigation systems have the capacity to compute solutions for almost every environmental situation [243].⁴⁵

Projection - Humans make predictions on the basis of their perception and comprehension of a situation. Projection is frequently the most highly demanding cognitive activity and various stressors (cognitive workload, fatigue, stress, etc) can affect a human's capacity to fulfil this high-level task. Appropriate automation might ease this burden.

Unfortunately, novelty and difficult situations may be indistinguishable to an autonomous UVS. As a result, humans and the UVS may need to share their individual perceptions of the environment by developing and maintaining a common situational awareness picture. Consequently, the UVS information must be filtered, manipulated, and then presented in such a way that a user can quickly assess the status of the UVS (or the cooperative) and the battlespace it observes. If the mental resources required to accomplish this exceed the task demand, system performance will remain above the required threshold. In a high workload environment, when the demand imposed by competing tasks exceeds a user's capacity to process information, performance can be expected to suffer⁴⁶ [110].

To this end, automation needs to be introduced primarily where it replaces the difficult or complex UVS task responsibilities and presents the residual cognitive or physical tasks to operators appropriately. The problem, of course, is identifying the difficult high priority tasks for what is a dynamic decision-making environment. Furthermore, this information must be collected, processed, stored, and disseminated appropriately to those who need it, whatever their geographic location. Additionally, the selection of these responsibilities is dependent upon a number of factors that include the nature and complexity of the task, operational tempo, levels of operator training, experience, and so on.

In this regard, it is well-known that situational awareness has an effect on humans' abilities to successfully complete missions [95]. However, at present most attempts at improvements in human situation awareness focus on providing better interfaces between the UVS and its supervisor, allowing the human to carry

⁴⁵ If a UGV needs to traverse complex terrain a solution may or may not exist depending upon the width and mobility characteristics of the UGV. Alternatively, even if the terrain is traversable, the ease with which the UGV is able to execute its solution may vary.

⁴⁶ There is also a predicted drop-off in performance for low workload environments [295].

out the processes of determining his situation awareness better rather than capturing the machine's ability to observe, comprehend or make predictions (i.e. enhancing the UVS' ability to develop its own self and situation awareness and indirectly and simultaneously enhancing that of the user).

For decision-making to be distributed between the UVS and the human, a high degree of shared situational awareness is required. In a manned environment the devices that deliver shared situational awareness include spoken and non-verbal communications, visual and audio shared displays, and a shared environment [166]. Unfortunately, the bulk of these delivery mechanisms are not yet viable for a UVS and the sensed data must be pre-processed to convert it to a common reference frame, fused with state predictions based on historical observations, transmitted through communications interfaces, assimilated with other sensed data that have passed through a similar process to that described here, and then represented visually for interpretation and use by the cooperative's supervisor.

The degree of system automation required is fundamentally defined by the relationship between the human resource supplied and the situational awareness task demanded [251]. In this regard we must take account of several factors pertaining to a human's capacity to appraise his situation [95], including the limited cognitive processing capabilities of the supervisors. Humans are able to divide, direct and select their attention capabilities, but their perception is limited by their capacity to parallel process sensory events, sensor modality and working memory constraints and by their sensory channels. Consequently, complex or dynamic environments can quickly overload a human's attentive abilities such that they selectively sample their sensory channels. As a result, they typically manage their attention focus based on events, sensory updates, environmental conditions or task dynamics. Given that UVS frequently "ignore" information that would cause their supervisors to re-direct their attention, managing the attention requirements of supervisors such that they optimally sense and understand their environment is critical for environmental perception.

Although significant advances have been made in this area [90] [160] [197] [266], most solutions treat the task allocation, decentralised data fusion, and sensor scheduling problems independently. For instance, the effective allocation of a particular UAV within a team at any instant may depend upon sensor scheduling constraints imposed upon the payload. Moreover, as the number of UVS in the cooperative increases determining the required behaviour becomes more computationally intensive and complex. Similarly, emergent behaviour and unforeseen circumstances also become more common [166].

This means that due to the difficulty of forecasting the (probably emergent) behaviour of UVS, particularly within a networked or NCW environment, it may be very difficult to detect that something is going wrong. Consequently, another issue is how to provide diagnostic and feedback support to the UVS supervisors and their commanders, who may themselves be distributed over a wide geographic

area, particularly as many supervisory functions are cognitive, hard to monitor, and embedded as components of other operations. As a result, rather than being able to monitor the individual tasks directly, the commander or supervisor may only be able to assess the systems outcome (i.e. the result of the autonomous cooperative's action).

In the context of defined tasks such as detecting and identifying targets, controlling and aiming a weapon, landing an aircraft and so on it is easy to understand what we mean by the phrase 'good' relative to an autonomous UVS; it is measured against the specific purposes of the designers and users. When we consider how to quantify the effectiveness an autonomous UVS making decisions between fulfilling a mission objective set by its commander and delaying achievement of this goal to satisfy other objectives (e.g. attacking an adversary en route to this objective) it is much more difficult to understand what we mean by 'good' as the metrics for such decision-making – particularly those pertinent to a context undefined at mission commencement – are usually poorly defined. The challenge for autonomous UVS is therefore as much based around the theory of work organisation as it is technical in nature; only after the basic patterns of stability and predictability have been thought through can UVS be productively applied.

For instance, a UVS control system needs to perform three basic tasks: avoid obstacles; avoid other UVS; and, operate the UVS within its performance envelope. Once these priority tasks have been accommodated, higher order tasks such as mission planning, surveillance, reconnaissance, target location, sensor scheduling, coordination, communication, etc. may then be undertaken. As stand-alone actions, the priority tasks are accommodated relatively easily as their goals are both decomposable and quantifiable in terms of physical quantities and closed loop control laws that relate to physical parameters such as lift, drag, thrust and so on.




Fig. 4.1 Approach of linking perception to action through cognition (Adapted from [58])

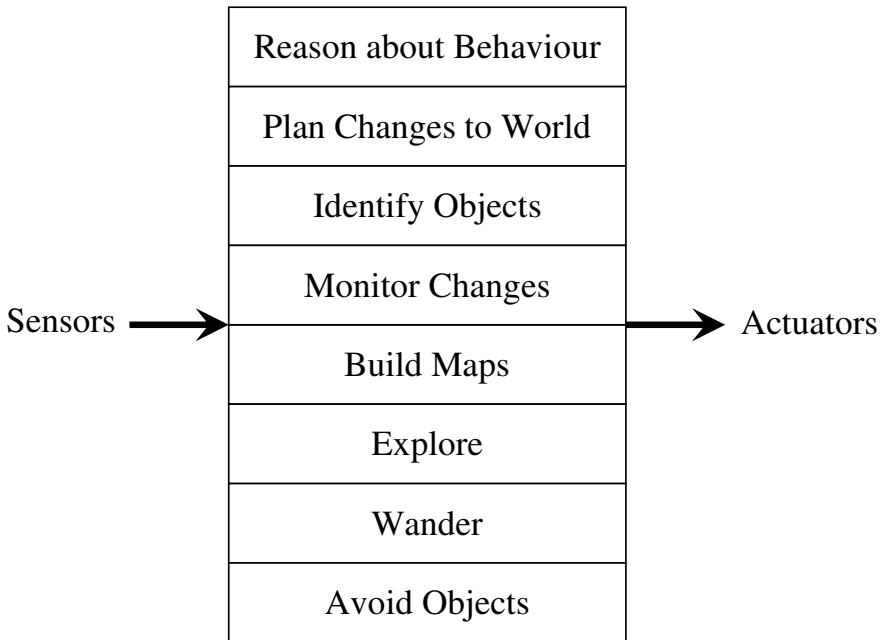


Fig. 4.2 Alternative decomposition of the perception-actuation problem (Adapted from [58])

Unfortunately, autonomous UVS also require the priority tasks to be closely integrated with the higher order tasks and there are at least two structural concepts for accommodating this integration. One in which perception is linked to action through cognition, where higher-level reasoning operates on the output of sensor-based perception to provide the necessary motion planning for actuation (see Figure 4.1); and another, bottom-up behaviour-based strategy where perception and actuation are more directly linked, without the need for detailed intermediate world models that relate one to the other (see Figure 4.2) [58].

There are several reasons why these higher order tasks are hard to instantiate:

- Higher order tasks are frequently more difficult for both humans and UVS to deal with as the decisions that they involve rely upon reasoning and judgement that are linked to the execution of higher order military command and control obligations.
- The problems are usually complex (for humans), which means that the problem is less well-understood and less structured and therefore harder to analyse or decompose into definable components. This means there is often a high likelihood of ambiguity, multiple possible courses of action, and/or the likelihood that one decision will impact a subsequent one.
- There are usually significant amounts of uncertainty, possibly conflict, both in terms of what is known a priori to any mission and what is observed during it. This means that determining optimality in terms of

any decision-making – that is selecting a course of action that has the highest probability of meeting any defined objectives – is much harder to compute with a high degree of reliability.

- It is difficult to accurately define suitable cost functions and metrics by which we quantify the benefit of applying one strategy or course of action over another. This adds to the complexity of the problem because we evaluate both the decision trade-offs and the quality of our decisions based on these.
- In addition to any ‘decision cost’ (the determination of one course of action over another), there is also an opportunity cost (the determination of a known course of action over one not considered), which is usually unknown or in-calculable.
- The prioritisation of goals is usually a subjective, contextual or interpretative task, for which it is not possible to anticipate all possible decisions and circumstances. Moreover, these tasks are often derived from direct or implied command and control strategies.
- From time-to-time it is necessary to re-frame an existing problem rather than interpreting the situation within the existing problem frame. Recognising this and dealing effectively with complex novelty and maintaining information regarding any decisions we may wish to make in this space (and in a form that is readily usable when we need it) is extremely challenging.
- Evolution has played a significant part in the development of human intelligence and its adaptation to the tasks to which it is suited.
- Other challenges include constraints imposed upon the problem space, such as time and mission constraints/obligations, mission complexity, and/or supervisory, environmental or adversarial interventions.

As [58] indicates, good decision-making agents reduce the complexity of the executive decision-making by breaking it down into component-decisions that are simpler to make. This also enables us to put greater structure into the verification problem (see Validation & Verification later) and to act on the sub-decisions and the information pertinent to each decision. It also aids in reducing the time it takes to make decisions by saving the UVS controller or the human supervisor from having to determine which information sources or sensors are relevant to any particular decision.

It may also result in information being presented in a more organised fashion, thus saving the operator valuable time when interpreting what might otherwise be a complex display. For example, if a controller or supervisor is only presented with sub-decisions that are possible at any instant, in-feasible decisions are automatically eliminated. If we then order the sub-decisions according to the hierarchy in which they have been determined we also have a mechanism for achieving traceability. In this way we can ‘walk’ a UVS operator through any decisions he needs to make while simultaneously recording any decisions made. This also allows interruptions (e.g. communications outages) to occur by providing continuity for algorithms and supervisors that need to bridge any such outages as they can return to the last relevant decision at the end of any outage.

4.2.1 *Planning in Dynamic Environments*

Military environments are inherently dynamic and if UVS are to maintain a persistent presence on the battlefield they must be able to adapt to their changing and often adversarial nature. In this autonomous UVS are no different from any other entities in the battlespace: they need to know what is going to happen next and what the best decision is now. Consequently, UVS require strategies not only for decomposing their missions into meaningful sub-tasks, but also for tracking progress towards mission goals and the changing nature of these tasks relative to the capabilities of the UVS. To do this they need to make plans and to establish a trade-off between the cost of any new plan (ideally compared to some global optimum) and the reaction time required to modify or repair any original plan given new information.

In this regard, there are two major steps involved in translating the problem-context into a solution framework, specification of the planning model and its evaluation function. In other words, the process of dynamic re-planning requires the creation of a model of the problem and use of that model to compute a solution. Consequently, when we solve a planning problem (or repair or modify a plan) we are actually only finding a solution to an approximation of the planning problem, which is a model of the real world. That is, we must find an incomplete solution to a problem that accurately represents physical processes, or a complete solution to a reduced (i.e. a simpler) problem [187].

The (processor-hungry) solution to this problem is to treat it as one of dynamic, constrained optimisation set in a time-varying environment and continuously recompute and execute plans over some multi-objective cost function. Unfortunately, there are limits to the processing capability that most UVS can carry, tasks are time-constrained, the constraints and the solutions to the cost functions are typically only those that provide good approximations to the parameters under consideration, and the optimisation is often application-specific and depends upon the real world variables being optimised. As a result, it is usually preferable to find an approximate solution to a precise model rather than a precise solution to an approximate model. This is because if our model has a high degree of fidelity we can have confidence that the solution will be meaningful [188].

A decision must also be made to structure the planning algorithms either as complete or partial solutions.⁴⁷ The computational advantages of using partial solutions are attractive, but they are not without difficulties. For example, the problem must be organised so that the component problems can be optimised efficiently and another evaluation function is required by which the relative value of the partial solutions can be determined. Furthermore, if the process is interrupted partial solution algorithms may not provide feasible planning strategies, whereas complete solution approaches should always be able to provide

⁴⁷ In complete solutions, all decision variables are specified and evaluation takes place by comparing two complete solutions; better ones replacing previous ones. In incomplete solutions, a complex problem is simplified by decomposing it into smaller, discrete problems that are easier to solve. When the partial solutions are all solved, they may then be combined and used as building blocks for the solution of the original problem.

at least one feasible plan. Unfortunately, while usually relatively simple to implement, complete solutions tend to be computationally expensive, as their main requirement is that the problem space must be evaluated exhaustively [9].

In this regard, there are many traditional approaches that can be applied to the problem, mainly because none is particularly robust to the broader problem space. In other words, if the problem space changes, so must the technique. As a result, many practical planning techniques tend to combine the benefits of reactive (local) and deliberative (exhaustive)⁴⁸ techniques, creating hierarchal systems that engage the low-level reactive planners under higher-level deliberative ones (or parallelised versions of the same thing). Depending upon circumstance, however, some scenarios are better serviced by deliberative strategies that execute closer to global optimality, whereas others achieve mission success faster using the reactive ones [9]. As a result, reliance on either choice can be poor in certain circumstances.

At a most basic level, a planner (e.g. for navigation) should prescribe a solution that is longer than the reaction distance of the UVS so that when new trajectories are computed it can avoid obstacles and other users of its environment. If this is achievable and the sensors onboard the UVS have sufficient range and resolution to perceive this environment then the UVS will at least operate safely within its environment while the higher level plans are computed. There are analogous safe planning solutions for other UVS behaviours such as weapons control, sensor scheduling and communications. For example, in regard to autonomous UVS weapons, we can postulate an acceptable generic architecture based on the major Principles of the Law of Armed Conflict (i.e. responsibility, military necessity, target discrimination and proportionality). This is discussed in greater depth in the section of Legal Issues, and in particular in the section on the Ethical Control of UVS.

A fundamental operating condition for most military UVS is that, once operating, most UVS systems cannot simply stop to compute a new plan every time the environment or circumstances change. Consequently, planning must be performed concurrently with normal system operation. There are several requirements:

- Robust plans are required to minimise the frequency with which successive calls are made to the planner
- When a call must be made to the planner, the repaired plan should only differ from the original plan by a limited amount
- In order to accommodate any limited deviation, the original plan should be readily adaptable to likely changes in the environment or mission
- When a call must be made to the planner, adapting previous plans and/or making completely new ones should take as little time as possible
- Despite the implied time pressures, the plans should be of a consistent quality

⁴⁸ Reactive techniques consider only recent and/or current information and produce local or one step-ahead strategies based on conditional responses. Deliberative techniques derive their recommendations based on all available information and strive for global optimality or a complete mission plan.

It is, of course, difficult to know a priori what any update rate pertinent to the dynamic decision-making timescale should be. Additionally, therefore, we will need to incorporate a degree of adaptive or reinforced learning into the prediction component of the planning algorithms to allow them to determine their own update requirements. That is, we will want them to have the capacity to learn from task and environmental changes in order to accommodate a better sampling frequency of sensor inputs and prediction outputs.

In this regard, evolutionary algorithms, which are essentially an adaptive combination of many techniques, show considerable promise. Regardless of the technique however most will compare newly generated solutions to existing ones and make some determination as to which solutions are feasible and/or preferred, and which need to be pruned or retained for further processing. This pruning process, however, is largely based on the evaluation function. Furthermore, it is usually assumed that the evaluation function is well-defined, whereas in reality problems are often set in noisy or uncertain environments. A key challenge in this regard, however, is whether to use an existing plan that is known to be sub-optimal or to wait for a better solution to be computed (the obvious solution – ‘wait and see’ – can lead to planning discontinuities).

Furthermore, each element of a planner must also detect that it has failed and inform any other components. These requirements are strongly linked to the desire to achieve persistent autonomy; that is, for the UVS to be able to determine conditions under which the prescribed mission tasks are unachievable, either within a required time frame or the broader capability framework of the UVS. By having this level of self-awareness, and notifying users of such limitations, the human-UVS system can then adapt accordingly.

For persistent autonomy, it is a fundamental requirement that the UVS be able to provide feasible solutions, and hence recognise those that are infeasible relative to mission time constraints, its own capabilities, etc. If the UVS determines to prosecute an infeasible plan, it has not really found a solution to its problems. That said it is acceptable for the UVS planner to work in infeasible space, defining solutions that it cannot achieve in order to determine those that it can perform. In this regard, however, there are some challenges for developers [187]:

- How do we compare infeasible plans
- Should we use an evaluation function for the feasible or the infeasible plans
- Are (or should) these two evaluation functions be related to one another
- Should we simply eliminate infeasible plans or attempt to repair them
- If we attempt to repair the plans, should we “move” them by the least amount
- Alternatively, is more radical “surgery” appropriate (i.e. more feasible solution)
- Do we need to find the delineation between feasible and infeasible plans
- Should we extract a set of constraints that define the feasible/infeasible boundary
- Having determined the infeasible plans, how do we translate them to feasible ones

As implied above, most planners cast their predictions as a binary problem for which the solution is either feasible or infeasible. Relative to the capabilities of the UVS, however, there may be areas of grey where solutions are simply difficult rather than impossible, particularly in relation to UUV operations near the sea floor or UGVs attempting off-road navigation. For instance, a bridge that is too narrow for a UGV might be considered a hard obstacle, whereas a steep slope might be considered a soft one; where a soft obstacle is one that can be negotiated by adapting UVS behaviour (e.g. velocity or heading).

One solution to such challenges is to explicitly compute cost functions that are defined in behaviour space (e.g. mobility maps)⁴⁹ [144]. These plans may then be treated as input to reinforced learning techniques that then learn by physically interacting with the environment. At present, however, even though the behaviour bounds of the UVS are relatively straightforward and well-understood, complex UVS-environmental interactions still lead to unknown and un-modelled factors. This means that the evaluation function is not crisp and application of reinforced learning strategies can therefore be complex.

4.3 Verification and Validation

“We’re sitting on four million pounds of fuel, one nuclear weapon and a thing that has 270,000 moving parts built by the lowest bidder. Makes you feel good, doesn’t it?”

Alan Shepard (Astronaut)

A number of studies have indicated that military personnel believe that only humans are capable of operating in a “free flowing environment of an offensive combat mission” [24]. However, trust and reliability really only guides rather than determines the reliance that humans put in automation and recent research has produced several seemingly conflicting findings [163]. What is clear is that many military personnel do not want UVS operating in the same environment as manned platforms, particularly in hazardous environments. This is illustrated by the current need for a number of highly qualified humans to observe certain UVS and take control of them if they feel uncertain as to what they are doing. On the other hand, several studies [199] [26] [61] [86] [193] have demonstrated the human tendency to rely on computer-based recommendations, even though there may be contradictory (and correct) information readily available. This is usually referred to as *decision bias* [68] and typically results from the use of heuristics that people routinely use to reduce cognitive workload involved in problem-solving. It can result from errors of omission (where operators fail to notice a problem) to errors of commission (where people follow an automated directive that is wrong).⁵⁰

⁴⁹ Alternatively, the degree of terrain ruggedness might be monitored through feedback from onboard inertial sensors and the UGV behaviour then adaptively controlled, as appropriate.

⁵⁰ Paradoxically, for imperfect automation the greater its reliability the greater the chance of operator over-reliance; this is because of the rarity of incorrect automation advisories with the commensurate result that the operator uncritically follows unreliable advice [220].

4.3.1 *Trust and Reliability in UVS*

Ultimately, fully autonomous UVS will need to achieve higher levels of reliability due to the very nature of these systems (for example, there will be no-one to change the UGV's flat tyre). Furthermore, the more a system is capable of doing autonomously the less human intervention is required and the greater the endurance requirements become. As a result, while endurance is typically measured in hours today, in ten years this may become weeks or even months. As a minimum, therefore, systems reliability must keep pace with mission endurance.

Fundamentally, there is likely to be a minimum threshold for reliability and autonomous UVS will start to be adopted by defence forces when they are cost-effective and have a proven, reliable track record. That said few technologies gain instant acceptance when introduced onto the modern battlefield as warfighters often inherently dislike or even distrust a new system. As experience is gained, however, reliable technology tends to earn the trust of its user community and the value of the capability enhancement is appreciated. Trust and reliability are therefore key issues that drive the level of confidence – and hence degree of automation – that we place in UVS. Moreover, trust in automation, and technology more generally, is a multi-dimensional construct that changes with time. It is influenced by the types and format of information received by humans, their individual approaches to developing and determining trust, and influences such as system capability and reliability. Moreover, users of UVS frequently trust malfunctioning equipment and/or mistrust equipment that is operating correctly.

These imperfect relationships are described by [219] as the “disuse and misuse” of automation. Misuse refers to failures that occur as a result of humans inadvertently or inappropriately relying on automation, whereas disuse refers to failures that occur as a result of them rejecting the advice or capabilities of automation. The processes of disuse and misuse are often described as a binary process of engaging or disengaging in reliance, whereas the practice is often more gradual, complex, and the combination of multiple factors. Nevertheless, and even though many studies indicate that humans respond socially to technology (and computers in particular evoke similar reactions to human collaborators [201]), this simplification makes the topic easier to discuss and the modelling of key parameters more tractable.

It is widely acknowledged that while humans are very good at issuing high-level goals, managing uncertainty, and injecting a degree of creativity and flexibility into systems, they are also prone to disuse and misuse; where these biases are heavily influenced by experience, the framing of cues, and the presentation of information. To this end, UVS that provide inappropriately framed information may inadvertently reinforce the human tendency to use heuristics, and hence the potential for decision bias [201] [86]. Humans are also prone to physical and cognitive errors and it may be reasonably argued (see Legal Challenges) that any UVS sufficiently complex to take decisions on our behalf will likewise be prone to hardware, software and/or algorithmic errors, mistakes and failures. Humans often also become frustrated and confused when a machine does not do what they expect it to. Moreover, uncertainty in humans frequently manifests itself

as hesitation or failure to act. Nevertheless, during this process, humans usually continue to gather additional information to improve their awareness of their environment or their confidence in a particular line of action.

UVS, on the other hand, rely upon their sensors, actuators and IDT to reduce uncertainty or improve their confidence levels. Unfortunately, their sensors can introduce or increase uncertainty as they often have narrow spectral or physical fields of view. Alternatively, the algorithms used by the UVS may employ heuristics to abstract data or events, which can introduce noise or erroneous data, thereby reducing confidence levels or introducing greater uncertainty [24]. As a result, algorithms are often fragile and variation in a sensor's data stream can result in poor classification or processing results.

Trust is essentially based upon a perception that is linked to organisational, sociological, interpersonal, psychological and neurological processes which should (but usually do not) influence the design, evaluation, and training approaches to UVS. Not surprisingly, therefore, the topic is complex and draws on a diverse range of research from a number of fields. Moreover, it has generated several definitions [69]. However, trust between humans and technology is essentially driven by a combination of the probability that humans can successfully predict the anticipated action of the technology *before* they can monitor such action and the reliance they have upon the technology [26]. There is a considerable body of work that shows that not only is trust important to mediating how people rely on each other in relation to task completion, but that this is also extended to the relationship between humans and automation. Furthermore, these studies also indicate that it can be observed and measured consistently [163].

The relevance of this is that, if we can measure trust, we can use it as a framework by which we measure the level of 'trust' a UVS might have in a human. That is, the extent to which it might be able to reliably anticipate any likely human behaviour. One of the complexities in this regard relates to the unpredictability of autonomous UVS in unfamiliar environments. Often, when working with humans we can anticipate their actions by vicariously placing ourselves in their situation, or we have trained with them and have gained knowledge of their likely actions through experience, etc. Autonomous UVS have a tendency to surprise even their developers, although this is also one of their greatest assets as they can provide unexpected solutions to problems that could not be pre-programmed into them. In hazardous environments (i.e. ones in which UVS are likely to be used), however, these unexpected actions can be very disconcerting for humans. Nevertheless, just as we would attempt to develop an understanding of a human colleague in such circumstances, say, based on past performance in more familiar surroundings we should be able to develop an assessment of the perceived capabilities of UVS.

Establishing reliable automation in UVS also brings with it the challenge of identifying tasks, task components, or periods for which leadership can (or perhaps should) be assumed by the UVS rather than the human. At the very least, given the high workload environment of the modern battlefield and the cognitive and processing limitations of humans, we will need to consider whether human

supervision of all tasks and at all times is optimal. In other words, should the supervision of specific tasks be replaced by a more equal relationship that reflects true human-UVS teaming, or (in certain circumstances) is it more appropriate for the UVS to write the human out of the decision-making loop entirely?

To some the notion of shared leadership may seem a little far-fetched. However, let us consider a scenario in which a semi-autonomous UGV can vary the level of trust it places in its user based on the user's level of attention, health, workload, etc, all measured by psycho-physical sensors embedded in the HMI. If we were able to observe these human states reliably, this would provide us with a mechanism for adapting the level of autonomy assumed by the UGV and thereby provide a means for varying its behaviour. Appropriately implemented, this would build increased trust into the relationship. In this regard [59] indicates that it is unlikely to be sufficient for humans to simply understand the UVS decision-making; the UVS must also be given a means by which it can understand the (potentially dynamically changing) intentions of the human. Ideally, both elements of the team will then have the capacity to adapt to this mutual flow of information, thereby building greater levels of trust in each other.

However, it is imperative that appropriate levels of mutual trust be established as any distrust will result in a 'fight' for control. In this regard the concept of understanding the limitations of the members of a team has been shown repeatedly to be more important than the establishment of a trust relationship per se; studies indicate that teams will often preferentially develop techniques for achieving outcomes with 'faulty' members they do understand over 'high-performing' members they are not familiar with. Consequently, it is more important for humans and UVS to understand each other's goals and limitations than it is for each to know other's capabilities and enjoy mutual trust. We will return to the topic of measuring UVS performance in the chapter on Force Integration of UVS.

Measuring the level of trust in humans relative to their UVS is a major factor in training operators to develop advanced skills in collaborative activities as it allows initial biases to be reduced, provides knowledge about system capabilities, and applies a risk-assessment based on the behaviour of the automation. Moreover, human-UVS teams only become truly effective when humans know how to appropriately trust (and hence rely on) the automation as they can then use this trust to direct the UVS accordingly within the relevant context [111].

Unfortunately, different human roles (i.e. commander, user, team mate) require different types of interaction with the UVS and hence potentially different levels of trust (and hence ways of measuring it). Moreover, while many may wish to interact with UVS at a high level (e.g. "Are there any targets over there and, if so, engage them appropriately") there will be many occasions when interaction is required at a lower level (i.e. a user wishes to control a specific payload on one UAV within a heterogeneous team). When this occurs the outcome can be either synergistic or counter-productive, depending upon the team relationship, the familiarity of the human with the UVS and their mutual understanding of the context. To this end, successful outcomes frequently depend upon the UVS (or teams thereof) to act predictably and to support varying levels and/or frequencies of user-UVS interaction.

Trust is ultimately built on system reliability and predictability, and to a very large extent it is the system's architecture that combines and defines the interactions between the sub-systems – and in particular the system's tolerance to faults and 'erroneous' data arising from real-world interactions. To this end, it is the architecture that drives our ability to define and grow our trust in UVS. In the cases where we are considering UVS carrying and using weapons the system must not only be trusted and safe, but enact behaviours that are seen to be safe and trustworthy. That is, users and observers must feel confident that the UVS will only use weapons within the constraints of the Laws of Armed Conflict. To accomplish this, the UVS will require a level of status reporting, the capacity to explain any ongoing or planned behaviour, and the ability to 'ask for help' when necessary. We return to this issue in greater depth in the section on Legal Issues.

Technological reliability is also a key factor in the development of trust. For instance, if the systems reliability is relatively high users may come to rely on UVS so that the occasional failures do not substantially reduce the level of trust – unless the failures are sustained. Another factor might be the degree to which failures are detected or particular behaviours are more generally observed. Similarly, the ease with which manual over-ride can be enacted, the degree of user self-confidence, and the overall complexity of the task may also all prompt different task strategies from different users.

One final note; and as a number of 'controlled flight into terrain' UAV accidents have demonstrated, human error is not usually a function of just the human, but of the system inaccurately or ineffectively facilitating user understanding of how the system actually works [4]. In other words, one cannot remove human error by simply increasing the level of automation and removing the human operator as the extent to which the UVS is made less vulnerable to operator error through increased automation makes it more vulnerable to designer error during the design and manufacturing processes.

Clearly, training has the potential to minimise or mitigate some of this, although it has been shown that training alone cannot overcome issues of trust arising from many aspects of poor design. Providing the users with interaction paradigms that they are familiar with, for instance 'natural' (i.e. human-like) interaction through gestures and speech (or even UVS that have identifiable 'personality traits') has been shown to improve trust between users and their technology [53]. Over-reliance on such technology, on the other hand, can also result in poor systems monitoring and a reduction in overall performance, just as too little trust can also lead to over-monitoring, which detracts from a user's capacity to carry out other tasks.

4.3.2 Systems V&V for Autonomous UVS

Developers naturally strive to achieve 'best practice' by implementing basic rules of thumb, keeping their designs simple, providing suitable documentation and creating initially stable designs. Empirical evidence suggests, however, even using a combination of peer review, static code analysis, subroutine and algorithmic testing, unit testing, component testing, functional testing (including human,

hardware, and software-in-the-loop), integration testing, system testing and qualification and acceptance testing that it is practically impossible to provide bug-free software on an unblemished processor past a certain level of sophistication [145]. This is due to a number of reasons:

- While the majority of production systems are built to a specified level of quality, they are also built to a budget and schedule;
- The complex, cluttered, dynamic and unstructured operational domain into which they are inserted differs substantially from their military test environments;
- The engineers and programmers cannot a priori anticipate all possible contingencies; and,
- The software involved usually contains many lines of code.⁵¹

Moreover, in a network of autonomous UVS we are considering the interaction between multiple software modules running on different processors, operating systems, and with architectures possibly unknown to each other in advance and in all probability across a range of UVS platforms and environments. In the final analysis, therefore, we can be reasonably sure that in addition to autonomous UVS operations entailing considerable fundamental uncertainty, at some level the system will probably malfunction and that we are unlikely to be able to predict the specific nature or timing of these failures.

Nevertheless, UVS are non-unique in that there are several examples of intelligent systems whose malfunction may have severe consequences. Such systems require a great deal of care in regard to their design, operation, and maintenance. Moreover, increased safety in these safety-critical systems must typically be traded against criteria such as usability, cost and performance. In order that safety is not inappropriately compromised at the expense of one of these other criteria, sound ethical judgements must be made. Typically, these cost-capability decisions are made on the basis of some statistically significant criteria such as “The life expectancy of a human shall not be altered by using such a system.” However, such criteria do not provide us with any absolute measure of what constitutes *safe* or *known malfunction*. Clearly the system will need to be designed to ‘world’s best practice’, but this on its own is probably insufficient; not least because such standards and practices are currently informal and therefore not legally binding (see Legal Issues).

As previously indicated, the sum behaviour of an autonomous UVS is a function of the interactions between its multiple interacting and independent elements and its human supervisors. Furthermore, as the degree of autonomy increases it becomes increasingly difficult to predict the sum state of the system. Moreover, the system is actually a function of a number of linked processing elements (hardware and software) and humans (programmers, engineers and users) and becomes significantly more complex as these systems are networked

⁵¹ Standard V&V techniques usually deliver between 97%-99% (or higher) overall code-defect removal for embedded systems if all steps are carried out correctly or perhaps 85% for application and information systems software [93].

either to each other or other technologies. As there are multiplicities of architectures, data formats, operating systems, programming languages, compilers and communications protocols, not to mention an almost infinite variety of hardware combinations.

From a technical perspective, however, we may consider two systems elements: hardware and software. Traditional hardware systems embody much of their functionality in the components that comprise them so they are relatively easy to model, they fail statistically through use or external damage, and their reliability is fairly predictable. Consequently, hardware systems can be analysed relatively simply and straight-forward tests can be formulated to prove their integrity before permitting their operational use. Furthermore, engineers can usually solve the problems of poor reliability with hardware redundancy.

When functionality is instantiated in software, however, the sheer number of states and a lack of regularity usually makes it much harder to bound the possible failure modes, and hence to devise tests against them.⁵² Furthermore, there are frequently many subtle and often unexpected interactions between modules. As a result, a complete analysis of all possible failure modes and their potential impacts may not be practical. Furthermore, redundancy does not usually solve software reliability problems as software fails almost always as a result of some latent design error. Hence, failure of a critical sub-component is often highly correlated with the failure of a duplicate backup system, unless different software designs are used. As a result, while it will increase cost, building-in redundancy may not improve the reliability of UVS.

An increase in the reliability of autonomous UVS will come from the development of affordable software Verification and Validation (V&V) strategies that reduce costs and compress production schedules. However, although there are a range of systems engineering and other analytical techniques available for evaluating the likely performance of software (e.g. [115] [186] [189] [167] [222] [223] [247]) and current certification practices have historically produced safe and reliable control software for many complex systems, verifying and validating software that controls the key functions of next-generation UVS poses significant challenges in terms of providing the requisite levels of confidence. As a result, current techniques are unlikely to be cost-effective for a number of reasons.

- Application of existing V&V strategies is a non-trivial undertaking [222] [223], and it is highly likely that modelling and measuring the reliability, usability, testability, portability, and understand-ability of the critical elements of the UVS software will be a major undertaking in itself. This is because almost all of the ‘intelligent’ functions in a next generation, autonomous UVS will be software modules that are likely to be distributed across a number of programs and processors with no one processor, program or programmer knowing the full extent of individual

⁵² Using the ‘rule of thumb’ (see Footnote 27 in Looking Forward) this means that there will be around 91,000 test cases required for an embedded software system that has 10,000 function points, although this number could be much higher (perhaps by a factor of ten) due to specific test-driven development [93].

outcomes or determinations. Moreover, if the UVS forms part of an NCW environment, software modules originating at another node in the network could be executed within the UVS that has potentially formed on some ad hoc or other ‘unpredictable’ basis.

- Although many of the functional components of an autonomous UVS are likely to be based on independently mature technologies, the sum behaviour of the UVS will be a function of the interactions between these components, the human supervisors (via an HMI), and a range of other external technologies that may include other UVS.
- UVS are polymorphic because they interpret data from a number of different perspectives and manipulate information in accordance with environmental conditions, the nature of the mission, and the problem at hand.
- Some sophisticated UVS may be designed to be intentionally unpredictable so as to inject a degree of creativity into the UVS mission, as predictable systems are not necessarily optimal for military operations. From a V&V perspective, consistency is obviously desirable, but in an adversarial context the capacity to predict exactly what the UVS will do may well be disadvantageous. Consequently, a balance will need to be struck between consistency and unpredictability that allows the programmers to understand, trust, and hence verify the software.

Testing software will need to be geared toward the verification of four key, high-level requirements: the loss of control, survivability, UVS performance, and safety (which includes compliance with the Laws of Armed Conflict). There are a number of issues:

- The software designer is not usually an expert on sub-component design;⁵³
- Next-generation UVS may replace human ability and judgement and our comprehension of higher order cognitive functions is not yet well-framed;
- Improvements in automated testing regimes may reduce labour costs and testing hours, but may not reduce them sufficiently relative to the emerging requirements.
- Most software requirements are incomplete (i.e. we will probably need to specify the unwanted as well as the desirable behaviour of the UVS and its intelligence);
- UVS will often be used in contexts for which they were not designed (i.e. we need to understand how the software operates across a broad environmental spectrum);
- Software is often changed (i.e. hardware fixes usually result in the recovery of a system’s functionality, but minor software ‘fixes’ may introduce new faults);

⁵³ There is evidence to suggest that embedded software engineers attend domain-specific events rather than mainstream computer shows or software engineering conferences [93].

- The impact of software changes are non-linear (i.e. a small modification may have significant – even catastrophic – results on system performance);
- The state of the UVS often depends on its past history in intricate ways that may involve several components or other sources of a non-deterministic nature;
- Adaptive learning techniques can adjust their own logic during execution and some software techniques have the potential for self-healing. Both would render obsolete any certification process;
- It may not be possible to exhaustively test the IDT software because the number of states is so large;
- Reliability criteria may be driven by payload or may take on other forms to accommodate the functionality of the remote user;

As pointed out in UVS Components, many of the functional elements of a UVS will be embedded software units with the potential for significant human impact. Consequently, the defect potential and removal needs to be monitored closely or there will be serious issues of liability (see Legal Issues). Furthermore, such quality control issues will clearly impact schedule, cost, and systems reliability. Finding and fixing bugs will therefore be the most expensive activity in UVS software development [93].

Having said all this, let us now return to the notion that a UVS control system needs to perform a number of tasks: avoid obstacles, avoid other UVS, operate within its performance envelope and, once these priority tasks have been accommodated, undertake higher order tasks such as mission planning, navigation, surveillance, reconnaissance, target location, sensor scheduling, coordination, communication, executive decision-making, etc. Now, rather than attempting to consider the autonomous UVS at the system level, we can apply techniques employed by [115] [167] and [186] and divide it into its constituent autonomous functions or categories of autonomous software. Analysing the individual requirements of these constituent functions then reduces the complexity of the V&V task somewhat.

Mission and Trajectory Planners

Planners typically make decisions by projecting action into the future on the basis of a model of the UVS, its current and potential behaviour and the environment, and then evaluating the outcomes according to a cost function or some other selected criteria. The evaluation function then represents the UVS objectives and constraints through the return of high values for plans that meet mission goals without violating the performance envelope of the UVS. Typically, this involves some form of search through a set of potential plans until an acceptable or feasible plan is found. Consequently, the key is to apply pruning techniques so that only successful plans are likely to be generated. There are four major risks (listed in order of increasing severity) [167]:

- The plan makes inefficient use of resources;
- The plan could not be generated;

- The plan generated is not feasible; or
- The plan places users of the environment, supervisors or the UVS at risk.

Clearly, we need to worry most about the last one. However, to assure ourselves that safety is not an issue we do not need to verify the entire planner, only its evaluation function. If the evaluation function is correct then the UVS and/or users cannot be placed at risk. Fortunately, the evaluation function is likely to be based on algebraic expressions or software techniques that rely on physical laws and/or techniques that have been used for decades. The main difference will be the absence of human oversight and ‘double-checking’ of results. As a result, we must hold the evaluation functions to higher standards of verification, but this is more a matter of degree, than novel concept [115].

The generation of infeasible plans is more serious than the absence of a plan as the latter is simply a fault. That is, the system should be able to detect that it has not generated a plan and can automatically invoke some sort of recovery procedure. The generation of infeasible plans requires the UVS to have an understanding of its state and may therefore be verified by evaluating some of self awareness criteria (see Dynamic Planning).

The use of resources is context-dependent and should be evaluated thus. For example, an autonomous planner might propose to do something a human would not. However, in terms of verification, this situation should be considered against the case where the mission may not have taken place rather than against more absolute conditions.

Navigation

Verification of autonomous navigation software is very challenging as it is mission-critical and relies upon the complex integration of algorithms in the context of a system embedded in a complex environment. Furthermore, some navigation capabilities depend upon the self-awareness of the UVS and/or its capacity to cue sensors to maximise its potential for observing certain types of data and hence its ability to perceive and predict its environment in the presence of uncertainty.

The inability to fully or exhaustively test software is not a concern in and of itself as many non-trivial systems cannot be exhaustively tested. Furthermore, exhaustive test is not required to produce reliable software. For example, UVS systems – like software programs – have structure and what often passes for exhaustive testing is in fact only sparse testing from the range of all possible states. In other words, the behaviour of the UVS in one state is not always independent of its behaviour in other states. As a result, testing the IDT in one state may provide information about other states, which can be grouped with respect to particular properties of concern. The key to making the IDT reliable is then to design it in such a way as to make its structure testable; or at least to allow its states to be decomposed into a tractable number of groups with respect to particular properties of interest.

Executive Decision Makers

A key challenge for an autonomous UVS is the application of work organisation. In other words, only after the basic patterns of stability and predictability have been determined can the UVS be productively applied. Good decision-making agents are therefore analogous to sequencing engines as they reduce the complexity of the executive decision-making by breaking it down into component-decisions that are simpler to make. This enables us to put greater structure into the verification problem and also to act on the sub-decisions and the information pertinent to each of them. However, such agents typically have complex semantics and comprise multi-threaded algorithms, which are prone to race conditions, deadlocks, and non-deterministic behaviour [115].

Another issue for V&V is that irrespective of the technologies used, the mere act of removing the human introduces risk because IDT, which have control over the UVS, can make errors that lead to mission termination or system failure. However, if we consider the problem in context, the autonomous function is likely to have been introduced because of some human failing: an inability to react quickly enough, the monotony of the observation task, etc. That is, the technology is embedded into the UVS because the processes involved are currently unreliable as a result of their involvement with humans. To design a robust IDT, therefore, we simply need to repeat the process and aim to design it with high levels of self-awareness. In other words, when the software fails the IDT must never fail to recognise that one of its software components have failed. In this way, unreliable UVS components can be combined into an overall architecture that has a fall-back recovery procedure. Acting on the component decisions in this way also aids in reducing the time it takes to verify decision-making software by allowing the information to be presented in a more organised fashion, thus saving valuable time interpreting what might otherwise be a complex situation. For instance, a verification strategy might order sub-decisions according to a hierarchy in which they have been determined to achieve traceability. This allows us to ‘step’ through any decisions that might need to be made [222].

4.3.3 *Simulation-Based V&V*

Simulation-based V&V is a flexible framework for simulating, analysing and verifying autonomous UVS. Essentially, an instrumented test-bed consisting of the actual control software and processors is embedded in a simulated operating environment. Conventional and model-based testing is then combined: the real software and hardware is executed and verified rather than an abstract model derived from the system; yet the simulated environment allows execution ranges over an entire graph of possible behaviours rather than a suite of linear test cases [115]. Ideally, each internal software state is marked to identify that it has been tested to avoid redundant testing or note any variation.

Simulation-based V&V avoids the need for developing separate models for verification purposes and, more importantly, the need to scrutinise each violation against the real system to see whether it corresponds to a real or a modelling inaccuracy. On the other hand, while simulation-based verification provides

important potential gains in scalability, automation and flexibility, it is generally less efficient than model checking verification techniques [247].

To enable controlled execution, instrumentation must be introduced into both the software under scrutiny and the test environment. Furthermore, if the test-bed is capable of iterating over all alternate events at each state, back-tracking to previously visited states and detecting states that produce similar behaviour it will constitute a virtual machine with a fully controllable state space. To constrain the state space, however, the environmental component of the test-bed must usually be restricted to a well-defined set of vignettes or scenarios. We may then use such a tool in three ways:

- By applying the simulation-based verification approaches described above;
- As infrastructure for developing a program framework for autonomous UVS; and
- As a framework for evaluating and diagnosing concepts of use.

Aside from the system under test, the tool will require three components [223]:

- A diagnostic component capable of interpreting the physical system;
- A simulator for the physical system on which the diagnosis is performed; and
- A driver for generating commands and faults according to user-provided scenarios

To verify the system, the tool should then run through all conditions specified in the scenario, back-tracking as appropriate to explore alternate steps and executions. At each step, the tool should also check for error conditions and, if an error is reported, record and report the sequence of events that led to the current state. Verification of the diagnostic software is also required as the key to the voracity of such a test-bed is its ability to accurately observe or infer information on the behaviour of the system under test. This is contextually dependent and must take into account the run-time conditions under which it should be possible to acquire certain information.

Finally, although not strictly the same as V&V, Accreditation must also be considered for autonomous systems. For example, how many hours and under what conditions should we test a UGV to ensure it does not lose control?⁵⁴ Furthermore, what protocols and safeguards must we instantiate and test to ensure that such systems cannot be intentionally or inadvertently subverted and do we even know whether this is a real issue, and if so, how to characterise this task? There are also issues of test infrastructure, such as whether or not the existing test facilities, designed mainly for manned systems, are adequate and (say) how test data will be collected when the instrumentation normally mounted on a vehicle is larger than the vehicle itself (e.g. a MAV or small UGV in a sewer).

⁵⁴ During the 2005 DARPA Grand Challenge, and without warning, one UGV that was performing perfectly well suddenly left the course and almost hit a building, only missing it because the chase vehicle activated the UGV's e-stop; not something that may be an option for vehicles engaged in combat.

4.3.4 Health and Usage Monitoring

Persistent autonomous military UVS operations will place great emphasis on health, usage monitoring, and fault detection, isolation and recovery systems as such systems must not only recognise that something has gone wrong, but also determine what has gone wrong; and leave the UVS in a safe state by restoring its functionality in the face of failure. In any system that interfaces with humans, however, the overall output will be affected by the physical or cognitive workload of the human and the limited physical and processing abilities of the UVS. Moreover, a UVS that is able to perceive its environment through limited sensor modality may induce or suffer from ‘cognitive blindness’ [118] when the UVS or its supervisor focuses (or fails to focus) on a particular environmental event; or attends to the demands of a particularly onerous task triggered by such an episode.

In other words, the separation between the operator and the UVS deprives the human of a range of sensory cues that are available to the pilot or driver of a similar manned vehicle. Furthermore, rather than receiving the sensory input directly from either the vehicle or the environment in which the vehicle is operating the UVS operator receives only that sensory information provided by onboard sensors via a data link. The sensory cues that are typically lost include visual, olfactory, auditory, kinaesthetic and vestibular input. For example, an actuator malfunction may be signalled to the pilot of an aircraft via visual, auditory, and haptic feedback. In contrast, for a UAV this failure may be indicated solely by perturbations of the camera image. This manifests itself in two ways. For tele-operated UVS this is felt in terms of the operator’s moment-to-moment control of the UVS; for more autonomous UVS, the vehicle’s health and status at any instant are unknown.

The end result is that a considerable amount of data must be relayed from sensors and systems onboard the UVS to the operators at the GCS. This data must also be processed and presented to the users in such a way as to simultaneously minimise their workload in regard to monitoring it and maximising their capacity to interpret and understand it, which is in addition to any information needed to maintain task situational awareness, control the vehicle or progress towards mission objectives. Furthermore, the potential for controlling, coordinating and monitoring the states of multiple vehicles using a single operator diminishes exponentially with the increase in the number of vehicles, unless the vehicle’s situational awareness is determined autonomously. To avoid network latencies and communications scheduling problems (that are additional to any required for mission completion), this processing must take place onboard the UVS.

As a result, the absence of an embedded pilot or driver promotes the need for a Health and Usage Monitoring System (HUMS) located onboard the UVS. Such systems must autonomously process, interpret, and deliver meaningful information about the status of the UVS platform, its sensors, and sub-systems. The key requirements are that it monitors the performance of UVS at both the holistic and functional component level in order to detect anomalous behaviour, characterise its nature, extent and seriousness, and report it to operators within useful timescales. Ideally a HUMS will also attempt to mitigate any potential damage, perhaps by affecting a repair.

Typically, HUMS will make use of analytical models of hardware sub-systems to provide estimates of the anticipated sensor observations and/or vehicle responses to actuator commands. They avoid the additional cost (and weight) of redundant hardware and can determine lost functionality at a sub-system level. They employ hypothesis-testing and robust estimation techniques to detect and isolate these failures, which can correspond to failed actuators, sensors or other systems failures that cannot be adequately assigned (e.g. a UGV has become bogged in wet mud). Typically, the statistical tests also look for changes in the statistical properties of any variables so that the HUMS can perform prognostic analysis on the likely failure trajectories or adapt maintenance regimes.

The detection of anomalous events requires an array of suitably placed and networked sensors, a strategy for acquiring and then processing the data, knowledge of the operating environment of the UVS, and the potential impact of likely threats and stressors. Based on this schema any damage must then be characterised and prioritised in terms of the vehicle and/or its mission in order for the HUMS to autonomously determine the urgency with which a response needs to be mustered. In an ideal system, the HUMS will also use its array of sensors to deduce information relevant to events leading up to the anomaly to identify and possibly isolate its cause. Finally, the HUMS should formulate a response option in the form of a sequence of actions or recommendations to operators that are achievable within the window of opportunity pertinent to the seriousness of the anomaly.

Clearly it helps to anticipate the type of events or anomalies that a UVS might experience, and these may be broken into two broad categories: external (environmental) anomalies and internal (vehicle-based) anomalies. External anomalies are likely to be dependent on the environment and therefore the type of platform or mission. For example, mud and water may enter the mechanical systems of UGVs and UAVs may suffer from icing on their wings. On the other hand, internal anomalies are likely to be broadly similar across UVS from each of the environmental domains even though their nature, frequency and severity are likely to be vehicle-specific and/or dependent upon operating conditions (and hence indirectly lined to their environments). Examples of internal anomalies include the failure of functional components (sensors, navigation/control systems, communications, propulsion, energy storage, etc) and the mechanical failure or degradation of materials, structures or interfaces. Clearly, in order to be of use a HUMS must measure a spectrum of mechanical, electrical, chemical and software-execution properties over a wide range of temporal and spatial scales and adaptive and reinforced learning techniques are particularly useful in determining the frequency and location of any sampling regimes.

Adaptive learning techniques are particularly useful for fault detection and diagnosis relative to unanticipated events. There are three primary categories of technique: model approximation, supervised learning and adaptation, and reinforced learning.⁵⁵ The regression techniques typically employ the use of

⁵⁵ These techniques have been used to model complex and non-linear systems such as aircraft flight dynamics, space vehicle control systems, jet-engine combustion, and helicopter gearboxes [202].

networks of radial basis or other functions to represent complex physical processes that are otherwise hard to model. These are then used to generate models for hypothesis testing within state estimators. Frequently, these techniques are supported by simulation data to provide an initial training set, whereupon they are then supported by data collected during field trials and operations.

The supervised learning techniques use a learning paradigm to select an optimal or good action to be implemented given the current state of the system. The learning is said to be supervised as the selection of a good action is based on a network of coefficients trained through human supervision or simulation. Once the system has been trained via this supervision, the system has the ability to generate 'good' actions given an arbitrary system state. Reinforced learning techniques are currently immature, but are capable of learning without a priori knowledge of a value function; that is, the technique learns the value function and evaluates goodness 'on the fly.' Reinforcement learning techniques are typically computationally intensive and are not usually able to run in real time on PC-based architectures. One such technique might learn a model of the vehicle and run it 'backwards' – i.e. take raw sensor data and commands sent to the hardware and find the most likely state of the given model that explains the observed measurements. However, in practice, the quality and robustness of this technique is likely to depend entirely on the accuracy of the model [291].

4.4 Multi-vehicle Systems

While it is relatively easy to build larger UVS that operate long enough and can travel far enough to perform useful military functions, these UVS are usually very costly to acquire, run and operate. The development process for many of these larger military vehicles also parallels that of their manned counterparts, which stresses longer life, higher levels of maintainability, multi-role capability and high reliability. The resulting systems are therefore more expensive with life-cycle costs and logistic complexities approaching those of manned platforms. Moreover, the continued drive for cost effectiveness, stand-off weapons delivery, precision engagement, the pressure for smaller operator footprints and higher workload environments, and the capacity for cooperatives of multiple UVS to accomplish tasks that are difficult or impossible for single UVS have all combined to increase interest in networks of smaller unmanned vehicles with increased automation.

As a result, Affordably Expendable⁵⁶ multi-UVS cooperatives are gaining prominence as they can be developed to carry out high value, high risk missions that are beyond the capability or justifiability of larger, single-vehicle systems. There is, of course, no free lunch. Even though smaller, less expensive, lighter systems lend themselves to being placed in harm's way, and their spatial benefits present opportunities not afforded single UVS, they are generally less capable than

⁵⁶ The concept of affordable expendability relies upon the notion that the useful life of the capability is a function of its constituent payloads and technologies rather than the physical life of the airframe [267].

their larger, more strategic counterparts, which tend to have longer ranges and carry more capable payloads.

For example, to have a 90% confidence in the classification of a target it is generally accepted rule-of-thumb that an image must have 16 pixels across the narrowest relevant dimension of the target. Consequently, to (say) recognise facial features (1 cm resolution) from a range of 1km requires a camera aperture of about 10cm. In other words, very small UAVs, which carry small sensors, need to approach their targets relatively closely, while larger UAVs are able to stand off considerably further and achieve the same end. Since the size of the camera aperture is proportional to the range to the target (for the same image resolution), to recognise faces from 10km requires a 1m aperture. As a result, if high-resolution images of significant swathes of the earth's surface are required, a high-altitude reconnaissance UAV needs to be relatively large to accommodate the necessary camera. Alternatively, a number of much smaller and lower flying UAVs must cooperate to achieve the same end – and must fly much lower. Then again, another effective operational combination is to have larger, high altitude detector/classifier UAVs cross-cue smaller “examiner” UAVs. This combination also lends itself to lower resolution imaging radars that can probe clouds, working with higher resolution optical imagers that do better at lower altitudes in clearer, cloudless atmospheres.

Furthermore, an equivalent problem to the above EO example exists for the acoustic sensors used on UUVs as the smaller UUVs cannot carry the larger, longer-range sensors. As with imaging radars, to some extent the physical laws limiting acoustic sensor resolution can be overcome by single or multi-vehicle Synthetic Aperture Sonar (SAS). Moreover, larger UUVs that cannot approach their targets closely enough to overcome the limited transparency of water can deploy smaller UUVs that carry optical sensors and can approach their targets more closely than their larger counterparts.

This lack of individual capability may be offset by the increased affordability of the multi-vehicle systems, our ability to derive process gain by networking the UVS and sensors (potentially achieving multi-aspect SA across the environments) and our capacity to withstand losses due to conflict or malfunction. Furthermore, a distribution of autonomy throughout multi-UVS cooperatives provides redundancy through the system's ability to re-allocate tasks and objectives, thereby increasing the number of objectives that can be met and the overall probability of mission success.

The endurance of a UVS depends upon its stored energy divided by its minimum power requirements and energy storage density for any given material is fixed. As a result energy storage scales (approximately) according to volume. Consequently, the range of a UVS is roughly proportional to the cube of its characteristic dimension, limiting our capacity to build arbitrarily small UVS. This presents practical problems of getting the (usually) slower and lower altitude UAVs to their required locations if they are not launched locally. One attractive option in this regard is – when they work with larger UVS – to have the larger

ones deploy the smaller (usually expendable) ones, that can then be used for final target confirmation.⁵⁷

The development of arbitrarily large networks of small UVS, however, is constrained by the requirements of internal communications essential for UVS coordination and network functionality as these smaller UVS must also contend with the inverse square law for omni-directional communications range requirements, at least for signal acquisition. A range of other key considerations for multi-UVS cooperatives include [108] [151] [174]:

- The number of assets in the cooperative could potentially be large
- Scalability is desirable as UVS may leave or join the cooperative
- Humans must be able to set goals for and interact with the UVS
- The health of the UVS and their sub-systems need to be monitored
- Each individual UV in a team needs to possess its own complex behaviour
- Each team within the cooperative should possess its own complex behaviour
- Humans can be supervisors as well as controllers of individual UVS or their payloads
- Supervisors and UVS are potentially distributed over a wide geographic area
- The integration may take place within a single environment or across them
- The cooperative should exhibit a highly fluid team-tasking and structure
- Operations occur in an environment that displays adversarial behaviour
- Situational awareness events can require a high speed response
- There may be various supervisors (of varying authority)
- There is a high probability of losing resources

Unfortunately, most UVS require the full attention of at least one and usually two or more skilled operators, and the ratio of personnel-to-vehicle rises to around 4:1 for even the small tele-operated UGVs when maintenance is taken into account [234]. This ratio is significantly higher for larger UVS such as Global Hawk, where the ratio is closer to 20:1. Clearly, given that most humans cannot manage multiple high-speed cognitive tasks in parallel significant advances in automation are needed if multi-UVS cooperatives managed by a small number of humans are to become militarily and economically viable.

To this end, there are a number of variables that must be considered when determining effective operator-to-vehicle ratios [108]:

- The spatial and temporal complexity of the environment
- The cognitive workload, training, experience, etc of the users
- The level of trust exhibited by the users and the reliability of the UVS
- The adversarial nature and/or temporal dynamics of any human tasking

⁵⁷ An example of this concept is the Finder UAV, developed by the Naval Research Laboratory which can be deployed from a long-endurance Predator UAV.

- The low-level (“navigation and mapping”) capabilities exhibited by the UVS
- The degree of high-level (“task-organisation”) automation exhibited by the UVS
- The capacity of the UVS to dynamically adapt these levels of automation
- The amount and nature of information passed between the user and the UVS
- The extent to which any decision-making may be distributed and/or centralised
- The capacity of the UVS to autonomously form into or dissolve from teams
- The degree to which the UVS is able to monitor/adapt to its own state and health
- The degree to which the UVS/humans are able to monitor systems performance
- The degree of network and/or processing latency inherent in the system

Multi-UVS research has its origins in the 1980’s and the field is still new enough for none of the topic areas to be considered mature, although some areas have been explored more extensively than others. Initially, a great deal of the research was based on the social characteristics and behaviour-based paradigms of biological systems such as ants, bees and birds. This early work demonstrated that the use of simple, local control rules allowed robots to mimic the foraging, flocking, aggregation and trail-following characteristics of these biological systems. Furthermore, the introduction of dynamics into the simulated ecosystems allowed the multi-UVS teams to demonstrate emergent cooperation resulting from selfish interests.

This work was then extended to incorporate studies in predator-prey systems, although much of this work was carried out in simulation and much of it focused on the development and evaluation of various pursuit policies. As a consequence, adversarial engagement between multi-UVS, such as that found in higher order biological systems, tends to have been studied in domains such as robot soccer (e.g. [153] [158]) or from the perspective of expected capture times and the sensing capabilities of the pursuers [152].

Much of the early work also tended to focus on using reactive or deliberative techniques (see Planning in Dynamic Environments). More recent work has used the benefits of each, creating hierarchal systems that engage low-level reactive planners under higher-level deliberative ones. Using modern, powerful processors these hybrid techniques are now sufficient to provide dynamic planning solutions for single UVS, but not for multi-UVS cooperatives. In part this is because many techniques “repair” their previous plan by optimising against information observed in the vicinity of the UVS location; a condition violated when multiple UVS operate in a geographically dispersed formation.

The challenges for multi-UVS arise predominantly out of determining the strategy that maximises overall systems performance, where such strategy decisions include whether the control should be explicit or implicit, whether the origin of the tasking should be distributed or centralised, the extent of the

communication, the complexity and power of heterogeneity versus the relative ease of homogeneity, and the nature of the individual motivation (i.e. selfish or socialised) [76].

To achieve task and resource allocation in dynamic, adversarial environments a number of researchers have used free market economic theory, auction strategies and biological inspiration [116] [45] [84]. Another classical approach is to start by building terrain or world maps and then develop and execute the relevant strategies in known environments. There are several techniques available for building maps, but most of the common ones are based on Bayesian estimation and Extended Kalman Filters (e.g. [48]). Unfortunately, even two-dimensional map-building processes are time consuming and computationally intensive. Furthermore, many techniques assume accurate maps and worst-case motion for the adversary, which with noisy observations and inaccurate maps usually leads to overly conservative policies for pursuing the adversary.

As a result, a number of researchers have now applied game theory to the problem and combined the map-building and pursuit-evasion policies into a single probabilistic framework, some with autonomous (UAV-based) supervisory UVS [283]. A number of researchers have also considered active evasion strategies based on partially observed Markov decision processes (POMDP's), usually based on vision-based sensors and executed in simulated environments [129]. Others have used optic flow to determine the number of moving evaders as well as their position and orientation [284]. All of these approaches, however, designate the roles of the UVS prior to the commencement of the games as either pursuer or evader; they do not provide, for instance, the evader with the policy option of countering their pursuers by becoming the hunter.

Game theory appears to provide this option, with another attraction being its capacity to model a multi-UVS task (such as search, surveillance and target tracking in an adversarial environment) within a framework that provides the flexibility to use different solutions or role-playing concepts: one based on the cooperative behaviour of the participants and another based on non-cooperation. Application of these concepts in the field of economics has accounted for the lack of altruism shown by participants, which has resulted in untenable cooperative frameworks – unless cooperation is enforced by a third entity. Additionally, as [263] has shown the non-cooperative Nash strategies perform better than the cooperative ones in the presence of noisy sensors, unreliable UVS or faulty communications. This is because the uncertainty maps derived from the contributions of each cooperating UVS changes with time in a manner unknown to the other agents. In such situations the cooperative decision-making breaks down.

Many of the multi-UVS coordination issues such as task allocation, path and trajectory planning, formation optimisation and pursuit-evasion strategies are now becoming well understood, although demonstration of them using real UVS in outdoor and unstructured environments (i.e. as opposed to simulation) has been rather rare. More recent research has focused on motion coordination within the context of behaviour coordination such as target search and feature-tracking behaviours. As a result, research into path planning and control, multi-UVS task/resource allocation, behaviour coordination and communications has become

coupled. This is largely because the structure of the multi-UVS cooperative changes with time and the properties of the cooperative change with structure. That is, the position of a UVS within its cooperative's structure and relative to its goals determines its projected sensing options, prospects for information gain, and capacity to (say) accurately explore, map and locate key features in its environment; just as its inherent capabilities, sensing options, scheduling of payloads, and so on impact the potential UVS trajectories, behaviours, feature tracking accuracy, communications strategies, etc.

Recently several have researchers attempted to address these coupled tasks as a single technique (e.g. [197]), whereas previously the problems of and approaches to communications and sensor scheduling, feature tracking and trajectory control were largely de-coupled and addressed using independent algorithms and strategies and then combined using some form of executive controller or architecture. More recent work attempts to manipulate the sensing process in order to maximise the information gain and feature location estimation, without using any a priori information. When the sensors are passive, this introduces a number of aspects that are not under sensor control (i.e. when precisely observations are made and what the observations are of), both of which have an impact on the development of longer-term scheduling strategies for the sensors and UVS.

Multi-UVS behaviour is often instantiated through the coordinated grouping of individual UVS into teams, the members of which take (or are instructed to take) a decision to commit to a particular task but who receive common reward for task achievement as a result of team decisions. The team members receive information about their environment and progress towards their task through observations and communications with each other, whereupon they take decisions based on their respective information. Teams can be self-organising or commanded through a centralised authority (and hybrid schemes exist also). In the case where the teams are self-organising, information may be explicitly or implicitly shared, where explicit communications is the specific act of conveying information from one UVS to another and implicit communications is the synchronism of UVS action through shared understanding.⁵⁸

4.4.1 Multi-UAV ISTAR Example

For context, let us consider the case of a multi-UAV cooperative tasked with surveying a potentially hostile region of interest⁵⁹. There are clearly a range of platform, mobility, propulsion, and energy issues that need to be addressed for such a system. As with the rest of this text, these are not dealt with here, except to note that the shortcomings and vulnerabilities of larger, slow-moving UAVs in this context are well known and have been described elsewhere (e.g. [102]). The cooperative must undertake a number of tasks:

⁵⁸ A classic example of implicit communications is lions stalking their prey. They do not communicate yet still synchronise their actions on the basis of their perception of the environment and a knowledge of the other lion's location, actions, etc.

⁵⁹ The example may be easily translated into a UUV, USV or UGV context.

- Based on a priori information about (say) target distribution and mission priorities allocated by a commander, the mission planning software must generate a series of near-optimal trajectories for each of the UAV to follow such that they visit as many regions of opportunity and interest as possible, while simultaneously avoiding as many hazards as possible.
- The optimisation of these trajectories must be based on (potentially time-varying) cost functions that allow for such things as: the distribution of payloads within the cooperative; the prioritisation of targets; the robustness of the proposed solution to operational and environmental uncertainties; the individual capabilities of the participating platforms; the benefits that derive from the association of the UAVs into teams; the communications and sensor scheduling requirements between the platforms to enable this cooperation, 'no-go' and 'difficult-to-go' zones and any UAV deconfliction requirements.
- Once underway, based on a change in the environment observed by one or more sensors onboard each UAV, the system must respond by dynamically re-calculating trajectories, re-allocating task/team associations and enabling payload and/or platform actions (within the constraints outlined above) based on the manipulation and fusion of the new data.
- Similarly, based on a change in the environment provoked by one or more of the UAV payloads or actions (e.g. jamming, UAVs joining the group), the system must dynamically re-calculate their trajectories, associations, etc.
- Based on a change in an operator's priorities or task objectives the system must respond by dynamically re-calculating their trajectories, associations, etc.
- Finally, all of the computational processing and communication must be achieved within the physical and electrical resources of the UAV and in real time.

In the mission-planning phases multi-UAV operations require multiple aircraft to be designated pre-defined flight paths, regardless of whether or not the UAVs have the ability to cooperate with one another. Irrespective of whether these pre-defined flight paths are generated using a route-planning algorithm or manually by an operator, the resultant trajectories must conform to acceptable levels of airspace deconfliction in terms of the temporal and spatial separation between the aircraft.⁶⁰

If the UAVs are networked and can coordinate their efforts then after the mission plan is uploaded one or more of them may dynamically and continuously

⁶⁰ Trajectory deconfliction and collision avoidance for multiple UAVs within a single environment implies similar route re-planning requirements separated mainly by their time scales. Deconfliction is a medium-long range task that attempts to avoid a collision while still allowing the UVS to remain within some predetermined navigation corridor, maintain time-on-target, conserve fuel, etc. Collision avoidance is a last minute, emergency manoeuvre aimed solely at preventing vehicle loss or damage – and does not take mission completion into account [268].

adapt their flight paths (e.g. in response to a target detection) to increase the effectiveness of their overall search strategy and/or capacity to prosecute the targets. Consequently, even if only one UAV needs to deviate from its pre-planned trajectory (autonomously or under the control of an operator), the rest of the cooperative must also have the capacity to dynamically adapt their trajectories safely. Additionally, constraints must also be placed upon the degree to which the manoeuvring UAVs are allowed to adapt their trajectories (for instance to fly within safe performance envelopes).

The deconfliction algorithms must also be able to accommodate ‘blunders’, where one vehicle in a cooperative deviates from its intended path for unforeseen reasons. In this case, other UAVs must then manoeuvre to avoid collision and maintain adequate separation. Generalised solutions to coordinated airspace deconfliction control and assignment problems for UAVs are non-trivial, particularly when there are cooperation constraints imposed (e.g. communications ranges and schedules, minimum or maximum airspeed velocities, collision avoidance, sensor field of view, scheduling, etc). Moreover, these generalised solutions do not usually lend themselves to an extension of simple two-UAV control and assignment problems [156]. Regardless, the environment must be monitored and the appropriate state information collected and disseminated within the cooperative so that an estimate of the current situation (i.e. UAV position, velocity, and altitude) can be provided to evaluate the likelihood of conflict and guarantee a period of conflict-free trajectory for each UAV while it carries out its higher order tasks.

Based on a priori knowledge an operator must also designate the location, dimensions, and orientation of an area of interest and whether or not it is known to contain objects of interest, targets, no-fly zones, communications dead-spots, areas of threat and/or terrain obscuration, etc. Thereafter, based on the payload configuration, target locations, priorities, etc flight trajectories for each UAV involved in the mission must be calculated such that (say) the probability of detecting the targets is maximised. Ideally, given an area of interest, an automatic route planning algorithm will calculate search patterns for the group (e.g. using probability maps divided into discrete cells [48]) optimised under constraints such as: maximise the probability of detection; minimise the time to detection, minimise the number of UAVs required; maximise the robustness of the search to aircraft loss; minimise the amount of network traffic required; and/or coordinate the timing of specific UAVs activities.

There are several techniques designed to search spaces for optimum solutions under multiple constraints. Broadly, the techniques fall into two distinct classes: algorithms that only evaluate complete solutions and algorithms that evaluate partial or approximate solutions. However, most traditional approaches suffer from either being too time consuming or getting trapped in local minima. This is primarily an issue for the dynamically unfolding component of the UAV cooperative’s task as during the pre-mission planning phase, centralised and potentially even computationally intensive team-based coordination and decision-making techniques can be employed. To this end, the required temporal and spatial allocation of payloads, tasks, and communications resources can be

computed and iterated using centralised optimisation techniques and hierarchal command structures (e.g. exhaustive search, local search, greedy algorithms, branch & bound, divide & conquer, taboo search, simulated annealing, A* algorithm, evolutionary algorithms, etc) [188].

As the mission data must also be conveyed to rear command echelons these optimisation algorithms can also be run continuously as a means of (latently) evaluating the progress of the UAV cooperative against the specified mission criteria. There is, of course, a danger that these algorithms will provide false insight into the UAV cooperative's progress as the data observed by the UAVs in real time may only be shared locally (in real time) and a sanitised or processed version passed to the rear echelons.

Once the UAVs have commenced moving along their pre-planned trajectories, they must then autonomously and dynamically respond to the detection of targets and threats that might 'pop up'. These responses may be stimulated by either their own onboard sensors or those onboard other UAV within the cooperative. Alternatively, the UAV may need to respond to a change in priority imposed externally by the supervisors. In order to do this, each UAV must have goals and priorities assigned, which in turn requires a set of metrics that enables them (individually and collectively) to evaluate situations and events. These metrics enable the executive controllers to autonomously choose between competing goals, to assign resources to the task, and generate priorities that maximise pay-off and minimise cost. Ideally, both predicted and observed situations are evaluated so that resources can be allocated ahead of time and 'stand-offs' and conflicts can be avoided.

One of the main concerns for the distributed instantiation of dynamic control is that the controllers can potentially implement different strategies for the same goal (based on different perceptions observed by UAVs in different environments). Another is that in a large cooperative, the controllers may implement multiple copies of the same plan [242]. The latter is usually due to an inability in large cooperatives to share all the information observed by each of the members of the UAV cooperative. However, strategies for identifying these duplicate plans within the network and then pruning them exist [166].

In addition to the UAV responding dynamically (as individuals or collectively) within the wider cooperative to sensed opportunities and threats, in order to maximise the system's effectiveness, the cooperative may also need the capacity to form teams. The motivation for team-formation (or more accurately temporal-spatial task assignment) is to improve the probability of attaining specific goals – i.e. the detection or geolocation of an emitter or the capacity of another team member to carry out their intended function more easily. In order to keep the supervisory workload to a minimum, however, the team formation needs to be self-organising, such that the formation of a team is the result of forces acting within the cooperative and between the member UVS, as opposed to being imposed externally by an operator. Two other attractive properties of self-organisation are that any formation can potentially perform self-repair and that it can respond appropriately when unusual events occur.

The fundamentals of team formation require decisions on which UAVs are associated with one another and which teams are allocated to which targets or missions. These team-level goals must then be mapped to individual UAV roles – and there may be more than one role for each UAV within each team. Moreover, as the processing capabilities of humans are limited, particularly in time critical environments, the instantiation of the teaming architectures may need to be adaptive to accommodate variations in the levels of automation. Finally, mission completion criteria must also be established.

4.4.2 Multi-UVS Coordination

At present, the algorithms that successfully support teaming activities are well suited to tasks such as search, locate, track, identify, and engage targets because these tasks are quantifiable in terms variables that can be maintained within certain bounds. The challenge is to adapt these team controllers – that are themselves often the superposition of multiple controllers – to variations in noise, dynamics that are inadequately modelled and incomplete or uncertain sensor feedback in order to minimise the deviation of the specified variables from some predicted trajectory or end-state [9].

Additionally, given the number, variety and speed with which UVS are currently being introduced into capability around the world there is a need to extend our experimentally-derived understanding of operator-to-UVS control capacity to a more generalised theory or model. This will allow predictive modelling to take place within the relevant capability context and suitable architectures to be developed. However, if possible, we must attempt to do this by following the example of [76], [215] and [216] rather than just through the expensive and time-consuming use of technology-force insertion or human-in-the-loop simulation-based experiments. That is, we must ensure that the predictive modelling and the observations match, not just in relation to decision speeds (which is a common metric currently used), but also in regard to decision quality.

Most theoretical techniques for predicting cooperative behaviour depend on the expected time that a UVS may be ignored (known as Neglect Time) [73] before its performance drops below some acceptable threshold and the average time it takes to for a human to interact with the UVS to ensure it is still working towards its mission goals (commonly known as Interaction Time). Nevertheless, as the automation is not entirely reliable and failures do not occur at discrete, neatly designed intervals we must also account for the impact of the human decision-making process on the overall system performance. In other words, as most humans can only process cognitive tasks serially we must also allow for the time it takes for the operators to appraise the general situation (i.e. to notice that there is a problem within the cooperative) and the time it takes for them to gain situational awareness by focusing their attention exclusively on the errant UVS to discern its specific problems. We must also account for any time spent on distractions generated by other incoming problems or cognitive demands.

For example, using techniques developed by [76] we can then model the latency of human interaction arising from the overlapping arrival of UVS-related problems using queuing theory and – assuming the human is a single-server network – determine capacity predictions for human-UVS interactions. We may also extend these techniques by applying optimisation strategies that allow multi-constraint optimisation (e.g. evolutionary algorithms) and then use the outcomes of these computations to determine mission and cost-capability trade-offs between the larger, more sophisticated, platform-centric UVS options and the smaller, cheaper, distributed, network centric ones.

A current limitation with the predictive modelling techniques is their reliance upon assumptions or estimates regarding the interaction, neglect and waiting times. There appears to be very little experimentally observed data as it is difficult to measure and interpret. In this regard technologies that measure the psychophysiological relationships may be of use, but the techniques need further development and significantly more investigation is required [74].

One of key elements in realising the goal of multi-UVS coordination is the capacity of the cooperative to coordinate the actions of the different UVS that carry a heterogeneous mix of payloads; and a significant impediment here is that many existing multi-robot coordination algorithms elicit emergent behaviour such that the individual robots follow simple coordination rules rather than complex teamwork models or goals [211]. These techniques then break down because the UVS cannot explain their actions or role to other members of their team or the humans.

The need for the warfighter to be retained within the decision-making cycle means that in addition to the integration of the sensors and platforms, the information must also be combined, suitably manipulated, and passed to a rear echelon, where it is further integrated with applications that are of service to the user, such as geospatial information, track data and imagery, and visualisation and document management tools. This fused, value-added product must then be disseminated to users in near real time to allow the monitoring and redirection of the UAV cooperative, as appropriate.

In addition to this, there is also a need for multiple levels of feedback control [294], which in turn depend upon the capacity of individuals and the cooperative to measure and prioritise their performance and actions against a number of metrics (which need to be adequately defined in the first place), and their ability to communicate (in a meaningful fashion) the success of these endeavours, both internally within the group and externally to human operators, who may be geographically removed from their location and/or of a different command echelon.

In other words, a critical supervisory element for an autonomous cooperative of UVS is the feedback mechanism that allows the human operators to understand what, how, and why the system behaves like it does. Furthermore, research indicates that when human decision-makers are put in the position of passively receiving interpretations generated by data fusion and hypothesis generation aid machines they are less able to recognise emergent problems [154]. Consequently, there is a need to represent a range of levels and types of feedback control, which

in turn depend upon the capacity of individuals and the UVS (or their cooperatives) to measure and prioritise their performance and actions against a number of metrics and their ability to communicate the success of these endeavours, both internally within any UVS groupings or externally to humans.

In addition to the more easily identified end results, the quality of the system and the team processes (i.e. the performance of the UVS, their constituent components, and the human interactions with them) needs to be taken into account. In this regard, [138] has used metrics such as efficiency (the percentage of a task completed vs. the amount of resources required), stability (the variability between plans and the degree to which different operators respond to similar plans), the degree of user engagement, the level at which the user delegates control to the automation, the extent to which a user's mental model of the system predicts the effect of adjusting the weight of a particular control loop, and comparative performance (how well does the fully automated system work against the human-automated system).

Before leaving multi-UVS coordination, we return briefly to the concepts of mission complexity and scale (introduced earlier in the section on Human Systems Integration).⁶¹ Multi-scale Complex Systems Analysis (MCSA) [34] makes use of mission complexity profiles to specify the dependence of UVS mission complexity on the scale of action required. In other words, MCSA links the variety of possible ways in which multiple UVS (or their sub-systems) can act to the number of ways and the level of scale that a particular mission can be addressed. Thus, success of the UVS cooperative requires sufficient complexity at each required scale of action. In this regard, while high complexity in and of itself does not guarantee success, even well-designed UVS or their cooperatives will likely fail if they are insufficiently complex.

This has implications for the command and control (C2) structures for multi-vehicle cooperatives of UVS as it highlights the potential limitations of certain C2 structures. For example, as [35] [36] point out if we assume that each individual has finite complexity⁶², in an idealised hierarchy only the leader can organise and coordinate the entire cooperative. As a result, the coordination between these UVS is limited by the overall complexity of the leader, which in turn means that organisational behaviours of the cooperative are limited by the complexity of an individual UVS. Since coordinated behaviours are relatively large scale behaviours, this implies that there is a limit to the complexity of larger scale behaviours of the cooperative, which means that hierarchal C2 is effective at amplifying the scale of behaviour, but not its complexity [34].

By contrast, a distributed or networked C2 arrangement can have greater complexity than that of an individual element; although it should be noted that while such a network is not guaranteed to have greater complexity than its individual

⁶¹ Mission complexity is the ratio of the number of incorrect ways to perform a task relative to the number of correct ways to tackle it, where the more likely the wrong choice the higher the mission complexity. The scale of a task is the number of actions that need to be undertaken for successful completion.

⁶² For example, in terms of their processing loads or capacity to communicate over fixed bandwidths.

components, it is possible for this complexity to exist. For high complexity tasks, therefore, we are likely to consider hierarchal C2 systems inadequate and will look towards more distributed structures. The recent tendency of military organisations towards more network centric operations and organisational structures and the evolution of massively parallel computing architectures also suggests recognition of the limitations of the hierarchal control structures. It should be recognised, however, that while distributed command and control is often discussed as a panacea for problems of hierarchal control, it does not actually correspond to a specific control structure. As a result, distributing control in and of itself does not lead to effective systems or the solution of identified problems: it is the instantiation of specific distributed architectures that are effective in addressing particular problems that provide functional advantage.

4.4.3 Autonomous Multi-UVS Task Allocation

Task allocation is the problem of committing finite resources to a number of coherent tasks based on the comparison and selection among a set of available alternatives. For military UVS cooperatives this must be attempted within dynamically changing and real-world constraints such as finite time or where the achievement of one task is a pre-condition for being able to undertake another. The tasking commitments may be temporal, spectral, or spatial in nature and while many constraints are usually understood a priori, the situational awareness of the environment at any instant or from any given perspective may only be partial or conflicted.

In the context of multi-UVS cooperatives, the task allocation process attempts to address the fundamental question, “Which UVS-payload combination should execute which task when (and possibly how) in order for the cooperative to achieve its global goal?” Fortunately, this problem is also central to problems in economics, biology (e.g. the division of labour in insect colonies), network allocation strategies, and multi-processor scheduling design. As a result we may take comfort from the number, quality and variety of researchers in the field. Moreover, economics, game theory and operations research all use the concept of ‘utility’ (also referred to as fitness, valuation or cost), which is based on the notion that each individual can estimate the value (or cost) of executing some action. Depending on the context, however, the utility may vary from simple directly-observable metrics to sophisticated planning techniques. The only constraint on such measures seems to be that they must each produce a single scalar value that can be ordered for the purpose of sequencing the candidate tasks.

For instance, in our multi-UAV case above, we might assume that each UAV is capable of estimating both the accuracy with which its payload is able to geolocate the targets and a resource cost (i.e. time of flight or the number of UAVs lost to enemy action). We may then combine these measures using an appropriate function. Regardless of the method of calculation, it is important to try to include all aspects of agent and payload state and their environment relevant to the utility function. Even so, the utility estimates will be inexact due to sensor noise,

trajectory and target uncertainty, environmental change, etc, all of which limit the efficiency with which coordination can be achieved.

Many formal command and control models still tend to target medium to large scale systems composed of simple, homogeneous vehicles for use in relatively structured environments. Consequently, though simple and elegant, these models are insufficient for complex military tasks that require precise control. For instance, in our multi-UAV scenario, the UAVs will likely need to carry a heterogeneous mix of EW payloads; broadband ones to characterise the electromagnetic environment and cross-cue narrower band ones that are better able to provide high resolution spectral observations. The UAVs will also likely locate and track their targets using a number of different techniques, each requiring different observations (e.g. for geolocation of radar targets scan-ranging, triangulation through line-of-bearing, time difference of arrival, etc); they might also have different, but over-lapping, spectral views of the environment, different threshold sensitivities and so on. Similarly, if an enemy radar were to 'light up' unexpectedly and is identified as a missile control radar (a high priority threat) decisions must now be made about which UAVs or payloads must be tasked to work together to locate the radar as quickly as possible, but still taking into account their original tasking and objectives. What sensor scheduling strategies they should employ and what trajectories they should navigate (taking account of required accuracy, time on target, airspace deconfliction, no fly zones, etc) must also be accommodated.

If the UAVs have the capacity to communicate then they can inform each other about the value that they each place on the task (relative to any cost they may incur) and thereby reach a consensus with one another about who is best placed to carry out the task. It is then a matter of allocating the mission and trajectory deviations accordingly. However, in certain situations it may be prohibitive or impossible for the UAVs to explicitly communicate their task evaluations with one another. Also complicating the situation is that a lack of situational awareness may result in the UAVs not knowing what tasks they are likely to confront in the future. For instance, the closest UAVs may have been autonomously tasked to locate the missile control radar only to find that another, even higher priority radar then lights up that matches the specific spectral characteristics of their payloads requiring them to 'drop' their mission control radar task. Other UAVs, now ill-placed relative to their initial potential, must now take up the task of locating it.

Ideally we will be able to treat task allocation as a problem in optimisation. However, we must first decide what exactly is to be optimised. Preferably this will be 'system' performance but this quantity can be difficult to define and measure at any time, let alone during the execution of a mission, particularly if we include humans in the system. Moreover, as outlined in a previous section, when we select between alternatives the impact of each option on system performance is not usually known. Consequently, some kind of unifying performance estimate is required.

Clearly, the task allocation procedure must be adaptive, but under what conditions should a UAV take tasks when the opportunity arises and when should it ignore opportunities because experience has shown that a more appropriate

opportunity is likely to arise in the future? Does this affect the number or nature of UAVs required to undertake the task? Or should we use more intelligent task allocation processes to distribute ‘commitments’ to each UAV and their payloads within the cooperative. Finally, how much of this task allocation process should be handled by humans and how closely should it be integrated with higher level military command and control functions?

A comprehensive review of task allocation procedures is beyond the scope of this text. Nevertheless, three key issues include: protocols, strategies and algorithms. At a protocol level we need to understand what type of transactions are possible and devise our message structures, communications scheduling and strategies accordingly. When designing the individual UVS we need to devise strategies that best exploit these protocols. This can include the provision of feedback to users or internally to the UVS cooperative in such a way as to provide incentives to the task allocation process to adopt a preferred profile of behaviour. At an algorithmic level, this means actually solving the computational problems faced by real UVS. In other words, algorithms that recognise that solutions are infeasible and call for simplifications of the task allocation strategies because a particular computational problem is too hard to solve within a finite amount of time. We also need to understand what commitments are to be distributed; how these are to be allocated (temporally and spatially); what procedure or mechanism should be used to distribute them; and, what the objectives are behind the distribution/allocation.

Task allocation between multi-UVS cooperatives fall into one of two categories (three if we include hybrid cases): centralised or distributed. Centralised task allocation systems tend to be hierarchal in nature with computational loads that tend to be very high and that usually increase with team size: a single entity allocates the tasking commitments, possibly after negotiating over preferences with the UVS (the central entity often acting as an ‘auctioneer’ in a form of bidding, e.g. [45]). The biggest arguments against using centralised techniques are the potential for single-point failure, the necessary centralised computational capability for large numbers of UVS, and the difficulty of (dynamically) assigning the master-UVS.

Fully decentralised systems have their computational loads spread across a number of geographically dispersed UVS, and tend to be communications-based. In these systems, tasking commitments tend to emerge as the result of locally negotiated steps, which are often restricted by bi-lateral communication (although systems that allow multi-lateral exchanges have been developed). Such systems also often suffer from a parochial view of their environment and tend not to be amenable to analysis so their precise behaviour is difficult, if not impossible, to predict.

Resources: A central parameter in multi-UVS task allocation is the nature of the resources themselves: some are perishable (e.g. fuel, bombs, etc) while others are static (e.g. payloads). Of the perishable resources, some are continuous⁶³ (fuel) while others discrete (bombs). This often influences how the resource can be

⁶³ The allocation of continuous resources has been studied in depth in classical economics.

traded. We should also distinguish between different types of resource. For instance, they may or may not be divisible (network access) or indivisible (payloads); consumable (fuel) or perishable resources (time); and, resources that do not change their properties over time are usually referred to as static. Finally, some resources are better understood as part of the allocation process as whether they are sharable typically depends upon the tasking procedure rather than on the characteristics of the item itself (e.g. sensor scheduling).

Priorities: Another key parameter in the task allocation process is the prioritisation of objectives, both in terms of how they are determined and then how they are represented and communicated between the UVS. Essentially, priorities express the relative or absolute concerns or precedence of an individual or group of UVS when confronted with a choice between alternatives. They are often closely tied to the context and hence the level of automation proscribed to the UVS. However, we need to understand how the priorities and objectives are determined and what techniques are suitable for coding and representing these priorities in terms of their expressive power, succinctness and suitability to task.

There are several options for representing and mathematically modelling priorities: evaluation functions comprising an ordered scale of quantitative (or qualitative) values; an ordinal relationship between alternatives (X is preferred over Y for $X < Y$ but not if $X > Y$); a binary set of good and bad states (i.e. reductionist ordinal representation); and, fuzzy expressions that articulate the degree to which X is preferred over Y . The second option allows comparison of the satisfaction between alternatives but does not express priority intensity. Nor does it allow intra-UVS comparison of priorities. Qualitative measures allow a weak form of intensity to be expressed, but are difficult for UVS to interpret autonomously. On the other hand, the set of alternatives for the first and last options is a possible value of a given set of variables. In these cases the alternatives are huge and it is not sensible to expect the humans or the UVS to be capable of ascribing priorities against such a set. For this reason, there is a need to develop strategies, protocols and languages that allow compact representation of priorities and preferences.

Other key issues pertinent to complex task allocation include [174]:

- **Synchronisation:** Which tasks require intentional (as opposed to emergent) cooperation? What are suitable measures of cooperative behaviour to assess the quality of a task allocation within a given context?
- **Complexity:** What is the overall complexity of finding feasible and optimal solutions? How much of this process can be solved locally by each UVS and how much information needs to be exchanged between the UVS to achieve this?
- **Negotiation:** For multi-UVS cooperatives that rely upon distributed processing, what are the appropriate negotiation protocols and what are the most suitable strategies for employing these protocols? For those that rely upon centralised techniques, how can we devise efficient algorithms to support complex negotiation strategies?

- **Accuracy:** How do we devise negotiation strategies that force the UVS to report their priorities truthfully, both to reduce global complexity and to enable a correct assessment of cooperative synchronisation?
- **Implementation:** What are the best practices for rapid prototype development for specific applications? What constraints does the real world impose on theoretical models and how do different coordination strategies perform in practice? What is the practical impact of allocating infeasible task commitments and how computationally intensive are theoretically intractable results?

Regardless of the task allocation policy, however, most multi-UVS research has focused on the construction, demonstration and validation of working systems rather than the more general analysis of problems and solutions. As a consequence there are now a large number of architectures, many tested in working systems or in simulation, but the field still lacks a theoretical foundation that can explain or predict coordinated multi-UVS behaviour. In this regard, [116] developed a taxonomy for robot task allocation:

- Single-payload UVS (i.e. capable of executing only one task at a time) vs. multi-payload UVS (capable of executing multiple tasks simultaneously, but from a single location)
- Single UVS tasks (i.e. each task requires exactly one UVS to achieve them) vs. multi-UVS tasks (i.e. tasks require more than one UVS to achieve them)
- Instantaneous task assignment (i.e. the available information on UVS, tasks and the environment permits only instantaneous task allocation, with no planning for future allocations) vs. time-extended assignment (i.e. more information or predictive models of what tasks may be expected to arrive in the future are available)

This taxonomy allows more formal studies to be conducted as it characterises a range of multi-robot task allocation problems, providing the possibility of provably optimal solutions for the simpler cases and insight into the more complex cases.

One final point regarding multi-UVS task allocation, the current approach to command and control is largely human and platform-centric. As a result, the scale and nature of interactions between warfighting entities has historically precluded an autonomous coordinated response to threats – except that instantiated through human-to-human interaction. This is particularly true for smaller defence forces, although the network-centric paradigm is changing this. In contrast, in an autonomous, multi-vehicle UVS environment, where each platform is potentially presented with an abundance of information derived from a range of external sensors, the assets must interpret, purify and apply this information in a manner that prevents rapid error propagation before allowing self-synchronisation of any response. Furthermore, this non-trivial undertaking must be achieved within a framework of finite resources so that the systems may autonomously coordinate their response options. It is reasonable to assume that the early instantiations of such enterprises may have a capability edge in data processing, fusion and even operational tempo, but they may not equate to the

levels of ingenuity, unpredictability and sophistication enjoyed by human-to-human command and control structures used by adversaries. As a result, such systems may have vulnerabilities that are exploitable.

4.4.4 Multi-UVS Navigation, Localisation and Mapping

For UAVs, with the exception of taking-off, landing and manoeuvring on the ground, navigation is usually relatively straightforward, albeit subject to the requirements of ‘see-and-avoid’⁶⁴. On the other hand, the navigation environments of military UGVs, USVs and even UUVs are usually unstructured and therefore much more complex and cluttered. Additionally, although many civilian UVS can theoretically navigate using EO sensors, battlespace environments can be expected to be opaque to such modalities (at least part of the time). As a result, all-source navigation estimators [11] that fuse multiple sensing modalities and, on the basis of those estimates, select sensing and navigation options that optimise information gain and UVS mission goals will be required.

The basic navigation challenge is to determine the location and orientation estimates of the UVS relative to an unknown number of environmental features (usually without an initial estimate of either), the location estimates of these features relative to the UVS, and the observed variation of these features relative to aspect, occlusion, UVS motion, time, etc. As a result, translating sensor data into maps for the purposes of navigation and mission-execution is an absolute requirement of a persistently autonomous UVS. Furthermore, as it is an integral component of the UVS control system, any errors in the world map reduce the reliability and safe-operation of the UVS and hence its potential utility.

There exists a large body of work that addresses these problems as they pertain to both single and multi-UVS navigation. One of the most successful techniques is Simultaneous Location & Mapping (SLAM) [166], which concurrently builds feature-based maps of UVS environments and obtains estimates of UVS location. These have been extended using machine-learning techniques for multi-agent systems, hybrid algorithms for multi-UVS control, multi-UVS localisation and map-building, and distributed sensor fusion [90] [238] [260] [266] [271].

⁶⁴ There is a general need for UAVs to fly in civilian, uncontrolled airspace. In order to achieve this, they will need to meet the requirements for visual flight rules at an equivalent or higher level of safety comparable to the ‘see-and-avoid’ or ‘detect, see-and-avoid’ requirements for manned aircraft (see Legal Issues). Detect, see-and-avoid is the process of trying to detect obstacles in the path of a UAV, determining whether or not they pose a threat, and, if necessary, taking measures to avoid them. There are a range of technologies (e.g. TCAS, ADS-B) that partially satisfy these requirements, but they only aid in avoiding cooperative aircraft. Other technologies (e.g. radar) are likely to be of use due to their all weather capabilities, but the weight, size cost and power of the equipment mean that they are unlikely to be considered practical solutions for small-medium UAVs. Such a system has two basic requirements: an ability to detect objects early enough to avoid them and an extremely low false alarm rate. It is also probably a requirement that the system function at a level superior to that considered acceptable for a human being as we have a tendency to accept ‘human error’ as a reason for failure, but expect autonomous systems to have a much lower failure rate.

More recently, techniques have been developed that are truly distributed and some of these techniques now take advantage of the properties of the distributed cooperative to achieve mapping accuracies unachievable with single UVS [261]. The majority of the work, however, operates predominantly in two-dimensional environments and relies upon environmental perception based on EO and LADAR. Regardless of the implementation, however, for multi-UVS navigation there are two broad classes of algorithm (or three if hybrids are included): one in which every control input and observation is passed between the UVS within the cooperative, and another where all the information is sent to a central node or 'mothership' running a single filter that estimates all the vehicle and feature locations [221].

The first of these techniques places significant bandwidth and scheduling requirements on the system and individual vehicles, while the second requires the central node to be aware of its own parameters (i.e. speed, orientation and position), as well as those of its subordinate units. As the structure of the SLAM navigation problem is characterised by monotonically increasing correlations between landmark estimates [91] and decoupling the state space is non-trivial, this places considerable computational load on the central node as the full covariance matrix must typically be updated with each prediction and observation of each UVS. Nevertheless, even though the mothership philosophy suffers from a systems level vulnerability of having one node that can be targeted as a single-point of failure, computationally effective algorithms capable of processing several thousand features in real time on high-end PC hardware have been developed [275].

A better approach is for the individual vehicles to build independent maps of their environment and for these maps to be fused together to form an aggregate, global map [291]. In a patch-work fashion, each vehicle can then add the current estimate of its local environment. The map data must then be correctly associated, both map-to-map for individual and between vehicles. This is usually relatively straightforward, even if a UVS joins the cooperative [292] as long as the location of the UVS is known either in global coordinates or relative to the other UVS. Under these circumstances, any new maps can simply be correlated to the global map. Furthermore, even if the location of the UVS joining the cooperative is unknown, the process is still tractable by building a map of the local environment and using this information to determine the relationship between the global map and the new local map [30]. It is also possible for the UVS to build maps that include feature estimates of the other UVS in their state estimators if the UVS are in the appropriate sensor's field of view. Once the correspondence has been established, the relative position between the reference frames may then be estimated. These techniques are also robust to communications that suffer from latencies or outages.

It should be noted, however, that all measurements have uncertainties so the location of the UVS and targets are only estimated as probability density functions pertinent to the regions where they are expected to be. As a result, single and multi-UVS mapping and localisation techniques tend to rely on the recursive use of distinctive environmental features or landmarks that, when revisited, aid the

UVS localisation process. This in turn also helps to keep track of the location of all landmarks. In order for features to be used as landmarks, however, their location needs to be estimated reasonably accurately. Consequently, multi-UVS map correspondence and location uncertainties arise from the use of environmental features that move, noise on sensor observations, and observations of features from differing spectral, spatial or temporal perspectives [275].

Motion tracking and estimation has been added to a number of mapping techniques [66] [197], although this is largely focused on sensors that provide both range and bearing information. For example, SLAM with Generic Objects (or SLAM with GO) [296] allows the addition of motion mode information (stationary, moving, move-stop-move, etc) for the landmark, although this has to be learned from the observations. The technique is straightforward, but computationally intense and not yet real-time. In another approach, called SLAM with Detection and Tracking of Moving Objects (DATMO) [66], each new moving object gets its own statistical estimator (typically a Kalman filter) and motion mode [109]. Vehicle state estimation then takes place separately and is used to update the Extended Kalman Filter (EKF) used for SLAM. This runs faster than SLAM with GO and is more suited to real-time implementation.

As the research stands, using the appropriate sensor modalities, as well as moving object detection and track initiation, the more advanced navigation techniques are able to determine when moving objects have coalesced, moved outside a sensor's field of view, or been temporarily occluded by a stationary object. They are also robust to long sequences of data and can adapt to false measurements and their extension to multiple vehicles has also been achieved (e.g. [273]). However, as yet, the techniques have not been extended to bearings only sensor modalities; are subject failure as a result of false observations arising from platform motion; and often struggle to accurately classify slow-moving objects or those that are temporarily stationary.

Image-based navigation, and scene and structure estimation derived from it through the fusion of external sensor information (e.g. INS), is also now solvable in real time [79] [226]. However, environmental dynamics have a deleterious impact on these techniques. That said if the dynamic features are characterised correctly they can be used to aid the mapping and navigation process, and vice versa [66].⁶⁵ When range observations and a priori information about the non-stationary objects are not available, however, it is not possible to determine their trajectories uniquely unless two or more sensors are used. When such observations are made from multiple moving platforms that do not have other means of localisation, stationary environmental features must also be used [109].

At present, most maps are usually classified in terms of statistical estimates of features described by data clouds, geometric returns, RGB pixel intensities, or through the use of occupancy grid-maps that are regularly updated. A more condensed approach relies on classifiers that interpret multi-modal sensor data in

⁶⁵ The algorithms must have already been initialised [81] and any recursive loop-closure already performed [81], which for bearings-only techniques is non-trivial. As a result, the mapping/localisation and motion-tracking problems have been currently only been solved separately and then integrated using range and bearing information.

terms of higher level descriptors such as ‘eucalypt tree’ or ‘bitumen road’. At present extracting such descriptors robustly and in outdoor environments is difficult and frequently dependent on aspect, background, lighting conditions, context, etc. Furthermore, the sheer number and diversity of potential objects – and hence the resultant searching of any hypothesis trees – means that a priori knowledge is usually required in order to classify any objects swiftly and correctly. Additionally, high resolution observations of the objects, preferably at long range (and hence large quantities of data), are also a pre-requisite for most techniques.

Currently, the error analyses associated with such feature-extraction techniques are also not well-understood. As a result, the efficient coding of these descriptor-based maps and their integration with navigation estimators are yet to be achieved. Such techniques have been researched, but do not yet run in real time [117]. Furthermore, the techniques usually require spectral and geometric correspondences to be formed so the features can be ‘fingerprinted’. There is a difference, however, between the processing requirements of an algorithm that can (say) classify terrain that is sufficiently flat and devoid of obstacles for a UGV of specified mobility characteristics to traverse, versus one that can classify environmental features completely at the descriptor level. While the latter is possible on current PC-based architectures, there is a need to use 0.1-1.0TFlop processors if such operations are to be executed in real time [117].

4.4.5 Capability and Systems Integration

Multi-UVS integration – and particularly when it crosses environmental domains – has many technological impediments. Another challenge, however, is perhaps best illustrated using the following example.

Many countries acquire their military capabilities from overseas. Typically, such acquisitions might include UAV-borne ISR or strike systems, which are effective at detecting and neutralising concentrations of enemy forces on the ground, but have much more limited effectiveness when an adversary blends with his surroundings. Hence, while major force concentrations might be eliminated, smaller enemy groups that can protract hostilities may remain. As a result, an acquisition focus might be given to the provision of theatre or tactical-level tools for optimally selecting, deploying and managing sensor assets or the development of onboard control, coordination and decision support mechanisms for multiple manned and unmanned force elements (e.g. autonomous UGVs integrated with the ISR output from the UAV feeds).

Such technology would likely be one of the key outcomes of a control and coordination research program, which could have application across several major capability domains. Such developments would also probably involve information integration for manned and unmanned systems; a key element of sensing and data fusion research, again quite possibly applicable across several capability domains. The integration might also involve an analysis of the appropriate reliability and resource allocation issues pertinent to the provision of a persistent autonomous presence on the battlefield; the major focus of research in persistent autonomy.

The indigenous development or adaptation of such technologies and operational concepts to allow the integration of these unmanned and manned systems to support operations in complex, hazardous environments might be a national priority. However, it is unlikely that knowledge of the systems at the level required for close integration or multi-UVS cooperation would be shared between nations without significant risk or cost. Solutions may, of course, be available from other overseas vendors, or by integrating the systems less tightly, however, both these solutions are likely to be unpalatable for reasons of cost or sub-optimality.

From the technology integration perspective, therefore, the designation of a single lead agency for UVS to oversee the general and cross-cutting matters pertaining to automation and introduction into capability could be beneficial. By way of example, such an agency might also oversee such matters as systems engineering, life-cycle cost management, software engineering, the development of an effective assessment methodology, and the use of modelling and simulation assessment tools as many of the lessons learnt in one environment will be directly translatable to others.

Unfortunately, the likely acquisition strategies of many defence organisations, which are largely still platform and/or environmentally based, means that it is more likely that Navy will take a lead on UMVs, Army on UGVs, and Air Force on UAVs (and for some larger defence forces each agency will likely acquire its own UVS in each domain). Given the likely focus on operational exploitation and the individual agencies' experience in each of these domains this is not a bad thing, it simply misses the opportunity to enforce a more cross-cutting systems discipline on the generic requirements of autonomous UVS.

However, as these intra-service benefits are not yet well-articulated, it may reasonably be argued that these factors are trumped by the need to exploit the service-specific requirements of UVS; just as they are for other environmentally cross-cutting endeavours such as ISR and EW. Furthermore, without the clear and focused requirements advocated by the single-service or platform-focused capability initiatives, the process is likely to suffer from diffusion and incoherence. Clearly, it is likely to take a strong advocate in high office to advance any such notion. If such an agency were ever stood-up, it should take on the role of:

- Identifying gaps in capability that can be filled by UVS,
- Identifying technology shortfalls in autonomous systems and UVS,
- Influencing the development and assessment of UVS-related operational concepts
- Providing support to UVS planning, investment, and programs,
- Influencing the direction and level of UVS-related technology effort, and
- Developing and fostering cross-environmental UVS technologies and systems

That said, the evolution of military systems towards high-tech networks of automated capabilities that are responsive to a range of information sources, and the commensurate move away from the use of humans as the command and control 'glue' traditionally used to instantiate such enterprises, will likely result in this systems integration becoming a problem common to many military technologies, not just autonomous UVS.

Chapter 5

Force-Integration of UVS

In this chapter, attention is drawn to the fact that some changes intended to increase productivity have resulted in performance declines because the technical systems and the social systems into which the technology has been embedded have been misaligned. In particular, since military UVS are expected to change and probably reduce the involvement of humans in certain tasks, and as modern military capability needs to respond to the threat of conventional force and the challenges of asymmetric conflict in terms of Effects Based Operations, a number of questions remain regarding what realistic effects can be expected from military UVS, how these systems should be integrated into a force, and how we might then quantify such benefits in the complex and non-linear battle space of the future. The intent of this chapter then is to try to highlight a number of the considerations, and to describe a process by which we are able to test the cost-capability trade-offs and value proposition for disruptive, next-generation UVS by testing their viability and evolving new operational and deployment concepts.

5.1 Capability Challenges

“The acquisition of military capability is driven by the development of new technologies (technology push) and the definition of new requirements by the operational users (requirements pull). The challenge is to match the requirements of users, who don’t understand the strengths and weaknesses of nascent technologies, to the solutions proposed by the technologists, who don’t understand the military requirements” [237].

In other words, due to the evolving nature of the modern battlefield and the revolutionary potential of autonomous and unmanned vehicles neither the capabilities of the technology nor the military requirements of the forces likely to rely upon them are fully-understood. Furthermore, as our adversaries are also investing in this space we must also try to understand and exploit their vulnerabilities. Fortunately, military organisations are inherently prudent and realise that it takes more than just technology development to field operational capability that is both supportable and affordable. As a result, the decisions made regarding the integration of autonomous UVS into the battlefields of the future will be based on a combination of technical competence, military advantage, commercial good sense and legal precedent, all matched against defence

capability gap analysis, concept assessment, product demonstration and in-service evaluation.

At present, however, there is considerable scepticism within some in the military community regarding the capacity of the more sophisticated UVS to withstand the rigours of the battlefield. This scepticism is often compounded as the benefits of these systems are sometimes stated in terms of replacing soldiers in the force structure, rather than aiding them to perform their missions, particularly as experience has shown us that automation does not supplant human activity, but changes the nature of this activity. UVS with higher levels of autonomy will start to be accepted when their roles, basic utility, robustness, survivability and value are identified and understood in quantifiable terms.

In addition to this, a number of researchers have found that technological innovations alone are unable to explain some of the observed improvements in systems performance for technologies introduced into the workplace, with some changes intended to increase performance resulting in performance declines; more detailed analysis revealing that the performance improvements actually resulted from an alignment of the technical designs and the social systems into which they were embedded [174]. As a result, just as it takes more than the development of sub-component technology to field operational UVS capability that is both supportable and affordable,⁶⁶ so it takes more than the development of new systems and their assessment against strictly independent criteria to realise the full potential of UVS.

To understand the value of UVS, we will need an observable framework that allows assessment in quantifiable capability terms against endorsed and projected requirements. This implies a systems engineering approach that allows us to develop robust use-cases, sound concepts of operation and appropriate organisational structures. Furthermore, we must do this in a manner that matches our understanding of future UVS potential against projected opportunities within realistic organisational structures and operations.

This will ultimately lead to an increase in the autonomy, performance, and affordability of UVS and a core investment strategy, thereby lowering the risk for transition to any future UVS warfighting concepts. The major capability challenge then will be to capture the next-generation capabilities, test the viability of these new technologies and evolve the new operational and deployment concepts (and cost-capability trade-offs) that emerge due to the availability of the new technology mixes as UVS technologies (and hence opportunities) emerge and evolve at a rate that is about an order of magnitude faster than the defence acquisition cycles can handle [102].

This force transformation process will require contributions to be made through a synergistic combination of concept demonstration, experimentation and analysis. In fact, the development of novel technologies simultaneously with doctrine and concept exploration through an experimentation-focused approach that places innovative UVS in the hands of the warfighters in the field (referred to henceforth

⁶⁶ In other words, it also requires a systems engineering approach to optimise the synergistic performance of the system (or system of systems) so that the assigned tasks can be accomplished effectively and efficiently.

as Technology-Force Insertion) will serve to validate and inform the parallel and complementary analytical approaches, which are reliant upon operations analysis and simulation environments and permit exploration of the broader operational concepts [102]. This process then allows us to measure how well a UVS or a human-UVS team performs within a military framework, achieves a specific task, etc, as well as for feedback to be applied, both in terms of capability guidance and technology stimulation.

As with the many technical challenges, a number of teams have dedicated considerable effort to identifying particular potential capability benefits of UVS and strategies for achieving them (e.g. [141] [202] [230] [234] [237] [245] [244] [268] [276] [277] [280] [278] [279] [280] [281] [282]). These documents have resulted in the creation of technology and capability roadmaps. Two in particular [278] [279], the U.S. Department of Defence (DoD) Unmanned Systems Roadmap 2007-2032 and Unmanned Systems Integrated Roadmap 2009-2034, acknowledge that they will aim to develop and employ “an increasingly sophisticated force of unmanned systems over the next 25 years and that this force must become seamlessly integrated with manned and unmanned systems.” The document also recognises that the U.S. “will pursue greater autonomy in order to improve the ability of unmanned systems to operate independently, individually or collaboratively, in order to execute complex missions in a dynamic environment.” Given the leadership of the U.S. in this area, it is likely that other nations will follow their broad direction.

The US DoD’s aspirations of “fielding transformational capabilities, establishing and implementing joint standards, ensuring interoperability, balancing the portfolio, and controlling costs” are articulated as broad goals [278]:

- Improve the effectiveness of combatant commanders and coalition unmanned systems through improved integration and joint services collaboration.
- Emphasize commonality to achieve greater interoperability among system controls, communications, data products, and data links on unmanned systems.
- Foster the development of policies, standards, and procedures that enable safe and timely operations and the effective integration of manned and unmanned systems.
- Implement standardized and protected positive control measures for unmanned systems and their associated armament.
- Support rapid demonstration and integration of validated combat capabilities in fielded systems through a flexible prototyping, test and logistical support process.
- Aggressively control cost by utilizing competition, refining and prioritizing requirements and increasing interdependencies (networking) among DoD systems.

There are strong cases for each one of these goals and from a capability perspective they represent some of the key practical objectives. However, none of these objectives address either the operational requirements for UVS or the

operational concepts for employing them. Furthermore, neither the goals nor the roadmap documents refer to methodologies, metrics or assignments by which achievement of these stated goals might be measured.

In fact, very little is reported on systems approaches, cost-benefit analyses, procedures or analytical tools that might be used for establishing a case for UVS. Additionally the US Government Accountability Office (GAO) reported that there was no apparent link between the cost of systems and munitions and their performance in conflict [113]. Consequently, as the expenditure on UVS increases such cost vs. performance-based analyses will be necessary for providing a sound basis for objectively assessing the operational effectiveness of UVS while simultaneously establishing realistic performance expectations relative to the emergent technologies and embryonic concepts of use.

Furthermore, to date many efforts could be described as ‘technology-driven’ rather than ‘system-driven’. For example, it is reasonable to assume that each successive demonstration of a UVS program has advanced the technology state-of-the-art relative to the previous iteration, but without a statistically valid or fiducially referenced data set there is no way to know. Demonstrations alone do not provide statistically significant data to assess the maturity, capabilities, and benefits of UVS technology, either at the individual functional component or at the systems or systems-of-systems level. For this we need systems-level evaluation for benchmarking the algorithms, sensors, architectures, systems integration techniques and so on. Additionally, if we are able to develop test and evaluation metrics and systems evaluation techniques simultaneously with the technology developments (or better still ahead them) we will not slow down the acquisition and introduction of UVS into service. Before considering this, however, let us briefly discuss two other key issues pertinent to technology-force integration: policy and training.

5.1.1 Policy Considerations

In addition to establishing the value of UVS in capability terms, the development of a robust use-case and sound concepts of operation, and addressing the many legal challenges outlined in the later Chapters, the degree of autonomy that can be introduced into capability will be a function of the state of the technology, the degree of system reliability and the conditions in which the UVS must operate relative to national policy considerations.

Unfortunately, at present there does not appear to be much policy relating to the use of autonomous UVS in the contemporary battlespace;⁶⁷ although absence does not necessarily imply a need. In particular, there appear to be no well-thought out concepts for the use of UVS in conjunction with manned assets. That said, like many technologies, autonomous UVS will be procured on their capacity to support specific military objectives so a set of appropriate ‘rules’ will almost certainly evolve as the systems are integrated into operational environments. Tactics, doctrine and concepts of use will then be captured accordingly for the purposes of training, capability planning, and so on.

⁶⁷ That which does exist relates almost exclusively to remotely operated systems.

For example, several analyses (e.g. [237] [278] [281] [282]) have identified the key logistics tasks likely to be carried out by UVS (e.g. anticipatory sustainment and improved distribution, UGV movement and tactical behaviour, improved delivery of supplies in non-contiguous operations, and improved inter-modal platform technologies and techniques). Furthermore, there is also speculation that autonomous air and ground vehicles may move warfighters and their supplies around the battlespace. While UGV and UAV-based delivery of ~ 100kg payloads is clearly feasible, such an approach may face many policy challenges. In particular, the safety of troops and civilians in close proximity to these vehicles, (dynamic) battlespace command awareness, integration with manned systems, and the timely reaction to unexpected conditions such as weather, obstacles and tactical conditions, etc are all issues that will need to be resolved. Other issues might include the theft of unattended supplies, sensor and network vulnerability, and the maintenance of specialised equipment [127].

Furthermore, many analyses also discuss the potential benefits of using UVS to evacuate casualties or to carry out remote, tele-operated medical procedures.⁶⁸ As a result, policy decisions may be required in relation to when casualties may be operated on using such systems or attended to by trained medical personnel. Additionally, even though it may be a relatively straightforward task to monitor a patient's vital signs (i.e. temperature, pulse, blood pressure, etc) using a manipulator arm or stretcher, it is much harder to know when and how to manipulate a wounded soldier in preparation for evacuation. Furthermore, a number of policy issues may need to be addressed if UVS are to evacuate wounded soldiers as some military forces currently require evacuated casualties to be accompanied by trained medical personnel.

5.1.2 Training and Maintenance

One of the most common problems for the smaller UVS is that they fail, not through a manufacturing or design flaw, but through routine dynamic loading, collision with an obstacle or operators using them beyond their design limits or performance envelopes. Clearly, improving their general robustness or providing real time feedback on UVS performance criteria relative to their component/system design maxima would help, but both would have an impact on UVS sophistication, weight, endurance, etc; and ultimately cost. Moreover, no matter what the level of system robustness, deployed users will probably find a way to exceed it. This is best demonstrated through the growing popularity of smaller tele-operated UVS which is leading to a proliferation of systems in military theatres and the often sacrificial

⁶⁸ At first glance it would appear that these technologies could be readily transitioned from those that exist within civilian hospitals. However, this equipment is large and costly, requires specialist surgeons and is usually dedicated to specialised procedures. Furthermore, just as with the control of automated weapons, bandwidth, latency, communications outages, interference, etc are all very real concerns for the mobile and tactical operation of such equipment. The surgeon also lacks sensor modality in terms of feedback (and his visual examination frame is 2-D); all of which also give rise to questions of viability.

manner in which the units are employed.⁶⁹ The end result is that as units are damaged, for practical reasons cannibalisation and make-shift repairs are commonplace, and specialist training and technicians' workshops are then needed to maintain the systems in theatre.

For the larger UVS this is less of an issue in that large numbers of dedicated staff and complex logistics chains are needed to sustain the capability anyway, which (in the case of UAVs) are frequently removed from the front. The support requirements for smaller UVS are often less clear. Typically, it is provided in the form of contractor support, but this is often precluded, either by policy impediments or by the diversity of UVS units deployed and the number and nature of the technical faults that develop in each proprietary system. Ideally, such support requirements will be considered early in the system's design and development as integration late in this process can be costly or impossible. Furthermore, determination of these requirements is usually based on experience or a review of previous systems and adjusting support requirements accordingly. The difficulty with determining future support requirements for autonomous UVS is that aside from the platform there is very little historical operational information to go by. There is, of course, indirect information from manned platforms and a number of tele-operated UVS, but these do not necessarily reflect the relevant considerations as they pertain to manpower, skill-level, training needs, maintenance, etc.

Fielding UVS will require an analysis of vehicle, system and force structure trade-offs that will impact systems design and life-cycle costs as well as deployment and manning costs. For example, design and life-cycle costs might be influenced by the selection of an energy source required to achieve some mission endurance, system size, weight, and power level. Furthermore, reliability and maintainability considerations may impact on deployed organisational structures and compositions, particularly for the smaller UVS which will almost certainly be used sacrificially, where current life-cycle support issues and procedures may need to be modified. For example, it may be necessary for teams of technicians to 'roam' the battlefield to recover and repair UVS in situ, rather than having them brought back to a maintenance depot [127].

One obvious solution would seem to be improved training for the operators. However, this is not without its challenges as, in common with most other organisations, militaries must operate within the finite financial and personnel

⁶⁹ Furthermore, it is becoming clear that not all companies have the capacity to support the development, production, replacement and post-production requirements placed on them during a sustained deployment. Many companies that have supplied the large numbers of small UVS to the US military, have demonstrated very capable production lines and post-production support for the small units that they have supplied to operations. However, these companies appear to have a current capacity to supply around 30-40 units per month, and can apparently increase production to 80 units per month (or higher) on demand [237]. As militaries around the world recognise the benefits of UVS technology, given the trans-national nature of the Defence industry, even at higher rates of production, significant delays might result for even a modest acquisition unless the products and support services are available to the buyer indigenously. Given the sacrificial manner in which these units are likely to be used, this limited production capacity may also have cost or capability implications.

resources of their mandate. In other words, when forces deploy into theatre they are subject to quotas and must comprise a mix of skills. The conundrum is then whether to allow 'generalists' who by definition may only have a scant knowledge of the specific capabilities of any given UVS but can be deployed on multiple military tasks or 'specialists' who have a more detailed knowledge of the UVS and its capabilities, but may not easily be re-tasked. These specialist tasks will include mechanical, electrical and software maintenance, although requirements may be simplified through the use of modular architectures and remotely connected technicians and support staff. Alternatively, they could be made more complex by the spectrum of unique demands placed on the technicians.

In the longer term, this will probably require the development of a suitable training plan and a standardised curriculum within military organisations, which will of course need to be maintained. However, it may be that the knowledge and training regimes required to field autonomous UVS differ only marginally from the knowledge, skills and abilities required to operate, repair and maintain other sophisticated systems such as avionics, electronics, computers, and manned platforms. While likely to be evolutionary, these training regimes will vary across the spectrum of UVS from tele-operation to the supervision of a network of differing UVS. The mix of skills required to maintain such systems is not yet clear. However, it is clear that there will need for multiple levels of maintenance and the less capable the UVS, the more capable the engineer at the front will need to be in diagnosing hardware and software faults. This may be eased by the development of autonomous diagnostic techniques such as HUMS – or, ultimately, self-repair technologies – that are capable of running of PC-based architectures.

The main policy objective of any training strategy will presumably be to ensure that the warfighters and their commanders have sufficient training on the use of these systems prior to their deployment in theatre. As a result, training with autonomous vehicles will need to commence well before the operators and systems deploy into theatre. This will require the establishment of a facility that allows real, virtual and constructive simulation and training, which will need to support the individuals, the commanders, small units and any staff-trainers. Furthermore, these collaborative strategies must be synchronised to avoid duplication. While somewhat self-evident, the cross-environmental nature of autonomous UVS may mean that many militaries ultimately make use of UVS from each domain (i.e. air, land and sea) that have functional components in common.

As vehicle automation becomes more sophisticated and integration (say) between air and surface platforms more routine there will also be a need for synchronisation across the environmental domains in training approaches and maintenance regimes. Furthermore, one operator may supervise several UVS in multiple domains. As a result, consideration will need to be given to supporting these UVS through a common training program. In each modality there are obviously unique environmental factors that influence or limit the support requirements. Initially militaries may be forced to support each through different

capability structures. However, over time these arrangements may place performance-limiting constraints on each class of vehicle and further needs may arise as the UVS are deployed operationally. One outcome of this is the need for interoperability, usually facilitated through adherence to standard message formats, architectures, and data protocols (for example the NATO Standard Agreements, STANAG 4586, STANAG 4609, or the Joint Architecture for Unmanned Systems, JAUS). While it is beneficial to analyse systems against their ability to adhere to such standards and their ability to be adapted to other unified standards, slavish or doctrinal devotion to them for their own sake will probably not be helpful.

5.2 A Systems Approach to UVS Analysis

“If you cannot measure it, you cannot improve it,”

Sir William Thomson (Lord Kelvin) 1883

Clearly, the most compelling reason for using UVS in a military context is that they can save lives by taking on some of the more hazardous missions currently undertaken by warfighters. However, this view in and of itself does not justify their insertion into capability. Moreover, we cannot assume that UVS effectiveness is synonymous with an overall enhancement in military capability. For example, it is conceivable that a UVS could be employed effectively and the related combat aspects fail for other reasons, or vice versa.⁷⁰ Hence, in order to understand the impact of UVS on the battlefield, one has to measure their effect on military operations, along with the effect of the components of the constituent systems. By way of example here are two simple examples.

Like UAVs, one of the most obvious benefits of UGVs is that they can extend the reach and access of the users into regions intentionally denied them. This can in turn bring them intelligence that can save lives. However, UAVs typically have ranges of tens or even hundreds of miles and an endurance cycle measured in hours. On seeing one an adversary may be aware of what it is and may take cover, but it does not indicate the presence of hostile forces in the immediate area. This is not the case with (say) the Dragon Runner UGV,⁷¹ which must be operated locally by an individual using a control pad and a video monitor (which in turn diminishes the operator’s capacity to maintain his own situational awareness). Consequently, the use of this UGV (in the wrong way) may increase rather than diminish the overall threat to users or the ground forces they are intended to protect.

⁷⁰ For instance, UVS can achieve their objectives simply through luck or fail having performed perfectly.

⁷¹ Dragon Runner is a small, four-wheeled, rear-wheel drive, front-wheel steer, man-portable tele-operated UGV designed to increase the local situational awareness of dismounted ground forces. It gives tactical ground units the capability to “see around the corner” in an urban environment.

Alternatively, in a US analysis that considered a deep strike scenario [175], in which UGVs similar to Gladiator⁷² were deployed deeply and aggressively and the Forward Observers (FO's) held back, the dynamics of the battle were changed significantly in favour of the Blue force. This was because the UGV's (or more importantly the sensors they carry) allowed detections of an adversary to occur earlier and more deeply in an engagement, thereby allowing fire to be called in deeper. This shifted the attrition further out, and made the close-in, direct fire battle more manageable. In this analysis sensor quality, property and distribution have a profound effect on the outcome of the battle: the better quality and better placed sensors made detections at longer ranges, which led to more detections and larger volumes of fire placed over deeper targets.

UGV size and speed also affected the outcome in the Gladiator study with UGVs better off in stationary hiding positions than attempting to withdraw slowly. As the UGV speed and capability was increased to a level comparable to the speed of the attacking force, however, many of the UGVs were able to standoff, which resulted in much lower losses [235]. The detections of enemy activity were influenced by particular scenarios, which influenced the preferred sensor height, sensor capability, etc, which in turn influenced the detection ranges. Moreover, the signature of the UGV also affected its capacity to be detected by an adversary and the link between the UGV signature and its size affected its capacity to traverse certain terrain, and hence to carry out certain missions.

Consequently, we need a systematic process for determining the value of UVS, which measures the capability integration of UVS holistically and not simply as a sum of a number of discrete elements. That is, a methodology that allows us to compare UVS technology with current (manned) approaches and UVS of fundamentally different design; for example, those that are tele-operated with those that are network centric autonomous and those that rely on adaptive learning techniques or reactive behaviours to those that rely on executive or hierarchal planning techniques.

Furthermore, while the goals of this analysis will be to understand the value of UVS, to assist in the development of investment strategies, and to stimulate and guide the technological development, a by-product should be to avoid the development of tactics, doctrine and organisational procedures that prevent the effective use of UVS. For example, it is common for technologists to introduce automation in a manner that places UVS operators in roles for which they are not well-suited (i.e. watching real-time video feeds from UAVs for protracted periods;

⁷² Gladiator is a small-medium sized UGV that performs scouting, ISR, direct fire, and personnel obstacle breaching missions. Essential Functions of the Gladiator system include: day/night remote visual acuity equal to that of human using current image intensifying or thermal devices; battlefield mobility capable of supporting dismounted units in all environments, including MOUT rubble; modular design and incorporation of standard interfaces for attachment of future mission payloads; and, remain operable and mission capable after being impacted by multiple 7.62mm small arms rounds at zero standoff distance.

or aiming and targeting weapons with open-loop controls on UGVs that have poorly placed sensors or suffer from delays and latencies). Indiscriminate application of automation without regard to the resulting roles of the operators (and/or modelling the system or the system impact appropriately) is likely to lead to complaints from the operators that have a negative impact on capability over time.

As a result, it is essential to include stakeholder feedback during any evaluation process. However, care must be taken to maintain a certain intellectual autonomy for the analysts so that they not to confuse the inclusion of stakeholder feedback with ‘populism’. For instance, while stakeholders must be invited to engage in the evaluation process and their points of view canvassed, it is not mandatory for these to be taken into consideration when the conclusions and recommendations are formulated. This is because many stakeholders rely heavily on their past experience or attempt to provide as many metrics as possible to gain a comprehensive understanding of the system. Both of these can lead to an inadequate understanding of the system, expensive experimentation and analysis and the possibility of errors.

Moreover, as [78] puts it, “Unlike medicine, evaluation is not a discipline that has been developed by practicing professionals over thousands of years, so we are not yet at the stage where we have huge encyclopaedias that will walk us through any evaluation step-by-step. As a result, while only a few people believe they can do research, many more think they can do evaluation. ... Good analysts have been trained in the practices of research or practical evaluation and how to capture user feedback, whereas a rudimentary training in evaluation methodology frequently provides a prejudiced perspective and does not qualify lay persons (however operationally knowledgeable) to the task of evaluation.”

This does not imply that militaries cannot evaluate the benefits of technologies and/or have not recognised the potential of UVS in capability terms. At present, however, most studies indicate that the warfighters expect many UVS operations to be performed in the same way as they are conducted by existing (i.e. manned) forces, which indicates the transformational nature of these systems is potentially being ignored. Overall, it is the degree of sensitivity to the full range of parametric variations that are usually overlooked in systems analysis. There are also a number of other deficiencies with many existing UVS assessment techniques:

- Analytical tools do not target UVS force integration, social or cognitive domains
- There is a lack of analysis pertaining to the key features of UVS force integration
- UVS development is largely stove-piped according to its application domain
- Data and information on military systems is not usually shared between nations

- Analysis timelines are often too slow to meet the pace of technology development
- The correct measures of merit are often not properly developed up front
- Analysis frequently does not include unintended consequences, risks and benefits
- Fielded UVS are not designed to capture data relevant to off-line analysis
- Data observed in the field is not usually of a standardised (or sharable) format
- Due to classification issues data observed in the field is not usually releasable
- Analysis focuses on military outcomes rather than civilian support activities

We therefore require a strategy that recognises that the nature and impact of UVS on current and future military operations are in general undefined and – in addition to any technology demonstrations, field trials or technology force insertion programs – attempts to develop quantifiable impact statements upon which informed decisions can be based. To this end, the analysis should have two main aims:

- To provide the decision-makers with a quantitative basis for making decisions regarding operations and acquisitions under their control, and
- To highlight those qualitative aspects of operations and acquisition that requires further consideration or the decision-makers to exercise their own judgement.

These may be further broken down into two communities: an executive level, which is predominantly interested in a means by which they might define and articulate capability requirements; and a user community, who are mainly interested in the systems at a more technical level, but who also need tools that simplify interactions between the users, acquisition office, designers/developers, industry, and test and evaluation teams. In other words, the latter require a framework that allows system-specific, specification-level detail to be articulated for the testing and V&V for UVS.

The approach described here (and depicted graphically in Figure 5.1) is one adapted from [32] and [100]. It is essentially a layered, systems analysis that recognises that each UVS component or system forms part of a ‘systems of systems’ and is designed to function both as a single entity (i.e. as a platform) and as a component of a complex system or network of humans and other UVS. The tools (see Simulating UVS) include manual and seminar war games, technology development and demonstration, technology force insertion into free-flow exercises, human-in-the-loop and interactive simulation, and operations analysis, each allowing abstraction at different levels of realism and hence resolution.

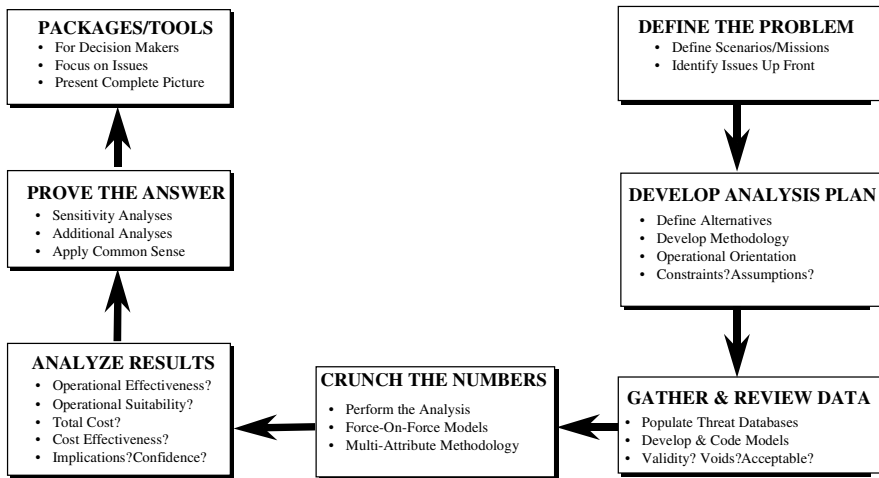


Fig. 5.1 Systems Engineering Approach to UVS-Force Integration [32]

As Figure 5.1 indicates, the first step in any analytical process is to identify exactly which problems need to be addressed, without stating how the solution is to be obtained. Our challenge is relatively straightforward to articulate, “Do UVS enhance the capability of military forces” and, if so “How and by how much?” Using the vernacular this might be phrased, “Can we use UVS to do things better; and, can we do better things with UVS?” More formally:

- Do UVS enable warfighters to do things that they currently cannot?
- Do UVS indirectly expose soldiers to more danger than it saves them from?
- Do UVS enable a military force to carry out its operations more cheaply, more quickly, or more productively than current or alternative technologies?

To achieve this, a multi-layered approach is necessary with the benefits of UVS evaluated through their impact on the fulfilment of the military objectives within scenarios and measured in terms of defined qualities that are relevant to these military objectives. This means that the operational value of UVS must be assessed in conjunction with its functional performance and the human-UVS team performance. This is because the UVS may accomplish its task perfectly, but negatively impact the system performance or military outcome in some way. Consequently, we will attempt to evaluate the following performance relationships:

- The force effectiveness of the introduced capability,
- The effectiveness of the human-UVS relationships, and
- The performance of the UVS relative to their identified tasks.

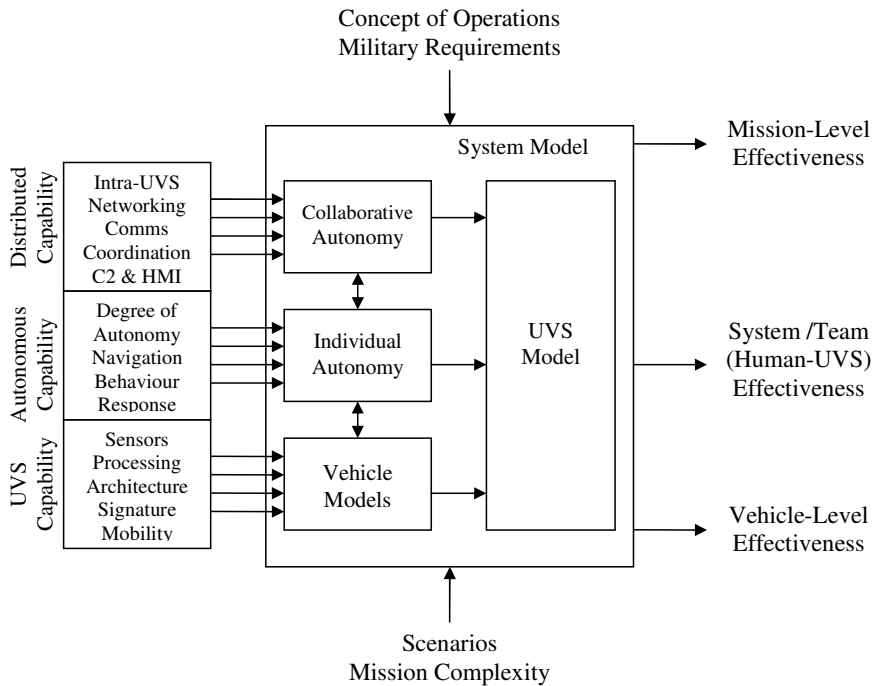


Fig. 5.2 System model for determining UVS effectiveness

The system model that allows these relationships to be addressed within the framework outlined in this book (networked versus individual autonomy versus mission complexity) is depicted graphically in Figure 5.2.⁷³ An additional benefit of using such a model – and which is also useful for the purposes of technology feedback, stimulation and guidance – is that we may determine the most appropriate level of UVS automation relative to the technologies available, requirements and concepts of operation. The model permits:

- Different types of levels of simulation and experimentation [3]:
 - Discovery – in which the potential military benefits of UVS are explored

⁷³ While it is recognised that the model shown in Figure 12 implies that the two axes of individual and collaborative autonomy are separated, in reality – and in particular for fielded systems or immersive and hardware-in-the-loop simulations – the architectures of the UVS may not lend themselves to explicit separation. For example, the data fusion and feature tracking algorithms may address several aspects of individual and collaborative autonomy (e.g. communications and sensor scheduling, target tracking and trajectory management) as a coupled problem. As a result, while the degrees of autonomy may be defined as separated, their instantiation in algorithmic form may in fact be practically indivisible.

- Hypothesis Testing – in which specific hypotheses are supported or refuted
- Demonstration – in which new technologies or concepts of use are shown to have value in a specific set of circumstances.
- Various degrees of human interaction with individual and multiple vehicles and the impact of specific human-UVS automation schema to be assessed;
- Various levels of collaboration between the UVS to be assessed relative to the indirect influences of interaction between UVS or their automation schemes;
- Both of the above to be assessed relative to the various degrees of scenario or mission complexity and scale (i.e. the number of mission-level activities that can be undertaken by the system regardless of whether they are undertaken by UVS, humans, or some combination thereof);
- UVS reliability to be assessed against all of the above through the introduction of specific modules that represent such aspects (e.g. comms outages, architectural or sensor failures, etc).

Other anticipated outcomes of such an approach include:

- The establishment of standards and expectation of performance;
- The establishment of the bounds of performance of the systems;
- The establishment of the effects of constraints imposed on the systems;
- Comparison of alternative systems designed to achieve a similar purpose;
- Assessment of a system's use in novel application domains or missions;
- Identification of potential weaknesses in force mix options or systems;
- Analysis on the impacts of organisational, force structure or system changes;
- Determination of the most cost-effective approaches to achieve objectives;
- Comparison of replacement systems against any predecessors or competitors;
- The capacity to trend and track on-going technological improvements;
- An ability to analyse returns on capability investment decisions relating to UVS.

5.3 Simulating UVS

Simulation is a staple technique for military or capability analysis and simulating UVS, just like simulating any other technology, can be categorised in a variety of different ways: technique, level of war, level of resolution, timeframe, geographic scale, and so on. However, little appears to have been done to integrate the simulation tools into the systems engineering process for assessing the impact of UVS technologies on systems performance, life-cycle costs, etc. Here the broadest interpretation is placed on the word 'simulation', with each technique offering

different degrees of reproducibility and abstraction and each loosely proportional to its operational realism and resource requirements (Figure 5.3). The techniques include [80] [225]:

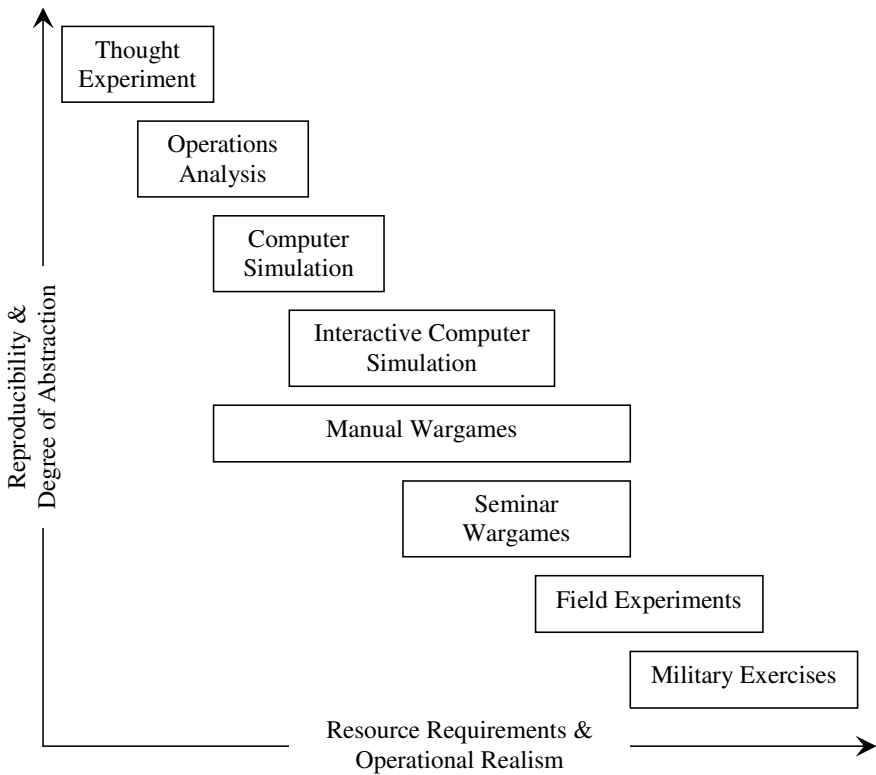


Fig. 5.3 Overview of Simulation Methods (adapted from [225])

Thought Experiments are usually used by subject matter experts (SME) to rapidly scope or assess specific problems or test a hypothesis or theory. The goal is often to test the potential consequences of the principle under consideration. Depending on the structure of individual scenarios it may or may not be possible to actually perform the experiment.

Operations Analysis (OA) or Operations Research (OR) is the term used to describe the process of applying science, in particular mathematical and statistical models, to solve military operational problems. Modern OR frequently relies upon models of complex adaptive systems to represent aspects of military conflict and involves a great deal of computer processing to carry out the ‘number crunching’. Strictly speaking OA and OR do not involve human decision-making. However, many of the ‘rules’ are derived from analytical devices, computer models of humans, or other mathematical methods.

Computer Simulation is a dynamic model created by adding a temporal dimension to it.⁷⁴ Whereas a model is often static, a simulation will change over time. For instance, it might show how a UGV performs in a combat situation, sustains damage, or travels across terrain. Many simulations are **Interactive**, meaning that they allow users to adjust the simulation in real time and influence the outcomes.

Manual Wargaming is usually a form of board game that attempts to model combat. The board generally represents a scaled map of some region and is divided into spatial units in order to allow movement. For tactical games, these spatial segments are often colour coded to represent various terrain types, which then further influences the rules of combat and movement for any units therein (e.g. a green hex segment may represent woodland, which may help to conceal UGVs, but may also impede their movement). Military units are usually represented by pieces of cardboard, plastic or metal and ‘play’ is governed through the use of a set of rules, cards, dice, etc.

Seminar Wargaming is an institutional exercise in which teams of players are grouped together in a room and presented with the role of a participant in a military-political scenario. The scenario, which frequently takes place in the future, is usually described by a team of analysts and designers. After discussing and assessing the situation each team decides on a number of actions, which are transmitted to the game’s controllers. The controllers then return to the teams with a situation that has resulted from the analysis of all of the moves taken by each team. Those who participate are often obliged to do so, regardless of their assessment of the game’s value. As a result, of (a) military organisations being hierarchal and professional, and (b) seminar wargames being designed by large teams, institutional ‘groupthink’ and compliance with doctrinal norms and hierarchal mandates tend to suppress novelty and creative thinking. As a result strong criticism and radical revision of extant thinking is not common. It is often difficult to ensure that participants ‘think like the enemy’.

Many defence forces engage in **Field Experimentation** and **Military Exercises**, which are essentially physical simulations separated mainly by degree. These simulations are usually extremely costly as they involve technology development programs and/or large numbers of troops. Often they do not lend themselves to the development or perfection of theoretical improvement as they either down-play the importance of individual soldiers or technologies in combat or because effort is focused on developing or perfecting well-known skills.

While realising the full benefits of UVS simulation means that its use must begin at the concept design phase, from an evaluative stand-point the value of the technique depends upon the goal of the simulation. There are generally five accepted goals [80]:

- **Model Validation:** At the very least any simulation provides an understanding of the model upon which it is based. Furthermore, to the

⁷⁴ A model is a representation of a real entity that represents the object or system at some level of detail, beyond mere reference. It aims to reproduce the various components and features of the physical world and/or the thing being represented.

extent that the model represents reality, corresponding insights may also be achieved. This is often useful for obtaining a deeper knowledge of the UVS model so as to be able to revise it and/or assess the assumptions upon which it is based.

- **Prediction:** Although the dangers of using simulation for the purposes of prediction are obvious – and most modellers warn against using their work in this way – the reality is that a great deal of simulation can be (and often is) used for this purpose. However, since many simulations are interactive with the outcomes dependent upon the actions of an opponent, great caution must be exercised when employing the outcomes of simulation in this fashion. In particular, strategic and capability plans cannot be determinate as the opponent's plans always constitute an unknown. Nevertheless, we are still able to use simulation to predict events in ways that are not possible from other sorts of endeavour. For instance, by mirroring the key considerations that drive the flow of a scenario we may be able to account for elements where other 'cleaner' models cannot. In other words, even though the analysis may itself be of dubious predictive value, it may reveal key tenets of technology, tactics, or periods of time that play a decisive role in a given scenario [225]. Similarly, hardware-in-the-loop simulation may be used in a predictive fashion to test (and hence predict) the performance of particular software upgrades, architectures, functional systems, etc.
- **Education & Innovation:** This is often the key goal of a simulation, as it is perhaps the only theatre outside actual warfare where militaries and individuals can experiment with new tactics, concepts and ideas. For instance, military novices and technologists are allowed to explore novel and unorthodox strategies without running the risk of loss of life.
- **Simulation Immersion:** The value of immersive simulation for UVS is related to the goals of the simulation. If the aim is to train users and commanders in the use of UVS then those who ultimately operate the equipment in combat conditions will have a feel for those experiences. Furthermore, if the aim is to test the human-machine interface or some other functional sub-component of the UVS then a great deal can often be gleaned through the use of human and hardware-in-the-loop simulation. However, as simulations fail to offer anything approaching reality, care must be taken to avoid making an obsession out of some elements while completely ignoring others.
- **Capability Investment and Policy Formation:** Ultimately, the goal of the analysis should be to foster consideration about the priority of investment strategies and policy decisions regarding UVS. Unfortunately, many simulation tools are not specifically designed for this activity, either because their goals are pre-ordained and not subject to revision or because their level of abstraction is not well-matched to the task.

As a result, the simulation needs to focus on the relevant functional components of UVS, the dynamics of battlefield engagements, the interplay of the relevant human decisions and the outcomes of those decisions depending upon the manner

in which they relate to the performance of the UVS relative to their identified tasks, the effectiveness of the human-UVS relationships or the force effectiveness of the UVS capability. Figure 5.4 shows the abstractive value of each technique as it relates to these levels of performance.

Category	Effectiveness in Evaluating			
	Effectiveness: Whole of Force	Effectiveness: UVS Section	UVS-Human Interaction	UVS Performance
Operations Analysis	✓	✓	L	L
Computer Simulation	✘	L	L	✓
Interactive Simulation	✘	L	✓	✓
Manual Wargames	✓	✓	✘	✘
Seminar Wargames	✓	✓	✘	✘
Field Experimentation	✓	✓	✓✓	✓✓
Military Exercises	✓✓	✓✓	✓✓	✓✓

(✓✓ ~ well-suited, ✓ ~ reasonably effective, L ~ Limited, ✘ ~ poor)

Fig. 5.4 Value of Simulation Techniques

The failings with respect to simulation are well known. For completeness, however, a number of standard pitfalls to simulating autonomous UVS are described.

The absence of reality from simulation means that the dangers of using such approaches are self-evident: simulated UGV's do not get stuck in mud, the weather does not change, and weapons do not misfire. However there is a trade-off to be struck between the accuracy or realism of a simulation and its usefulness. If the aim of the simulation is as a predictive tool or as a strategic training instrument, basing a simulation on a faulty series of models is like 'building on sand' [80] and any achievement made through simulation is compromised by the flawed assumptions on which the models are based. Indeed, almost all of the goals are harmed by faulty models.

Achieving the requisite level of accuracy for a model, however, is non-trivial. Moreover accuracy is characterised by several dimensions: if a UGV carries a weapon that can fire 800 m in reality it should be able to fire the same distance in a simulation; if a UUV can travel 30 miles in a day it should be able to do the

same in the simulation. Nevertheless, quality models do not have to be replicas of reality as they may aggregate a range of complex phenomena into simpler systems. As a result, more useful models of UVS are often less detailed. For instance, it is not necessary to model each and every physical or functional component (i.e. all elements depicted in Figure 5.2) as this would require significant resources to be applied to elements unrelated to the goals of the analysis. Instead the outcome of several interactive processes may be readily represented in terms of UVS mobility, payload action, etc in response to inputs such as sensor range, feature abstraction, a level of autonomy, and so on.

The decision of where and how to aggregate complex phenomena into simpler systems should reflect the aims of the simulation, which in our case will vary in accordance with the level of analysis conducted, the performance of the UVS relative to their identified tasks, the effectiveness of the human-UVS relationships, or the intended force effectiveness of the introduced UVS capability. So, for example, higher level simulations may require us to represent phenomena that correspond to less obvious elements of using UVS on the battlefield such as logistical footprints or maintenance processes.

The value of much of our simulation therefore lies in its ability to highlight the key human factors involved as by its nature it must attempt to explore the messy and often unquantifiable questions that technology development and force-insertion programs ignore. The aim of using these tools is then to teach us ‘what we did not know that we did not know’ [223]. Indeed, the potential value of such tools is predominantly as education and training aids; devices to help explore and new operational concepts of use and an aid for explaining the value of these new concepts and ideas to executive decision-makers. Moreover, as an exploratory tool, the above techniques provide technologists, analysts, users and other stakeholders with insights that can lead to further investigation into the validity of their views.

In particular, as [258] points out, if a tool incorporates the critical factors into its models and procedures the analysis can lead to the discovery of factors or other issues of concern, which have previously been unexpected or undervalued. Thereafter, by allowing human decisions to influence events made under the pressure of time and on the basis of imperfect or incomplete information – and by incorporating a degree of randomness or luck – the UVS simulations may come closer to illuminating the dynamics of the modern battlefield, which in turn allows other factors to be quantified in more concrete terms.

The power of these analytical tools in communicating a message, however, is also a potential danger [80]. The tools – particularly immersive and computer-based simulation techniques – create the illusion of reality, which can be a powerful and insidious influence on those without experience in evaluation or operations analysis. To this end, a poorly designed methodology or simulation can reinforce or provide an ineffective or even false picture of the value of UVS.

It is also important to understand that “the simulation is not the analytical process” [80]. It does not produce rigorous, quantitative or logical dissection of the problem or define the MOE and other metrics that allow comparison of the alternatives and options. Nor, despite the beguiling nature of many computer

simulations, is it real. When applied correctly the utility of combining distributed ‘gee-whiz’ technological capabilities and virtual reality techniques are extraordinarily valuable, but they are not a substitute for experience with actual systems in real environments. To this end we must use a combination of the techniques outlined in Figure 5.3 to understand the value of UVS. Additionally, the more and more realistically we attempt to reproduce the details of real environments the greater the amount and level of detailed data with which we must populate the models. Furthermore, data for creating and populating models, particularly that relating to one’s adversaries, is usually difficult to acquire as the sources, classification, accessibility and release-ability are often highly restricted.

A final and significant vulnerability of simulations is that they sometimes prove what they already assume. Care must be taken, therefore, not to accept or reject a model’s outcomes purely because they agree with intuition.

5.3.1 Interactive Simulation

Simulation for the purposes of research, such as that focused on adaptive automation, scalable human-robot interfaces, multi-modal and intelligent feedback and control, weaponisation, advanced perception, the instantiation of tactical behaviours and a great deal of human factors analysis, is both possible and effective. At an algorithmic level, there are many obvious similarities between achieving success in a synthetic environment and achieving it in the real world; except that considerably more signal or sensor conditioning is required in the real world, mobility tends not be much of an issue in simulation and the environment is fully definable and under-stood in the synthetic world. Moreover, given the need to model the many and often somewhat unpredictable tactical behaviours of humans under battlefield conditions, in some ways it is easier to simulate the conditioned (albeit polymorphic) responses of a UVS in similar environments.

This sort of simulation allows experimentation and evaluation of novel technologies, concepts of use and even force-mix options prior to acquisition. Additionally, it allows experimentation that may be physically, ethically or environmentally prevented. Fundamentally, however, it allows repeatable scientific exploration at lower cost than can typically be achieved using the often prohibitively expensive technology demonstrator or force insertion programs alone. Even so, such simulation techniques should not be considered low cost or simple development options.

Many of the more effective interactive UVS simulations need to rely upon more realistic immersive experiences which incorporate information limitations inherent in real combat roles. The success of any such an endeavour will then depend upon the involvement or representation of humans within the decision-making loop, as they will be needed to issue mission-level directions, manage uncertainty, and inject a degree of flexibility and creativity into the system. In addition to the simulation of any sensors, platforms and autonomous functions, etc., therefore, the sensed information must also be combined, manipulated and passed to a ‘rear echelon’ where it usually further integrated with user applications such as track data and imagery, geospatial information, visualisation

and document management tools. This fused, value-added product must then be disseminated to human users (real or simulated) in near real time to allow them to monitor or redirect the UVS, as appropriate. Information provided by the UVS must therefore be presented in a form that can be rapidly understood by humans; and information provided by humans to the (simulated) UVS must be in a form that is readily interpretable by them.

At the force level the representation of human behaviour is often difficult to achieve and is more readily represented at a lower level. For example, the ability to represent a battle unit's perceptions built, updated and validated from information available to the unit from its UVS requires each unit to have its own perceptions, gaining knowledge from superior, subordinate, or adjacent units only when appropriate. Alternatively, simulating a commander's decision based on his unit's perception of the battlefield requires each unit to act based on what it perceives the situation to be, not based on ground truth available within the model; in other words, when an entity takes action based on inaccurate perceptions, it suffers the appropriate consequences. The second of these is easier to represent than the first [100], but neither are particularly trivial undertakings.

Modelling any force level process as a group of agents favours the capture of the decision making nature of tasks, which can be implemented, in an object-oriented environment, as either objects (e.g., actor or applet type of agents) or aggregates of objects (coarse-grain agents) [160]. However, a structured hierarchy of models will likely need to be used to create an audit trail from UVS systems and processes through to battle outcome. The aim should be to create supporting performance level models of particular aspects of the process, which can be examined at the performance level. For example:

- A detailed model of an autonomous UVS (or network thereof) can be very complex if we attempt to take into account the flow of a priori intelligence requirements, tasking, collection processes, fusion processes, and supervisory interaction. In order to analyse the impact of information flow in a realistic manner, it is important to have all of this detail, but it does not necessarily have to be represented explicitly in the main model. It may be possible to use a supporting model, which captures all of this detail, to produce outputs such as speed and quality of information. These can then form inputs to the main model, which will in turn take them into account in producing its own outputs at the higher level;
- Similarly for the transmission of information across the battlefield. The question arises as to whether all the communications transmission media (radio, satellites, data links etc.), with their capacities, security level, data rates, error rates etc., should be represented in the main model explicitly, or whether this aspect should be split out as a supporting model of the overall process also. Again, supporting models may be able to provide sets of input data for the main model. The main model would generate demands on the communications systems. The supporting model could check if these demands could be satisfied. If not, communication delays in the main model could be increased, and the main model re-run. This usually has to be done a number of times to bring the main and

supporting models into balance [238]. Nevertheless, such an approach can generate valuable insights, although the high rate of services that may be required to support the main model can involve a long analytical process, which can become critical with a large assortment of UVS-related parameters, a long scenario period, and/or a number of sensitivity analyses to perform.

- Technology demonstrator programs will need to simulate UVS insertion for a range of platforms and sensors using both human and hardware-in-the-loop simulation. The infrastructure to accommodate this (which can also be used as a software test bed for the demonstrator programs) is essential. Information gained from such HWIL models will provide inputs into the higher level force-on-force model, ensuring that it will not become overly complex.

To attempt to model sensors, the cognitive processes of the supervisors, the UVS and/or their autonomy at a component or functional level with any degree of fidelity requires a very significant level of staff and computing resources/commitment and can defeat the aim of redressing the technology tempo and improved value for money. To this end, it is the functionality of the sensors and/or the effects of IDT or HRI that need to be modelled. In other words, we should not seek to model the individual functions of the sensors, sensor processing, perception, navigation, location, communications, architectures, data fusion and target location algorithms, human cognitive process, etc. where the output of each acts as the input to another. Instead we should seek to model the total capabilities of the system as parametric profiles – i.e. we might model the detection capabilities of the UVS (e.g. the sensitivity, target location error, and target velocity error of its sensors) as a function of (say) range, time and clutter – and its autonomy as a high-level function that allows robust navigation, corrupted by a time dependent, stochastic failure mode. The integrity of these parametric profiles can then be validated offline against data obtained from more detailed models of the UVS sensors, HWIL simulation, technology-force insertion trials, demonstrator programs, etc [98].

There is obviously an overhead to pay in implementing effects based models of this type when one takes into consideration the need to model the interaction between an individual sensor and different types of environment. For instance, concerns will be raised about the validity of any analysis that is based upon such models and techniques. We must remember, however, that in many instances we will be attempting to validate concepts for events or technological developments that have not yet happened or that may only occur in ten years time. Consequently, in the first instance, it is more important to apply common sense to the models than to expend significant resources on the validity and basic integrity of every component of the simulation. This does not diminish the importance of the validation, verification and accreditation processes, but highlights the need to carry out an appropriate sensitivity analysis.

To a large degree the techniques and the level of performance they attempt to replicate are separated by the way in which they deal with time. All, however, must proceed in stages, or ‘moves’ in the case of the higher level techniques. At

each step, therefore, the participants (i.e. the real/simulated humans and/or UVS) must assess the situation as they perceive it and framed by the scenario. This information must be suitably manipulated and processed and provided back to the simulation kernel/control for the next stage to be processed. Given that one of the potential benefits of employing UVS, particularly multi-UVS networked to other assets, is that they produce responses and decision compression at levels unachievable by manned assets it may be necessary to accommodate mixed time scales to various phases of some types of simulation. Real time and ‘moving clock’ simulations would be appropriate, but such techniques are generally intensive and very costly to resource.

Currently, the propagation over time of a simulation is typically based on the notion that plans and orders originate from senior commanders and more junior echelons refine the orders from above into more focused orders for their subordinates’ execution [225]. This cascade of command and control is readily represented in the form of rules and conventions (and hence missions for UVS). In a force involving autonomous UVS, the decision-making may well be ‘flatter’ and dependent upon the IDT and characteristics of the UVS themselves. As a result, the need for intermediary command and control may become superfluous. Unfortunately, should the UVS be deemed to add value to capability, selection of the optimal command and control strategies will also require simulation.

As highlighted in the previous chapter, many of the principal requirements for achieving capability value-add through the use of UVS depend upon our capacity to derive and maintain shared situational awareness. To this end, we must do considerably more than simply simulate widened communications bandwidth and multiple linked platforms. We must also include the interpretation and application of information of uneven quality and timeliness, the purification of information to preserve the quality of the network and prevent error propagation, and the delivery and interpretation of information by people and UVS to allow coordination of assets to obtain synchronised response. It is this enhanced coordination and synchronisation of assets that has the potential for improving UVS-related operations. However, it also increases the matrix of options available to us.

Clearly distributed and agent-based simulation has a potential role to play. However, simulation of a suitable tactical network, the cost of such endeavours, and the need to develop shared situational awareness techniques means that we have a ‘chicken-and-egg’ dilemma regarding which should come first: the technology development or the simulation? Similarly, it may be necessary to alter the command and control doctrines to allow multiple autonomous UVS to be effectively simulated – and what the training needs of such an endeavour might be are, of course, an open question [240].

Interactive simulation, however, provides us with an opportunity for addressing many of these issues. Traditional simulation consists predominantly of a single layer in which all entities are immersed. Conceptually we may divide UVS into two layers: a manned layer and a UVS layer.⁷⁵ By separating out these two layers,

⁷⁵ As the level of simulation complexity increases to include (say) multiple distributed entities, we may also introduce other layers such as real command and control systems and structures.

we may then use real systems in our simulations and allow the humans and many of the UVS decision-making functions to interact just as they would in real environments (or interact with other simulated entities as they would real ones). However, such simulations are only truly meaningful if the underlying models are sufficiently accurate and (for example) the spectral representations of natural or man-made objects, the effects of weather, etc. appropriately detailed. This will become even more of an issue as smart materials allow novel sensors and actuators to be embedded so that the material serves multiple purposes simultaneously (e.g. structural rigidity and sensing).

Finally, the traditional view of simulations is that they are played out on maps, boards, or computers and are conducted through the movement of assets and systems around geographic areas, with their movement and interaction governed by rules. For UVS simulation, perhaps the key ‘board’ may not be the traditional geographic display, but the shared human-UVS situational awareness picture because it is the information flows between the networked (simulated) entities that enable the commanders and UVS to become appropriately oriented. Just as would be the case in determining the requirements for Causal Analysis (see Legal Issues) for UVS-related accident investigation, deriving and demonstrating optimality for UVS in capability terms may then require the simulation tool to provide a ‘who knew what when’ list as an IDT onboard a UVS interacts with the information flowing around the battlefield. Furthermore, to a very large extent, it is these bits of information that have become the new ‘players’ in a simulation.

5.4 Measuring the Effectiveness of UVS

As indicated in the previous sections, there are several ‘customers’ for the outcomes of any performance evaluation or analysis: developers, sponsors of the development work (who are often executive decision-makers and/or capability planners) and users. This means we must provide each with quantitative feedback that they can use to make judgements either relative to areas of work that need more attention or to allow what would otherwise be an ‘apples and oranges’ comparison. For example, the analysis might need to be focused on assisting the executive decision-makers to determine how much military performance can be traded for the sake of greater force protection and less loss-of-life; not something that can easily be communicated through the language of technology. By contrast, messages conveyed to the users and commanders should allow them to understand how to effectively support their operations and integrate the novel systems into their planning process. These clients must understand when such systems best support their operations in a variety of environments, cultures and missions, plan for their support or replacement in the event that they are damaged or break down, and – perhaps most importantly – understand when not to acquire or use autonomous UVS.

It is important to maximise the understanding and the outcomes gained from any analysis. However, due to limited resources it may not be possible to collect data against all of the required outcomes. As a result, the analysis should be interpreted and evaluated based on how much it assists customers understand key

phenomena of interest and matters over which they have developmental, operational or investment control. It is also important to try to interpret the metrics in terms of how they help explain the underlying technical reasons for what other metrics measure.

Additionally, the outcomes of an analysis might well have an impact at several levels. For example, a cost-benefit analysis might suggest a modernisation program should be conducted to anticipate efficiency gained through the introduction of UVS. Such an analysis should include not only the forward operating contingent that makes use of the autonomous vehicles, but also the training, sustainment organisations (see Training) and other organisational structures, allowing the conclusion that (for instance) force elements that currently have organic maintenance capabilities or small logistic footprints may need to operate and be structured differently when employing autonomous UVS.

In the final analysis, however, it is imperative that the results of any analysis be presented in a form that is intelligible to its customers. In other words,

- The analysis must focus on issues that are relevant to customers;
- It must highlight aspects of the analysis that require more detailed consideration;
- The outcomes must be conveyed in language that the decision-makers understand;
- It must provide a complete, yet concise, picture of the analysis actually conducted

5.4.1 Measuring the Force Effectiveness of UVS

For any analysis, we need to define specific missions and scenarios of relevance to likely users of UVS (the missions will influence the boundaries of the scenarios and the level of detail needed therein). For example, to analyse the effect on land capability resulting from changes in size, mobility or endurance of a UGV relatively detailed, task-specific scenarios will be required that cover the full range of relevant system applications. It is expected that in most cases, scenarios will require a high level of detail to adequately address such issues. The missions and scenarios need to be developed by the agency conducting the analysis in close conjunction with the technologists developing the UVS and the users likely to be responsible for their integration.

It is important that only scenarios endorsed by senior user communities be used. This is pivotal to the success of the overall objective and considerable time and effort needs to be spent ensuring that the correct questions are being asked and addressed. Moreover, in order for the analysis to remain focused towards specific user requirements, the scenarios may need to be developed such that they are relevant to one or more major procurement programs of the host nation. In essence the role of a scenario is to define a set of conditions and restrictions to enable ‘good’ analysis of the issues, as well as to create a structure within which the results of the analysis can be understood and interpreted.

Due to resource constraints not all concepts, missions, scenarios or areas of interest can be examined thoroughly in the first instance. Consequently, a range of

scenarios and tasks will probably need to be investigated in a preliminary fashion so that the results of these analyses can be used to identify concepts, technologies and areas of operations that require more detailed analysis. This approach carries with it inherent risk, but mitigates the much greater risk of expending valuable resources on unproductive avenues of study.

Many performance tests, competitions, demonstrations and analyses have been devised to measure the intelligence or overall effectiveness of UVS, ranging from robot soccer (football) to measuring the performance of a UVS in computer simulations. As a result, many of the systems designs are narrowly focused and optimised against the specific tasks under consideration. As military UVS will be used in contexts for which they were not specifically designed or tested, this raises several questions [211]:

- Should performance tests be domain specific
- Do we need calibrated facilities to make quantitative observations
- Would standardised test data with geo-registered ground truth be useful
- What kind of data should be collected and who should perform the tests⁷⁶
- Can sub-systems and components be meaningfully tested in isolation
- Will the adaptive nature of future UVS negate any meaningful testing

Also, due to an historical absence of universally agreed and quantifiable metrics for UVS performance, most of the research results associated with artificial intelligence, robotics, UVS, etc have been in the form of specific missions and demonstrations rather than experiments with data that is quantitative or fiducially referenced [24]. Moreover, as complex concepts often require multiple measures to provide valid information, no single measure or methodology exists that satisfactorily assesses the overall force effectiveness of UVS technology. As a result, to link the performance of a system as a whole to the performance of its components, any metrics must correspond to critical tasks.

In regard to force effectiveness, specific objectives, issues, and operational outcomes must be identified so that the relevant metrics can be established and then quantified. A layered and domain-specific approach is required for developing and defining metrics to measure the performance of UVS at the force level. Loosely this translates into Measures of Force Effectiveness (MOFE), Measures of Effectiveness (MOE), Measures of Performance (MOP), and Dimensional Parameters (DP). We also note, however, that the development of such measures requires a good understanding of fundamental parameters and their relevance to systems performance, which is usually the domain of technologists and operational users rather than experienced analysts [100].

Relationships between the different types of measures are often difficult to establish authoritatively. Linkages between lower-level measures (e.g. DP and MOP) are easier to determine than those between higher-level measures (e.g. MOE and MOFE) as MOE and MOFE are more dependent on the operational context and are hence more scenario dependent [203]. This again stresses the necessity for a

⁷⁶ User, acquisition organisation, developer/manufacturer, or an independent agency.

selected range of scenarios for the proper analysis and interpretation of the measures. MOFE tend to be few in number and difficult to obtain without the assistance of Subject Matter Experts (SME). However, they should, whenever practical, be used in context with MOP and MOE that also provide diagnostic information about the dynamics of the process. Figure 5.5 depicts characteristics of such measures. In selecting measures, they should also attempt to meet validity and reliability criteria (e.g. Figure 5.6 and Figure 5.7) [100]:

Measures	Focus	Scenario	Effort Required	Number	SME
MOFE	Force Outcome	Dependent	High	Few	Military
MOE	System Effectiveness (Equipment & Platform)	↕	↕	↕	↕
MOP	Asset (Equipment or Platform)				
DP	Process	Independent	Low	Many	Technical

Fig. 5.5 Tendencies of characteristics of measures (adapted from [203])

Validity Criteria	Definition
Is it Well-Defined?	Includes data collected for each metric
Is it Relevant?	Relates to UVS, mission, etc
Is it Realistic?	Relates realistically to the UVS and any associated uncertainties
Is it Appropriate?	Relates to acceptable standards and analysis objectives
Is it Inclusive?	Reflects those standards required by the analysis objectives
Is it Simple?	Easily understood by users

Fig. 5.6 Validity criteria of measures

Reliability Criteria	Definition
Is it Discriminatory?	Identifies differences between alternatives
Is it Flexible?	Helps to identify new requirements
Is it Measurable?	Able to be computed or estimated
Is it Quantitative?	Can be assigned numbers or ranks
Is it Objective?	Defined or derived, independent of subjective opinion
Is it Auditable?	There is a clear cause and effect trail
Is it Sensitive?	Reflects changes in system variables

Fig. 5.7 Reliability criteria of measures

The fundamental parameters are commonly referred to as Dimensional Parameters (DP) and are the properties or characteristics of the functional components of a UVS. DP's are the 'inputs' to a simulation or model such as probability of detection, probability of false alarm, time to process information, and so on. Depending upon the level of model, however, DP's in one model may be MOP's in another.

MOP's are measures of the performance for UVS, (e.g. communications range, mobility, endurance, etc.) and are often linked to DP's through an algebraic expression. MOP's indicate the degree to which an asset performs a task or meets a requirement under particular conditions. At the force level they are not generally used in the analysis. However, this does not reduce the relevance of the MOP's because they are an important element in the validation chain. They must be considered to ensure a logical progression into higher-level measures dealing with effectiveness, outcome and capability. Moreover, they are also essential elements in the process of identifying aspects of the operations that merit more detailed consideration.

MOE's should be quantifiable and are tools that assist in discriminating among a number of operational alternatives to answer a critical operational issue. They show how the alternatives compare in meeting functional objectives and mission level needs. In essence, they are 'predictions' of how an operation will progress when executed, are often expressed as a probability or percentage, and are usually derived from a combination of MOP's.

MOE's reflect the degree to which the employment of UVS and the other assets meet their requirements; the extent to which a system is well-matched to its mission; they focus on the impact of UVS within the operational context, e.g. number of threat systems destroyed, etc. MOE are often related to a number of MOP's through a model such that a change in a MOP causes a change in the MOE. In many instances, several

different MOE's of the operation may be pertinent and the choice of 'optimum' operation will depend upon the choice of MOE.

MOFE refer to the outcome of a military action and focus on how a force performs its mission or the degree to which it meets its objectives (e.g. lethality, survivability, etc). They are usually a function of a number of MOE, where each MOE is given a weighting in its importance to the mission. This weighting is typically based on military judgement. As in previous phases, considerable interaction is needed between the technologists, the users and the analysts to ensure that the correct MOE's and MOP's are used. This interaction is also useful in ensuring that appropriate weights are allocated to the relevant MOE's and MOFE's when combining them to highlight qualitative aspects of the studies that merit deeper consideration.

In addition to using the MOP's and MOE's to assess the operational value of UVS they must also be used to assess UVS systems and human-UVS team performance, and to provide feedback for improving the human-UVS relationship as it relates to workload, training needs, and so on. In this regard the MOE's should measure the quality and execution of UVS-related warfighting tasks and should focus on the effectiveness of mission and team-behaviour. Similarly, MOP's should be observable and derivable measures of the UVS operator's skills, strategies or steps to accomplish specific tasks. We will return to this in the next chapters, which are dedicated to measuring the effectiveness of UVS and human-UVS performance.

Many questions may be answered subjectively or qualitatively very easily. Moreover, it is important to note that ascertaining even a single quantitative measure may be non-trivial as we must identify specific values and parameters that represent numerical values. As a result, while the discussion hitherto has only made reference to quantifiable and hence objective metrics, it is important to note that they can also be qualitative or subjective.

For instance complex MOP's such as a commander's confidence in the information that he is receiving or the operational efficiency of his decisions when he knows that the information is imperfect may be useful. However, it may be necessary to involve human factors specialists to develop appropriate assessment procedures for such information. For example, uncertainty associated with situation awareness cannot be measured directly but it can be quantified by measuring parameters that drive it. Clearly the performance of such a system may vary with operating conditions (propagation, terrain, countermeasures, visualisation, data manipulation, stress, operational tempo, etc) and these variables must also be accounted for.

In order to ensure that the results of any analyses are meaningful, it will be necessary to populate the scenarios and models used with representative data and parametric values. To ensure the resultant simulation and analysis is tractable, however, it will also be necessary to make acceptable assumptions about some UVS and their functionality. This means that considerable information on each of the modelled systems will eventually be required for any

overall analysis to have sufficient fidelity. For instance, threats, orders of battle (ORBATS), tactics, rules of engagement, courses of action, deployment strategies, concepts of operation (CONOPS), technology and force mix options for both the blue and red forces all need to be included. To this end, a database containing information relating to a range of potential threat systems will be required, which may prove time-consuming and complex. Constraints imposed through various limitations, any assumptions made, and the variables also need to be carefully documented.

Most military UVS will be required to perform over the full spectrum of conflict from peacetime to full scale military operations. As a result, multiple scenarios covering different missions, blue and red force capabilities and behaviours, and possibly other actors (such as neutral parties and coalition partners) will need to be considered to ensure that all issues are fully addressed, and also to prevent inappropriate detailed analysis of a single, perhaps anomalous, situation. Figure 5.8 provides a list of key factors that each scenario will need to address. The goal should be to identify those that must be included in the scenario. For example, if the issue involves coalitions, both blue and allied forces must be considered.

A scenario may be too large for reasonable analysis and may need to be decomposed into smaller vignettes, each dealing with a certain aspect of the overall scenario. This allows the analysis to focus on the important segments of the scenario space or UVS mission. Each vignette will still need to be executed within the framework of the overall scenario, so it is important to note that any findings are only valid within the limitations of the assumptions and constraints of the scenario.

It will also be important to understand and record which scenario assumptions, boundary conditions, etc are driving factors in the analysis as military decision-makers need to be made aware of the degrees of uncertainty in the scenario and the robustness of the overall conclusions. For instance, the requirements of the IDT, platforms, sensors, algorithms, weapons control, etc are dependent upon the scenarios in which the UVS are employed and analysed. Consequently, in addition to defining metrics by which we evaluate the system and its potential design, we also need to understand scenarios in terms of their detailed components (targets, environmental effects, vehicle motion, and so on).

In order to state a level of confidence in the interpretation of measures, the underlying assumptions must be clearly stated and uncertainties recognised. Uncertainties manifest themselves in several ways that affect measures [203]:

- Uncertainties in the scenarios (e.g. the relevance to the purpose of the evaluation, uncertainties in the military objectives, knowledge of red force CONOPS, intentions, capabilities, weapon performance, uncertainties in terrain data, etc);
- Uncertainties in the model, (e.g. parameters used); and
- Uncertainties in outcomes, (e.g. sensitivity to input variations, model fidelity, etc).

Category	Key Factors
External	<p>A description of the external factors including national security interests; the political, economical, cultural, and military situation and the acting assumptions; mission objectives, constraints, and limitations related to red force and coalition partners; warfare domains, level of violence etc. is required.</p> <ul style="list-style-type: none"> • Political/military/cultural situation • Mission & mission objectives • Mission constraints & limitations • Rules of engagement • Military scope & intensity
Capabilities of Blue, Red and Allied Forces and non-combatants	<p>Military capabilities and available resources need to be included.</p> <ul style="list-style-type: none"> • ORBAT (force organisation, force components) • Platforms, equipment, weapons • Tactics, CONOPS, etc • Logistics, resources
Environment	<p>A description of the mission environment (whether generic or specific) is required. Scenarios should include logical assumptions about the environment under analysis.</p> <ul style="list-style-type: none"> • Geography/Region/Terrain/Accessibility/Vegetation • Climate/Weather • Infrastructure

Fig. 5.8 Scenario Framework (adapted from [203])

Crunching the numbers is usually seen as the heart of any systems analysis. In fact, it is usually the most straightforward and least time-consuming component if the other elements have been carried out properly. Fundamentally, the analysis needs to evaluate the MOFE and MOE for the UVS, their component technologies, human-UVS teaming, force integration, and concepts of use. To this end, performance deficiencies and potential operational and technical requirements and improvements are identified by evaluating the Mission Level Effectiveness (MLE) of the technologies, concepts, and techniques against the MOFE, etc. To this end, the analysis must:

- Evaluate the MLE for a current (i.e. manned) capability
- Evaluate the MLE for a proposed basic UVS capability
- Evaluate the MLE for a proposed advanced UVS capability
- Evaluate the MLE for proposed advanced capability w/modified CONOPS

Stages (1) and (2) highlight deficiencies in current capability, whereas stages (3) and (4) identify potential improvements. In time, the collection of data from field experiments and technology-force insertion programs will enhance the analysis process. To this end, it is usually beneficial to involve personnel in war games, trials,

and demonstrations so they can build an appreciation of the issues of ownership and operational functions of UVS and develop objectives and novel concepts of use. Additionally, each scenario must be played a number of times to obtain a spread of data and the data placed into a database where it can be readily extracted into appropriate tables for further analysis. Statistical analysis will need to be used to analyse the data further. Other issues that have to be considered include operational suitability, cost effectiveness, confidence in the results, and so on.

Analyses of operations are largely scenario specific. Furthermore, in many instances, several different MOE's may be pertinent. Several methods then exist for aggregating a number of different MOE's into MOFE [32] [100] [105] [203], but in the first instance the analyst should rely upon their own common sense and the judgement of the senior military decision-makers – and seek feedback early. Figure 5.9 shows some typical MOE's and MOFE's, although the list is far from exhaustive. Furthermore, metrics identified in one category do not imply that they can be of use in another. Additionally, each metric or set of metrics has advantages, limitations and costs. The key is to efficiently select a set of metrics (among a large set) for a given context that maximises value and minimises cost in terms of the overarching goals of the analysis.

It is also necessary to ensure that the assessments made are robust to operational and parametric variations. In other words, the dependence of the outcomes to the most sensitive operational parameters also needs to be identified. In this regard, all complex operations depend upon a large number of parameters and the less sensitive ones must be eliminated in the early stages to ensure that best use is made of the limited resources available. Sometimes it is possible to statistically analyse the history of operations to pinpoint likely candidates for parametric sensitivities. However, given the novelty of autonomous UVS technology it may not be possible to statistically analyse the history of operations to pinpoint likely candidates for parametric sensitivities in this case.

Nevertheless, the most essential part of 'proving the answer' involves performing sensitivity analyses to ensure the validity and reliability of the results and thereby reduce uncertainty. By varying the assumptions and input data within the plausible ranges, excursions in the analysis provide insight into the effects of uncertainty. The goal is to establish the regions for which any results are valid and to isolate those factors that may be introducing uncertainty. Such analysis begins with a determination of the types of uncertainties involved, including those that may arise from the scenario or the models adopted. In particular, assumptions and limitations built into the scenario, the model and the data structures should be considered. The results of sensitivity analyses should be included in the final report to the clients.

Some additional analysis may also be required (e.g. statistical, to check that outputs are within acceptable bounds). Results can be double-checked via further runs of the models, with new seeds. Input variables can be sampled from defined distributions and these samples used to generate "sampled distributions". Sampled and defined (i.e. observed) distributions can be compared for each variable as an indicator of the representative nature of cases. There may also be a need to perform a qualitative analysis in conjunction with a Subject Matter Expert (SME).

Typical Measures of Force Effectiveness	Typical Measures of Effectiveness
Lethality	Total threat systems destroyed Total high priority threat systems destroyed Total (high priority) threat systems destroyed versus time Total threat systems destroyed by each category of friendly system Indirect and direct fire loss-exchange ratios (incl. over time)
Survivability	Total/percentage of friendly systems destroyed Total/percentage of high priority friendly systems destroyed Total (incl. high priority) friendly systems destroyed over time Total friendly systems destroyed by each category of threat system Fratricide ratio
Situational Awareness	Timeliness of information (time between the arrival of the information within the friendly force and the exhaustion of the response options) Latency of information (time between the arrival of the information within the friendly force and the dissemination of the information to the relevant platforms) Age of information (average time between information updates) Accuracy of information (% detections correctly/incorrectly ID) Relevance of information (% detections/messages acted upon) Completeness of information (% threat & friendly assets detected) Dissemination of information (% informed friendly assets) Surveillance volume (% time, spectrum, area under surveillance) Area of threat/friendly SA coverage versus engagement area Area of threat/friendly threat coverage versus engagement area Accuracy/movement of areas & target locations versus time
Tempo of Operations	Time required by friendly forces to accomplish mission Average time allowed for friendly/threat decisions to be made Number of response options available to friendly/threat forces Average time to transfer messages between/within C2 echelons Time required for a target to be acquired or sensor-shooter Operational availability of main combat systems

Fig. 5.9 Examples of MOFE and associated MOE's

In this situation military knowledge and common sense must be applied. Critical results can then be tested in a number of ways:

- By using simpler, more focused models,
- By checking for self-consistency, and
- By comparison with actual events if this information is available.

In addition to the need to highlight the operational parameters that impact most heavily upon the operational outcomes, it is necessary to assess the fidelity and integrity of the core assumptions made earlier. Considerable emphasis is often placed upon the validation process and we must remember, initially at least, that validating core assumptions against events that have not yet happened is a rather nebulous concept, particularly when the real events may only occur in the future. Consequently, in the first instance, it is more important to ensure that the core assumptions and results are 'filtered' using a combination of knowledge, experience, and common sense than to expend significant resources on validating the basic integrity of every component of the simulation. This is not to denigrate

the importance of verifying and validating the simulation's integrity, but simply highlights the importance of carrying out an appropriate sensitivity analysis.

In addition to the high-level validation process, where operational and trials outcomes are compared to MOE's, it should also be possible to conduct a system-level validation in which experience and outcomes obtained from operations, field trials and laboratory experiments are compared to MOP's. This falls into two broad categories:

- That obtained from the operational deployment of units, and
- That obtained from field experiments, etc. of more autonomous systems.

In this regard, it should be noted that it is not possible to get a true assessment of the suitability and usability of a system without testing it against representative users. Moreover, most UVS that have been immersed in operational engagements are tele-operated or semi-autonomous systems. Where the users are warfighters we may consider the tests representative. However, in field trials of the more autonomous systems, not only are most scenarios merely approximations of their operational counterparts, but the users are often the developers of the UVS, who have significantly more experience working with technology than the warfighters.⁷⁷ In these cases, the feedback may only be considered to have come from domain experts rather than representative users.

5.4.2 Measuring the Systems Performance of UVS

Key to the success of military UVS is their ability to work as partners with their human supervisors, leveraging the most useful capabilities of each. Basic mission tasks, regardless of the environment, will demand close collaboration. However, because cost pressures and the need to minimise operational risk will keep operator teams small, the effectiveness of the interactions will have a major impact on the effectiveness and performance of future UVS missions. To assess the systems performance of UVS we must therefore measure how well the humans and the UVS perform as a team where these interactions may be for mission planning, plan and task execution and monitoring, plan and problem diagnosis, and the authorisation of mission and plan execution.

There are a number of well-known techniques for evaluating human-robotic teams (HRT) and human-robot interactions (HRI) (e.g. [89] [121] [138] [211] [233] [259]). However, there is not yet a consensus on a standard framework, a major difficulty being the heterogeneity; the physical differences between the vehicle types and the diversity of military applications, which can range from HALE UAVs with relatively good communications to UUVs that may have low bandwidth communications or must work for extended periods without any communications. Furthermore, it is usually possible to carry out tasks using some combination of humans and UVS with the key issue being to achieve an optimal mix of automation for each task.

Some methods for assessing systems performance are based on decomposing scenarios into 'functional primitives' [236], allocating these primitives either to

⁷⁷ And considerably less experience in how such systems might be employed doctrinally.

the supervisors or the automation, evaluating execution of each primitive and computing the ratio of the performance benefit resource allocation. Other methods [216] measure interactive effort. That is, degree of autonomy and/or the overall effort required by the human to work as a component of the supervisor-UVS team. This technique is particularly effective when the overall task requires a mix of competencies within the team. However, the effort needed from the human component is a function of the complexity of the task, the circumstances under which it is performed, the training and skills of the human, the perception available through the 'window' of the HMI, the level of trust and confidence enjoyed, and the degree automation imbued within the UVS.

As a result, many of the techniques may provide a misleading insight as their metrics tend to focus more on the human work component or the UVS component [82]. Also, whereas UVS of duplicate design offer repeatable performance, many of the above degrees of freedom vary across individual humans. Additionally, many assessment techniques focus on single UVS problems, where the discrete levels of autonomy allow direct comparisons of the system's overall performance to be made against one another. When there are networks of UVS the problem becomes more complex as it is necessary to assess the coupled impact of the levels of human-UVS automation, the effects of various levels of collaboration between the UVS, the indirect influences of interaction between the automation schemes, and the impact of mission complexity.

Nevertheless, techniques do exist to identify the decision-making roles in which supervisors are most influential and effective relative to the capabilities of UVS [73]. For example, observations can be made under varying levels and types of human intervention and the speed and accuracy of decisions and actions, the time to respond to critical events, the duration of tasks/mission activities, and the ratio for completion of mission-critical vs. secondary objectives all used to estimate the operator-to-vehicle ratios for any given task/mission [211]. This can also be drawn out from the speed and accuracy of the task completion for different levels of task demands associated with the mission (e.g. the number and rate of the required tasks for successful mission completion; the complexity of the mission; etc). These objective measures can then be used to identify (say) the point at which the operators start to shed other tasks or fail to achieve accurate task completion. Ideally these metrics will form elements of the MOFE or MOE in higher level analyses.

Fortunately, the methodology outlined in the previous section describes not only a systematic approach for measuring the force effectiveness of UVS it also describes a methodology for developing metrics in general. We may therefore use it as a toolkit to determine metrics that stress objectivity, repeatability, fidelity, and real-world validity, thereby allowing the comparison of human-UVS relationships across different systems, levels and types of autonomy, tasks and missions. Using appropriate components of the defined missions and scenarios we may then frame the assessment of the interaction of UVS and humans by focusing mainly on 'teaming' arrangements such as situational awareness, information exchange, communications, team leadership, etc.

This usually requires us to examine how well a human-UVS team accomplishes some task and a number of metrics often used in this regard include: how quickly was

the task accomplished, how many mistakes did the team make, how much of some overarching goal was achieved, and so on. We must also allow for aspects of the automation and teaming which are imperfect: incorrect task allocation strategies, decision inaccuracy, incorrect recommendations, false alarms, missed alerts, etc.

As described in the Technology chapters of this book, at the outset an operator will task a system by specifying high level tasks, constraints, and priorities through the HMI. These tasks range in complexity from (say) data acquisition to maintaining surveillance over a region and responding to perceived threats. There are also a large number of constraints that are task dependent: e.g. no-go zones, no surface zones (for UUVs), communications ‘dead-spots’ and time constraints, which may be ‘hard’ or preference-based. In this regard, we must start by defining the complexity of the mission, where we might consider

- The geographic size of the area under surveillance
- How many vehicles (manned and/or unmanned) are required
- The impact of weather, terrain or other environmental factors
- The disposition or placement of the UVS within its environment
- The number of sub-tasks undertaken relative to each of the above
- The numbers of each type of UVS involved in each scenario
- The effect of mission re-planning time on each of the above
- The nature of missions and sub-tasks required in each scenario

We can then categorise missions in terms of the degree of risk involved. For example, having defined the degree of risk severity, we may choose to avoid or allow specific risks or degrees of risk – or in the case of multi-vehicle systems, allow certain vehicles to undertake certain risks. Alternatively, we may choose to allow UVS with certain levels of autonomy to undertake certain missions or risks.

In terms of how well a human-UVS team addresses its task there are a number of MOE that we can then use, such as planning time, mission time, mutual situational awareness, operator and UVS workload, usability and mental model (Figure 5.10). These may be metrics that stand alone or we might apply or combine them to determine how well a human-UVS team is able to carry out a mission or sub-task associated with one of the MOE in Figure 5.9, such as “total threat systems destroyed” or “number of response options available to red/blue team.”

In general, more work is required to refine both the measurement process by which we observe the key metrics and the distillation process by which we reduce the large number of available metrics to select an efficient set for any given evaluation. For example, while the previous section on mission level effectiveness provides a framework by which we might attempt this, the extant techniques are still deficient in regard to measuring mental models and trust. To date, such techniques tend to focus on observing the operator’s spatial mental awareness model of the task environment and understanding interface features [151], rather than (say) testing the operators on their capacity to predict how systems will react in particular circumstances.

Typical System MOE	Typical Measures of Performance
Planning	Time from presentation of task to execution of mission Percentage of mission spent planning (by UVS and human) Ratio of operator-initiated input vs. UVS-generated prompts Accuracy and completeness requirements of UVS inputs Level of operator assignment (mission or objective level)
Mission Tasking	Time to complete mission, tasks or sub-tasks Number and nature of tasks undertaken simultaneously Mission points at which key tasks or sub-tasks were dropped Task completion time vs. required mission completion time Accuracy of task completion vs. task completion time
Training	Quality of training needed for effective task completion Quantity of training needed for effective task completion Nature of training needed (interactive, presentation, book)
Situational Awareness	Accuracy of operator understanding of past & current events Completeness of understanding of past & current events Effectiveness of operator communication wrt these events Accuracy with which operator can predict future events ⁷⁸
Workload	Mental, physical, temporal mission demands on operator NASA TLX or Cooper-Harper workload ratings [124] Degree of effort or frustration experienced Impact on mission effectiveness (task shedding) Overall & critical component workload profiles Number and nature of tasks undertaken simultaneously
Usability	Likert scales of usability, consistency, reliability, etc. ⁷⁹ Lee & Moray trust scale rating [162] Degradation in task effectiveness over UVS neglect time How frequently & meaningfully do UVS/user communicate Appropriateness of info exchange between human & UVS Timeliness of information exchange between human & UVS
Mental Model	How much mental computation needs to be performed Temporal & spatial correspondence between human/UVS Ratio of operator taxed vs. fully occupied but handling tasks Percentage of time spent interacting with UVS

Fig. 5.10 Typical Measures of Effectiveness for Human-UVS Teaming

Other considerations useful in developing metrics for human-UVS performance assessment include such things as command and re-planning frequency, decision accuracy, error recovery, error impact, [124] [183] [259]. For example:

- Was an appropriate mission strategy selected?
- How appropriate was the strategy and how well was it executed?
- How quickly are decisions made and how 'correct' are they?
- How quickly can any decisions be submitted and interpreted?
- Did the system do what the human expected? Was this appropriate?
- How well does the automation handle unforeseen events or errors?
- Is the human taking control when and where appropriate?
- Is the UVS handling its duties when and where appropriate?

- How many feasible alternative strategies were derived and presented?
- How varied were each of these alternative plans?
- What is the quality of communication within and between teams?
- What is the complexity of the information flows?
- How much interpretation is necessary for the human?
- How much work is needed for the human to direct the automation?
- How much human resolution of ambiguities and uncertainties is required?
- Do the team members exhibit good feedback and backup behaviour?
- Is there agreement between the human and UVS perception of the situation?
- To what extent does the human accept the operational picture (or never accept it)?
- How many and what mixture of UVS were required for a successful mission?
- What was the impact of mission re-planning time on mission success?
- How did reactive or creative manoeuvres impact on successful missions?
- Did the loss of assets affect achievement of objectives and mission success?

Many metrics for evaluating team performance are useful in monitoring the ongoing progress of UVS missions, for example, relative to mission success or failure criteria. In regard to the achievement of tasks, however, progress towards them is usually best measured after the task has been completed [134] as progress on sub-tasks can advance and provide the impression that progress is being made towards the overarching goal, without this actually being the case. An example for a UGV might be the task “prosecute a mine” with the sub-task “advance towards target.” If the distance between UGV and the target is reducing, the UGV is providing an appearance of making progress towards its global goal. If, however, there is a ravine or an obstacle preventing the UGV from advancing towards X then the progress towards the global goal is nugatory.

The overall performance of the human-UVS relationship can also be measured by observing competencies that effective teams possess: knowledge, skills, leadership, task distribution, communication, decision-making, adaptability, and so on [130] [251]. For instance, trust between humans and UVS is essentially driven by a combination of the probability that the human can successfully predict the anticipated action of the UVS *before* he can monitor such action and the reliance he has upon the technology. Therefore, as well as measuring the systems performance of the human-UVS team, the metrics and methodology described here also provide a mechanism for monitoring and characterising user trust; and perhaps identifying tasks, task components, or periods when leadership can (or perhaps should) be assumed by the UVS rather than the human – or at least instances where the supervision of specific tasks is replaced by a more equal relationship that reflects true human-UVS teaming.

One of the complexities in this regard relates to the unpredictability of autonomous UVS in unfamiliar environments, which is compounded by the high workload environment of the modern battlefield and the cognitive and processing limitations of humans. As a result, as UVS increasingly possess more awareness of their own states and indirectly more awareness of their users workload, if the

information contained in these metrics is suitably manipulated it can also be provided as feedback to the UVS so that it too can use the framework to reliably anticipate likely human behaviour, thereby allowing a kind of ‘trust’ to be developed by the UVS in humans. One way to assess this aspect of performance and to regulate whether the UVS or the human has the control-initiative is measure the percentage of requests for assistance made by the UVS and/or human and the number of interruptions rated as non-critical [259].

The key to these tasks is perception. For a supervisor to have effective control of a task or system it is necessary for them to have situational awareness relative to the progress of the mission and there are a number of measures that can be used in this regard. For instance, [26] suggests that displays can be ‘blanked’ and the users asked to answer questions relating to the key features of the mission or make predictions about expected mission progress. Some subjective human performance measures in this regard include:

- Operator comprehension of the mission complexity
- Situational awareness during the mission
- Correct use of automation capabilities
- User trust in automated capabilities
- Efficiency and accuracy of decision-making
- Operator effectiveness in task prosecution

In UVS reliant upon human interaction, perceptual inference can be performed by the UVS, by the human supervisor, or by some combination of these two where (for instance) the UVS directs its supervisor’s attention to features of interest, but leaves the deductive reasoning to the human. Many UVS also abstract information from sensed data and fuse this with information about their state to complete higher-level situational awareness functions or to plan and execute other tasks intelligently. For instance, a Micro Air Vehicle (MAV) flying inside a room might make inferences about objects in its camera’s FOV based on its bank angle. Alternatively, a UGV might make inferences about the terrain ahead of it based on a combination of the data arriving from its environmental sensors and the degree to which its wheels are able to maintain traction. Based on its understanding of the environment and its current or potential states, the UVS might then seek particular new information, thereby increasing the perceptual confidence that it has in other previously observed data or information and/or its current state. As a result, we need to measure components of tasks:

- The ability of the human-UVS team (system) to observe its environment,
- The ability of the system to unambiguously interpret this information,
- The system’s capacity to disseminate this information to those who need it,
- The system’s capacity to fuse its sensed data with other relevant sources of data,
- The system’s capacity to accurately project likely environmental conditions, and
- The system’s ability to seek out particular new information pertinent to improving its ability to carry out any of these functions.

We also need to evaluate the impact of the human on the team decision-making process. Techniques (e.g. psycho-physiological methods) that focus on the human element are not discussed here, rather the reader is referred to a number of works on the subject (e.g. [45] [56] [124] [233]). Such functions, however, tend to be measured against metrics pertinent to effective critical decision-making, operator performance and workload on numerous dynamic control tasks when there are multiple competing goals and multiple, simultaneous task demands on a user's attention. Parameters typically varied include:

- Human factors (fatigue, stress, training, tasking, experience, etc)
- HMI (screen size, colour scheme, image contrast, screen resolution, etc)
- The target (type, size, shape, pose, range, camouflage, contrast, motion, etc),
- The sensor (type, spectral band, FOV, dynamic range, sensitivity, noise, etc),
- Environmental effects (visibility, clutter, weather, solar angle, attenuation, etc),
- Vehicle motion (vibration environment, sensor/weapon stabilisation, etc)

Yet another human factors technique used to measure operator expectations and stimulus-response compatibility is to assess the accuracy of mental models of device operation and the reduction in mental transformation of information, faster learning and reduced cognitive workload [95]. Measuring human situational awareness in this way can also be valuable in diagnosing performance successes and failures and identifying effective training or design intervention regimes.

Another technique is to measure the performance of the human-UVS system, assess the performance of the UVS (see the next section) and infer the impact of the human. In this regard, subjective metrics are often useful in qualitatively assessing the usefulness or value-add of information in a situational awareness display, the usability or consistency of information, or more complex matters such as standards compliance. Alternatively, in order to 'normalise' the tasks and avoid task-specific comparisons we might also attempt to measure user workload (e.g. [124]).

We can do this relatively easily if we allow users to subjectively rate their experience of mission difficulty and cognitive demands for the overall mission workload or for critical components of the task or mission, although objective measures are preferable. In this regard, we can employ two approaches: we can overload the operator with tasks and see how many of them remain. Alternatively, we can progressively introduce tasks until the operator's workload reaches saturation. In this regard, irrespective of the UVS operations undertaken, requiring the operator to respond to a stimulus (e.g. a buzzer) that requires him to enter a sequence of (say) eight digits is very effective as the frequency of the required input can easily be varied.

These metrics are also useful for identifying the task distribution workload of the human-UVS relationship and/or organisational structure within a team of operators [236]. In particular, they aid in determining where additional automation may be desirable. For example, humans fuse and manipulate situational awareness information very effectively, but tire of routine and mechanical tasks. Moreover, there are aspects of an HMI that can impede or simplify tasks. By measuring how well a system supports (say) the human workload we can derive valuable information about the HMI.

For example, we might attempt to measure the capacity of the HMI to allow users to accurately and directly perceive, comprehend and predict UVS activity and related significant events; its capacity to prompt for or autonomously allocate resources, monitor uncertainty and risk, and thereby provide automated assistance; and/or its capacity to correspond multiple sensory perspectives into a readily interpretable common operating picture. This may be achieved by using a range of categories that include [76]:

Design & Display Real Estate

- Was all of the key information readily available?
- Did the display omit any information?
- Was information difficult to access?
- Did information require time to retrieve?
- Was the information all clearly marked?
- Were font sizes & colours appropriate?
- Were there moving map displays?
- Were the displays/windows easily navigable?
- Was information 'buried' deep within a display?
- Were appropriate displays/windows located next to each other?
- Was a significant amount of cognitive resource required to navigate the displays?

Operator Attention

- Were operators distracted from their primary tasks by secondary tasks?
- Did individual displays (e.g. 'pop-ups') obscure an operators view of the HMI?
- Did the HMI alert the operator to any key events (or did it fail to)?
- Did the HMI make use of audio and visual cues/modalities?

Cognition & Workload

- Was the operator's mental workload high/low?
- Did the operator have to search for information?
- Did the operator have to mentally integrate/manipulate information from multiple screens, sensors, etc or was this executed autonomously?
- Did the displays support direct perceptual interaction?
- Did operators have to remember key information or was this embedded in the system (e.g. speed limits, separation zones, etc)?

Change Analysis & Situational Awareness

- Was the information up to date?
- Were events and UVS status reported?
- Were fuel/battery life/etc information reported?
- Were potential and past target trajectories shown or available?
- Were UVS options/paths presented?
- Was a risk analysis presented or available?
- How was novelty dealt with?

Planning & Execution Tasks

- Did the HMI permit effective and efficient decision-making?
- How was event and information cueing delivered wrt sensors, UVS & humans?
- Did the HMI permit constraint violations (i.e. allow UVS into a known threat area or provide plans that exceeded battery life)?
- Did the HMI attempt to negate any inadequacies in the operators' partial or conflicted understanding of information (e.g. if the UVS detected an 'unexpected' target, how did this affect its battery life)?
- Did the HMI's functions/display vary according to human or UVS workload?
- Were users prompted for or receive autonomously feasible resource allocations?

5.4.3 *Measuring UVS Performance*

To measure the performance of UVS we must attempt to measure the functionality of the system and its components in terms of their key aspirational drivers. For example:

- Persistence, low cost, stealth, ready deploy-ability and retrieve-ability;
- The capacity to detect, locate, track, identify & engage targets autonomously;
- The ability to gather, disseminate and act on several types of information;
- The capacity to network together and to the higher-value, manned assets;
- The capacity for individual platforms and sensor elements to self-organise;
- The systems' flexibility in relation to their deployment options;
- The level of risk or burden imposed upon the users/operators;
- The portability and robustness of the systems design; and
- Their sustainability and reliability in theatre

To do this systematically, we need metrics that allow evaluation relative to categories of tasks typically carried out by UVS but irrespective of the application or domain. For example, we might try to quantify the performance of a UVS in terms of its efficiency and effectiveness, where effectiveness is measured as the percentage of a mission completed with the UVS in autonomous mode (with the number, frequency and duration of operator interventions used as suitable metrics) and efficiency the time taken for the UVS to complete its task [259]. As UVS will finish many tasks if allowed time, however, we should also measure such things as the number of tasks completed regardless of human intervention versus the number completed autonomously.

There have been several attempts to develop taxonomies for measuring the performance of UVS (e.g. [89] [121] [259] [293]). As with measuring HRT/HRI performance, however, there is not yet a consensus on a standard framework. Furthermore, many of the frameworks strive to establish performance against what individual UVS were designed to do and hence what constrained their design.

Regardless of their design, role, or domain, however, military UVS operating in real world conditions have at least two things in common: they are designed to meet a set of requirements that are subject to constraints such as budget, schedule, etc, and they are used in environments that differ from those for which they were designed and tested.

We must therefore categorise key UVS functions: navigation and mobility, system awareness, situational awareness, management and coordination, and effector tasks. An additional category of Engineering Quality and Design is also needed to allow evaluation of the UVS against basic criteria such as design, construction, support, etc. This is not discussed here, but allows evaluation of the many mechanical, electrical and software engineering design, quality assurance, customer support, documentation, disposal requirements, etc as a function of parameters such as cost, availability, time to delivery, etc. A range of standard evaluation techniques are available for assessing such criteria.

Some studies (e.g. [259]) have also added a 'social' category (i.e. human cognition and interaction). This category is designed for robots that interact with humans socially or simulate social intelligence. In this regard, this category would measure interaction characteristics such as persuasiveness, trust, engagement or compliance, all of which can provide significant insight into the effectiveness of robot design. For military robots, however, the typical modalities of HMI do not yet include this sort of interaction. Furthermore, we incorporate measures of trust, engagement, and compliance within the Systems Level evaluation.

By measuring how well the UVS undertakes each of these, either directly or in combination with the systems performance measures, we may infer how well the UVS performs at various levels. Furthermore, when used in conjunction with the systems performance measures we may feedback the outcomes to perfect the UVS design options against mission techniques, technology components, CONOPS, etc.

Navigation and Mobility Tasks

As a fundamental task carried out by an autonomous UVS navigation is usually measured by how well the UVS determines where it is, where it needs to be, how it gets from where it is to where it needs to be in terms of its trajectory management and resource usage, and how it deals with environmental contingencies along the route. This is accomplished by measuring such things as:

- Global navigation capabilities (e.g. where the UVS is in its environment)
 - Success rate in reaching navigation waypoints
- Local navigation capabilities (e.g. what potential hazards exist in its locale)
 - Comparison of internal, external and environmental state estimates
 - Number and accuracy of terrain features tracked
- The efficiency with which navigation is carried out (e.g. time to complete a route)
 - Distance travelled to reach waypoints
 - Steps required to reach waypoints

- The effectiveness with which obstacle avoidance functions are executed
 - Time taken for computation of hazard detection
 - Number of obstacles detected, avoided, etc
 - Number of obstacles that could not be avoided
 - Percentage of navigation tasks completed
 - Deviation from planned routes
 - Coverage of area
- The effectiveness with which obstacles are encountered
 - The characteristics (size, hardness, etc) of obstacles that can be negotiated
 - The degree of environmental difficulty encountered
- Operator confidence in the navigation capabilities
 - Number of operator interventions per unit time (planned or unplanned)
 - Ratio of operator input time to UVS navigation time
- Average and maximum speed over terrain
 - Ratio of average to maximum speed
 - Percentage time spent at maximum speed

Many of these navigation and mobility tasks are best measured comparatively against courses or arenas specifically designed to stress the techniques under consideration. For example, a maze of walls, elevated floors, slippery surfaces, complex terrain, and sensory obstacles intended to confuse specific perception techniques can all be used to test the capabilities of the UVS in this regard.

For autonomous UVS, many navigation tasks are based on the ability of the UVS to perceive or sense its environment. As a result, some metrics for navigation are also useful for measuring aspects of (say) situational awareness tasks. Similarly, many other aspects of UVS operations (e.g. path-planning) are measured indirectly either through the navigation tasks outlined above or those that follow.

For example, a UVS may navigate by determining the absolute or relative distance to an object (i.e. how long will it take to reach a landmark based on its size). Alternatively, the accuracy with which ego-motion or the movement of objects within the environment can be determined allows absolute or relative estimates of the UVS velocity to be made (i.e. how long before the UAV collides with terrain or another aircraft). Both these examples may also be used to infer indirect measures of situational awareness.

System and Self Awareness

The extent to which a UVS can accurately understand its current and likely future states and capabilities directly impacts its ability to efficiently interact with its supervisors and to make appropriate decisions relative to its current or likely context. The less aware the UVS is of these aspects the less likely it is to understand when it is having trouble. For instance, a UGV that knows when it is lost may be able to prompt for human intervention or ignore information from particular sensors. Self-awareness is also very useful when attempting to

determine when reliance upon human input is beneficial. A number of metrics for measuring self-awareness have been proposed [293]:

- The capacity to understand a UVS' mobility, navigation, or sensor limitations;
- The capacity to monitor health, resource usage, task progress, response options;
- The capacity to recognise deviations from any nominal trajectories and plans; and
- The capacity to detect, isolate and recover from faults during both the mission-planning and task execution phases of any operations.

Another measure of the system's self-awareness is its ability to operate independently, particularly in the presence of environmental difficulty. For instance, environmental difficulty is typically evaluated against the concept of a 'solution ratio', which is typically the ratio of the number of total possible choices the vehicle can make versus the number of solutions that meet the mission or task objectives [134]. When UVS use these metrics for the purposes of decision-making, thresholds can be set on the degree of difficulty, based on cost/benefit/risk factors, and then a determination made as to whether to accept or reject a solution. For example, the UVS may attempt an undertaking beyond the physical capabilities of the UVS: an entrance is narrower than the width of a UGV or its mobility characteristics are insufficient to traverse certain terrain (i.e. the identified navigation solution is infeasible), there is clearance for the vehicle or it can cope with the mobility, but it requires high level perception, planning, and execution capabilities (i.e. the solution is restrictive), or there is open space or terrain that does not require advanced computation (i.e. the solution is unrestricted).

There are several application-specific methods for measuring self-awareness. Neglect Tolerance [216] measures the performance of a UVS (or more accurately the degradation in its performance) over time when a supervisor is not attending to it. Several methods are available for measuring this parameter (e.g. [76]). However, while the metric itself serves as a good indicator of UVS autonomy it is not decoupled from factors such as task complexity, UVS capability (i.e. individual or collaborative autonomy), the HMI, or human performance. As a result, the metric may only be used to gauge an overall measure of UVS autonomy rather than specific details such as failure modes [211].

Finally, as mission accomplishment is often driven by a combination of the probability that the human can successfully predict the anticipated action of the UVS *before* he can monitor such action and the reliance he has upon the technology, in addition to measuring the UVS system's self-awareness we can attempt to measure the degree to which it is able to understand the intentions and state of its human supervisors. As previously alluded to, this also provides us with a mechanism for monitoring and characterising user trust and perhaps identifying tasks, task components, or periods when leadership can be assumed by the UVS rather than the human, or at least instances where the supervision of specific tasks can be replaced by a more equal relationship.

Under these circumstances, the UVS may need to be sensitive to its user's expectations, constraints or intentions. Clearly, its level of awareness will depend upon the degree of autonomy that the UVS is expected to achieve and the roles played by the humans [76]. Furthermore, UVS awareness of humans implies human competency, the proficiency of which must also be measured. For this, [293] proposed the metric of 'awareness violations' (information that should be provided that is not) to determine the level of human-oriented perception held by the UVS.

Situational Awareness Tasks

Most military UVS conduct some form of situational awareness (SA) tasks. However, we need to distinguish between SA carried out by UVS and that carried out in conjunction with or on behalf of a supervisor (e.g. for a specific military purpose) as in the latter case the human performs the interpretation.

Measuring the capacity of a UVS to observe its environment is relatively straightforward although performance depends on the combined capabilities of onboard sensors and any related processing. Performance is typically assessed in terms of:

- *Detection*
 - Determination that an object may be of military interest
 - Possibilities might be "take closer look"
- *Classification*
 - Object can be discriminated by class
 - Possibilities might be "tracked vehicle" or "a human"
- *Recognition*
 - Object can be distinguished by category within a class
 - Possibilities might be "military tracked vehicle" or "carrying an object"
- *Identification*
 - Object is distinguished by model
 - Possibilities might be "M1A2 Abrams Tank" or "rifle or axe"
- *Feature Identification*
 - Targets distinguished by model, individual elements of clothing, etc
 - Possibilities might be "M16 or MK47" or "US or Australian uniform"

Common measures for these include

- *Probability of detection* – the probability of correctly discriminating an object in an image from background and system noise.
- *Probability of classification* – the probability of correctly determining the class of a detected target.
- *Probability of recognition* – the probability of correctly determining the class membership of the target.

- *Probability of identification* – the probability of correctly determining the exact identity of the target
- *Probability of false alarm* – the probability of an error in detection.
- *Probability of re-acquisition* – the probability that, once acquired, a target can be correctly re-acquired if lost for some period of time

Using these metrics, we can then evaluate functionality against metrics such as

- *Signal-to-noise*: The observations are corrupted by noise, clutter, interference, object orientation, contrast, etc.
- *Receiver-Operator-Curve (ROC)*: The performance of a system degrades as the signal-to-noise ratio decreases. It is also useful to plot probability of detection against false alarm rate as a function of signal-to-noise ratio to produce an ROC.
- *Confusion matrix*: This is a 2D array that indicates the identity assigned to an object by the ATR algorithm (i.e. how often a rifle is confused with an axe; in the case of detection, the confusion matrix reduces to the classical false-alarm rate).
- *Consistency*: This is a measure of how often an ATR algorithm gives the same declaration for successive image frames in the same context. The difference between frames is noise from the atmosphere, the sensor electronics, etc.
- *Efficiency*: This is a measure how much resource, effort or time was required to perform target detection, identification, etc.
- *Improvement*: This is a measure how and how quickly the above performance measures improve over their initial detection.

We commonly interpret these in terms of performance metrics such as detection range, the time needed to process and/or disseminate critical ISR information, and the ability to communicate with other manned or unmanned elements in the force. We also need to measure the accuracy with which a UVS might control the pan and tilt of its sensors relative to the orientation and motion of the UVS or in relation to a particular surveillance or targeting operation. For instance, in regard to the above performance measures we need to establish the time and the resources required to detect and recognise the targets, the amount of sensor motion, the accuracy with which the targets can be located, etc. These targets might include a specific shape (e.g. a human form), a heat source, motion, sound, gaseous emissions, and standard signs (e.g. eye charts, Snellen's tumbling E's, or hazard signs) [139]. Generally, these will be characterised against the time and resources to search for a target, the sensor coverage as a percentage of potential area coverage, operator confidence in sensor coverage (i.e. the number of correctly/incorrectly identified targets, the number of targets missed, and so on).

We might also wish to evaluate the ability of the system or its components to achieve or contribute to higher levels of situational awareness. For example, we

might seek to assess the capacity of the UVS to accomplish each of the following in terms of accuracy, speed, efficiency and cost-benefit:

- Sense the environment and the internal state of the machine
- Perceive and recognize objects, events, and situations
- Remember, understand, and reason about what is perceived
- Attend to what is important and ignore what is irrelevant
- Predict what will probably happen in the future under a variety of assumptions
- Evaluate what is perceived and predicted
- Make decisions, plan, and act so as to achieve goals
- Learn from experience and from instructions

Accuracy can be measured in terms of variance between goals and achievements; speed in terms of bandwidth and latency; cost in terms of resources consumed, risk and the consequences of failure; and, benefit in terms of pay-off for goals achieved.

These higher-order perception tasks are often the outcome of data or information fusion tasks and algorithms as single modality UVS are usually incapable of identifying and recognising complex environmental structure at a detailed level.

Management and Coordination Tasks

Management and coordination tasks require us to measure such things as the vehicle-to-human ratio, the human-robot interaction, systems performance, problem recognition, teaming, degree of heterogeneity and information sharing. Essentially, this relates to the difficulty in handling the UVS during use and the traditional way in which this is measured is through the use of the *fan out*, *intervention*, *attention demand*, and *free time* metrics [118]. This is a measure of how many robots (with similar capabilities) can be effectively controlled by a human and the amount of operator intervention that needs to be devoted.

The **fan out** measure directly relates to the logistical demands of the UVS deployment, its handling difficulties and the cost-benefit equation. For example the Global Hawk UAV requires around 20 people to operate it, whereas cooperatives of much less expensive (experimental) UGVs require only a few operators. This measure is also a good indicator of the upper limits of the physical and cognitive workloads of operators and the degree of intervention required to manage the UVS. Additionally, when the numbers of UVS are large and they are supervised by teams of humans, metrics and methods pertinent to the management of air traffic control may be useful [231].

The **intervention** metric usually measures the physical or cognitive intervention response time either from when the operator first recognises the problem or from when the UVS first requests assistance. Response time can also allow for specific details of the intervention or latencies to be measured, such as the time to deliver the request from the UVS, time for the user to notice the request, situational awareness or planning times, and execution time [259]. The

attention demand metric measures the amount of time a UVS requires; that is the fraction of the time given by a user to operating the UVS, whereas **free time** is a measure of how much time is left over from other tasks. The aim is clearly to maximise free time. However, we will wish to ‘fill’ this free time with multiple UVS operations, which means that free time alone is insufficient for characterising efficient human-UVS interaction.

Increasingly, UVS are being designed with variable levels of autonomy, some being more appropriate to specific tasks or missions than others. In these cases, by measuring the human’s ability to rapidly and accurately identify the appropriate levels of autonomy for task execution, we may draw conclusions about the UVS’ ability to activate autonomy correctly and the HMI’s capacity to communicate information relating to management and control tasks. It should be noted, however, that the performance of the human is coupled into this metric. Furthermore, as per neglect tolerance and UVS system self-awareness, while the metric itself serves as a good indicator of system efficiency, it encompasses several factors such as situational awareness, user trust, etc.). However, for particular missions, where the optimal levels of autonomy are known, the metric may be used to determine whether the UVS responds appropriately. With respect to the UVS mission complexity, we might consider [130] [146] [216]:

- What is the level of human-UVS collaboration required
- What are the time or precision constraints for the mission
- Are any coordination or synchronisation behaviours required
- What is the level of resource/asset/payload management required
- What is the authority hierarchy, for data access, plan execution, etc
- What is the degree of adversarial conduct and/or rules of engagement
- What are the mission risks and requirements for survivability
- What types and amounts of information are required for the mission
- How much and what type of information is available a priori
- What degree of uncertainty is associated with this information
- How easily can the mission outcomes be predicted
- What are the sensory and processing requirements

We must also try to establish metrics with regard to task objectives [8]

- How does the complexity of the environment affect performance?
- How does the number of UVS affect the level of performance?
- How does the distribution of UVS (or targets) affect performance?
- What tasks/environments require multi-UVS cooperation?
- What tasks/environments are improved by multi-UVS cooperation?
- What cooperation emerges as a result of interactions between UVS?

In regard to multi-UVS missions the following could also be measured [130]

- What were the number of simultaneous tasks or missions that could be handled
- How many different task or mission types could the system handle

- How many tasks were dropped vs. the total number of tasks undertaken
- What proportion of the tasks/missions were not achieved (for feasible tasks only)

Effector Tasks

Many UVS intentionally have an impact on or interact with their environment; they may carry weapons, manipulator arms, electronic jammers, etc. Alternatively, they may move objects, drop off or collect payloads, or cause action to be taken by an adversary. As a result we must try to establish what the intended effect is, how this is to be done, and then evaluate the outcome.

Unfortunately, these tasks often take place through sensors that have limited spectral, spatial or temporal modality, over communications links that have latencies, or through HMI that provide 2-D orthographic projections of a 3-D world perceived from an unnatural perspective. As a result, mental workload – which is strongly influenced by demands made on short and long-term human memory [110] – is often stressed during effector tasks and is heavily influenced by the number and nature of features in the environment. Consequently, key measures in almost all effector tasks are the ratio of intended to unintended ‘victims’ and the physical or cognitive processing requirements placed upon the human operators. For example, mental ‘manipulation’ activities that must typically be performed by operators might include object identification, association, rotation, and target-tracking.

We must therefore measure the effect at either an outcome or a ‘contact’ level. For example, a UGV attempting to remotely or autonomously defuse an unexploded bomb (UXB) may be assessed against either the number of inadvertent contacts its manipulator arm makes with unintended elements of the environment or the positional accuracy with which the arm or tools can be placed. Furthermore, the type of contact can be measured: glancing, hard or soft, each against a backdrop of the environmental complexity [259]. Alternatively, the overall outcome or effect of the manipulation errors can be measured, although metrics such as mission success/failure would normally be measured against force effectiveness criteria.

Equivalent metrics obviously exist for other effector tasks such as electronic jamming and kinetic weapons, etc. For example, metrics for quantifying the effectiveness of the safe arming and firing weapons are likely to be closely related to those used for safety-critical systems, where the list of potential ‘inadvertent contacts’ might include: the degree of unintended system operation, the number of inadvertent firings of the weapon, and the number of unintended targets engaged. Similarly, the impact of a weapon will probably be related to the trajectory and impact of its projectile, the target, the target’s environment, any vehicle dynamics, the accuracy and resolution of sensors, the accuracy of target tracking algorithms, the latency and jitter in any feedback control loops, the inertial response of the weapon, weapon recoil, and so on, all of which are reasonably straightforward to characterise.

Chapter 6

Legal Issues for UVS

Many of the advances in computing and technology have increased the decision-making capabilities of UVS to the point where they may now make truly independent decisions regarding by which route they might travel, what constitutes a target, and whether and by what means these targets should be engaged, possibly with lethal force. As a result, it is now widely acknowledged that, if permitted, the use of such systems would mark a sea-change in the role of technology in warfare as the human could potentially be removed from the decision-making loop. This introduces a number of issues that the need to be assessed in relation to the Law of Armed Conflict. In this chapter we discuss these issues.

Additionally, as UVS will need to work in environments shared by people, property and other vehicles, we also discuss the legal status of UVS, their operation in a shared environment, how they fit into the existing legal frameworks, who might be responsible for any infringements perpetrated by them, whether or not they should (or could) hold some sort of legal personality, and what the implications of such concepts might be.

It is recognised that such considerations will not resolve the issues over whether military UVS will one day operate in complex adversarial environments in a more reliable manner or autonomously target and control weapons. Moreover, there are many level-headed people who will consider it futile to discuss such concepts when we clearly have trouble making relatively simple UVS work. Regardless, such a discussion contributes by providing a way of thinking about the issues from the perspective of legal responsibility. Additionally, such deliberations also help to avoid technology outpacing the regulatory regimes and allow us to develop an understanding of existing technologies in the light of certain hypothetical possibilities, as well as exploring established legal concepts. This in turn enables us to consider what the responsibilities of users, owners, developers, etc. might be given the autonomy of UVS.

Furthermore, without a full understanding of such issues developers may be forced to focus on making supervised systems rather than fully autonomous ones, or users may refuse to accept liability for the unsupervised actions of UVS, thereby creating capability gaps. Additionally, if we do not frame the responsibilities of the UVS and those responsible for their construction and use, they will likely develop unfettered, with some humans progressively absolving themselves of any liability. As a result, adversaries who may be less restrained in

their interpretation of their international obligations may accord UVS unlimited responsibility for their actions. In the case of the Law of Armed Conflict this may have far-reaching consequences.

6.1 Legal Issues for UVS Platforms

Regardless of whether or not UVS carry weapons they are usually large objects that move at speed and must avoid people, property, and other vehicles. They are also constructed of sub-systems and materials that are inherently hazardous. At one level this means we can treat UVS as “unremarkable technological artefacts, similar in nature to toasters or cars” [24] and for which the law has a highly developed set of principles that apply to product liability. That is, we can apply the law to UVS as purely commercial products and as there are many examples of intelligent, safety-critical systems whose malfunction may result in death or injury (e.g. medical equipment, railway signals, air traffic control systems) it is believed that the majority of the concerns usually touted as being possible dangers for UVS or their weaponised counterparts will fall under this mundane interpretation.

At present, however, as a result of the novelty of the technology, the difficulties with allocating responsibility, and an unwillingness to burden the relevant agencies responsible for drafting such laws with additional work, the use of most UVS is inadequately regulated. Nevertheless, there is a great deal of interest and investment in building autonomous UVS of increasing complexity and it does not seem an unreasonable prediction that within a decade we shall see a fully autonomous unmanned combat aerial vehicle (UCAV) in service. Moreover, as occurred with cruise missiles and GPS, where within a decade of their first use during Operation Desert Storm, they will likely then become vital components of military considerations for national defence forces. Consequently, as UVS become more sophisticated, unlike tethered or remotely operated vehicles, which are effectively extensions of their human operators, issues of liability will become more complex. Sometimes accidents will happen, but if UVS do not demonstrate sufficient and predictable capacity to obey the *rules of the road*,⁸⁰ matters of liability will be raised in the courts. It is perhaps instructive therefore to quote from [239]:

“In 1936, a Duke University law student published an article summarizing the path of automobile liability law. He observed that in 1905 all of American automobile case law could be contained within a four-page law review article, but three decades later, a “comprehensive, detailed treatment [of automobile law] would call for an encyclopaedia.” That law student was Richard M. Nixon, who would later become President of the United States. His conclusion was that courts were mechanically extending ‘horse and buggy law’ to this new mode of transportation in most doctrinal areas. However, some judges were creatively crafting new doctrine in certain subfields of automobile accident law by stretching the legal formulas at their command in order to reach desired results.”

⁸⁰ The phrase ‘rules of the road’ is used as an expedient way to describe the legal framework within which a vehicle must comply when traveling from point-to-point, regardless of whether it is on or off-road, in the air or at sea (on or under the surface of the water), or operating within the Law of Armed Conflict.

Richard Nixon's observation appears to apply equally well to the age of unmanned vehicles as, at present, the use of UVS falls within a regulatory gap. That is, the technology appears to be under a loose legal framework, is self-regulated and/or only allowed to operate in restricted areas. There are many legal precedents connected to the use and production of UVS as purely commercial products and the related principles of responsibility and product liability are discussed in the relevant sections on Accountability and Liability. Before we do this, however, let us consider the legal status of UVS and their operation in the presence of other users of their environments.

6.1.1 Maritime Vehicles

As we have seen from earlier sections, UUVs have existed for many years, although their development and usage in the last 20 years has increased significantly. There is a recommended Code of Practice for UUVs, which was produced by the Society for Underwater Technology [60]. Despite the comprehensive nature of this document, it is not legally binding and the safety and legal framework for UUV use has not yet been formally adopted. As a result, the growth of the technology's use appears to be evolving under a loose legal framework such that the UUV community has had to be self-regulating. In this section, we present a condensed version of that contained in [60] [126] and [253] and readers are referred to these documents for more information.

There is a long history of International Maritime Law (also known as Admiralty Law) governing maritime issues and offences and the relationships between entities operating vessels on the oceans. It is distinguished from the Law of the Sea (also a body of public international law) that deals with navigational and mineral rights, jurisdiction over coastal waters and international law governing relationships between nations [148].

Although not endorsed or ratified by all nations, the 1982 United Nations' Convention on the Law of the Sea (UNCLOS) [274], which proscribes rules of navigation for vessels at sea, and the International Regulations for Avoiding Collisions at Sea (COLREGS) are the most widely used legal references in this regard. These conventions are usually enforced through local Coast Guards and courts. However, they were specifically written for guiding human behaviour and are not suitable for direct input to UUV control systems.

6.1.2 The Legal Status of UUVs

First, it is not clear that UUVs would be classified as "vessels" as the majority would fail the test against most formal definitions⁸¹, either because they are used

⁸¹The term "vessel" is undefined in UNCLOS [271]. COLREGS defines it as "every description of watercraft, including non-displacement craft and sea planes, used or capable of being used as a means of transportation on water." The International Maritime Dictionary [148] describes a vessel as "a general term for all craft capable of floating on water and larger than a rowboat. The term vessel includes every description of watercraft or other artificial contrivance used or capable of being used as a means of transportation on water".

to explore the ocean environment for scientific or military purposes or due to their size and design or because they could not be used as a means of transportation (although the question may be posed, “transportation of what” [126]). Unfortunately, the clearest definition (“a human-made device, including submersible vessels, capable of travelling at sea”) is non-binding and therefore has no legal status⁸² [40].

Despite this, from a purely practical point of view, it is likely that UUVs will be considered some form of vessel and will be required to comply with international maritime law. Following the “rowboat rule” [148], however, it may be that small UUVs fall under one set of laws and larger ones under another. That is, there are legal precedents to the effect that small watercraft are considered to be under control from where they were launched rather than nearby support vessels (i.e. from where their operators were receiving instructions). UUVs may therefore achieve some vicarious status on the basis of their size or launch platform.

Another possibility is that UUVs might be considered “warships” (which need not be armed). According to UNCLOS, however, a warship is a ship “belonging to the armed services of a State bearing external markings distinguishing such ships of its nationality under the command of an officer duly commissioned by the government of the State and whose name appears in the appropriate service list or its equivalent, and manned by a crew which is under the regular armed forces discipline” [274]. UUVs, which do not have a commander and crew, would not seem to qualify. Those deployed and/or remotely operated from a support vessel might be considered to qualify under extensions of that ship, but would presumably enjoy the same level of sovereignty as the support vessel (i.e. they would be immune from seizure, etc⁸³).

If considered vessels in their own right, but not warships, UUVs might be considered Auxiliaries, which are “vessels other than warships owned or operated by or under the exclusive control of the armed forces” [126]. Unfortunately, as previously mentioned, UUVs may not qualify as vessels so this is still not clear. Regardless, auxiliaries also enjoy sovereign immunity, which means that UUVs might still be protected from seizure by a foreign state.

Finally, UUVs may be considered weapons, particularly if they are weapons delivery platforms (i.e. a torpedo that carries torpedos). This has particular relevance to where and how such systems may be used in foreign waters [60].

⁸² This definition is provided by the American Branch of the International Law Association (ABILA) Law of the Sea Committee (www.ambranch.org), a non-governmental association with UN consultative status.

⁸³ It is a well-established principle of international law that warships are an extension of their respective states and cannot be seized, boarded or searched without the permission of their commanding officer. They are, however, legitimate targets in a war. Somewhat perversely, if UUVs do not qualify as warships or vessels they may also not qualify as legitimate military targets under the LOAC, even though an armed enemy merchant vessel, a merchant vessel acting as an auxiliary, one conducting intelligence operations or directly belligerent acts on behalf of its armed forces also constitutes a legitimate target [126].

6.1.3 Operations in Ocean Areas

There are different classifications of ocean areas in which maritime vessels operate, each of which determines the degree of control that a coastal nation is able to exercise over foreign vessels operating in those areas. They typically include: Internal Waters, Territorial Waters and High Seas, although more recently have come to include Exclusive Economic Zones (EEZ) and Archipelagic Waters [274].

Users of the High Seas enjoy complete freedom of navigation. An EEZ extends 200nm from the low water mark of a country. Navigation in an EEZ is not restricted as long as foreign vessels do not interfere with resources. Consequently, as long as UMVs can navigate with respect to other users and pose no territorial threat to the coastal nation, they would seem to be able to operate with freedom in these zones.

Sailing through International Lanes and Archipelagic Zones requires vessels to “proceed without delay” via “continuous and expeditious transit” [126]. This may prevent a support vessel stopping to launch or recover a UMV and/or a UMV performing sweep searches. Nevertheless, the protocol does allow submarines (and presumably, therefore, UUVs) to transit submerged. Unfortunately, research and survey activities are prohibited so REA and ISR-related activities may not be possible without consent. That said, the protocol allows self-defence against suspected threats, so a warship would be entitled to deploy a UMV ahead of its path for the purposes of force protection, say for a UUV to undertake mine countermeasures or a USV to guard against a RHIB⁸⁴ terrorist threat.

Vessels wishing to navigate through Territorial Waters, which extend 12nm from a coastal nation’s low water mark, will be subject to even more stringent regulations. Like those passing through International Lanes and Archipelagic Zones a vessel’s passage must be “continuous and expeditious” [126]. However, its transit must also be for a specific purpose and must not be “prejudicial to peace, good order or the security of the coastal state” [272]. This includes “any exercise or practice with weapons; any act aimed at collecting information to prejudice the defence of the coastal state; the carrying out of research or survey activities; the launching, landing or taking onboard of aircraft; the launching, landing or taking onboard of any military device; or any other activity not having a direct bearing on passage” [272]. Additionally, as submarines must both travel on the surface in these waters and show their flag the same must be assumed for UUVs.

From a practical standpoint, the use of UMVs in foreign territorial waters could face considerable legal challenges, particularly if they carry weapons even though it is the conduct of a vessel that defines the nature of its passage, not its class.

6.1.4 Navigation Environment

Safety requirements refer to vessels operating on the surface of the ocean being required to carry internationally agreed lighting and signal equipment (e.g. a white masthead light, visible for up to three miles for inconspicuous or partly submerged

⁸⁴ RHIB – Rigid Hulled Inflatable Boat (i.e. a rubber dinghy).

vessels or multiple lights for vessels longer than 12m) [253]. The vessels are also required to carry equipment, which varies in accordance with the size of the vessel, so they can be heard and can respond in a particular fashion to allow safe passage. Presumably UUVs (including UUVs navigating on the surface) would be required to carry the same equipment and to operate it according to these regulations.

Strict interpretation of these procedures also implies that a vessel must maintain an “able lookout” at all times [274], although for UUVs it is not clear what the term lookout might mean. While in their experimental stages of development lookouts might be considered to be on the support vessels, but it is an open question as to whether the separation between the support vessel and the UUV constitutes able, particularly as the separation distance between them grows.

In addition to the basic capacity to navigate and avoid other vessels and obstacles, UUVs will also require knowledge of signs, rights of way, etc. and an understanding of the guidelines consistent with standard maritime operating conventions. For instance, there is a hierarchy of vessels at sea so that mariners can operate safely and understand their responsibilities in regard to ‘giving way’ to one another. From highest to lowest, the list is: not under command; restricted in ability to manoeuvre; engaged in fishing; sailing and underway; power-driven and underway; and, sea-plane underway [41].

It is not clear where in this hierarchy a UUV should sit. In one sense, declaring the UUV to be “not under command” seems appropriate, but unless there was some technical malfunction⁸⁵ this would not be accurate as the classification describes a vessel without command. An alternative would be to declare the UUV “restricted in its ability to manoeuvre” or “an obstacle” although given the level of sophistication likely to be onboard and the fact that it is motorised and autonomous, it seems more sensible to at least attempt to get it to follow the relevant regulations. In the final analysis, if a vessel is incapable of avoiding a collision it is considered “unseaworthy” and therefore negligent; and liability will be imposed in any collision. A similar set of considerations will be needed in regard to UUVs becoming entangled with nets. In the final analysis, the specific definitions and conveyance of risk and liability are likely to be defined by civil law arising through normal operations: conclusive determination will need case law.

Benjamin [39] makes reference to an interesting case in which a US warship (which was blacked out) collided with an Australian ship in harbour at night. The Australian ship was not blacked out and was displaying the requisite navigation lights. Even though the warship maintained it had right of way due to the “starboard rule”⁸⁶, because it also had the ability to see other ships without being seen, it was held to be at fault as it had the last opportunity to avoid a collision.

⁸⁵ It is not clear how the UUV would recognise any such malfunction and, to comply with regulations, then declare itself not under command.

⁸⁶ The COLREGS make it clear that vessels should not assume a “right of way” – it is either the “stand on” or “give way” vessel. In other words for two ships on a collision course, the ship on the left (the *give way* vessel) must give way to the ship on its right (the *stand on* vessel).

This may have bearing on the safe operation of UMVs, which may be deemed not to be able to ‘see’ other ships.

Environmental and regulatory regimes also exist to protect marine mammals from noise and harassment. For instance, one may not “harass, hunt, capture or kill or attempt to harass, hunt, capture or kill any marine mammal” [253]. It is thought highly unlikely that a UMV will collide with a marine mammal as the UMV moves slowly, but its operation may constitute harassment. However, it is not clear what operations, beyond those of manned vehicles, UMVs might undertake in this regard except that they may well be used in shallower waters.

Finally, certain operations will require that some UMVs carry weapons (UCMVs). Just as with UCAVs, then, procedures will need to be developed and legal arrangements made regarding emergency operations for UCMVs. For instance, in the event of engine problems, loss of a command and control signal, weapon malfunction and so on, the UCMV will need to follow precise and fail-safe procedures. These may include pre-planned trajectories or self-destruct points, but at present are not universally agreed or enforceable.

6.1.5 Air Vehicles

The foundations of international civil aviation and the related legislation were laid during the Chicago Convention on International Civil Aviation held in 1944 (which was ratified in 1947). This convention charged the International Civil Aviation Organization (ICAO), an agency of the United Nations, with coordinating and regulating international air travel. It also established rules of airspace, rules for aircraft registration and safety, and codified the principle that each nation maintains sovereignty over its airspace [136].

Article 8 of this convention states “no aircraft capable of being flown without a pilot shall be flown without a pilot over the territory of a contracting State without authorisation by that State and in accordance with the terms of such authorisation. Each contracting State undertakes to ensure that the flight of such aircraft without a pilot in regions open to civil aircraft shall be controlled as to obviate danger to civil aircraft” [136]. Article 3 of this convention recognises that “state aircraft⁸⁷” are exempt from civil regulations; although the article also stipulates that national regulations for state aircraft must “have due regard for the safety and navigation of civil aircraft” [136].

Military UAVs⁸⁸ are considered state aircraft [161] regardless of whether or not they carry weapons. Those that are not state aircraft must abide by ICAO and national regulations. Beyond the technology of the air vehicle, therefore, even though their status may not necessitate them to abide by civil regulations, the fundamental problem for most UAVs is that their practical and safe operation

⁸⁷ The Chicago Convention does not define the term “state aircraft” but suggests that the term is determined functionally by the use of such aircraft by the military, customs services, or the police.

⁸⁸ Some nations actually rent some of their military UAVs (e.g. Australia and the US) from civilian corporations. As civilian aircraft these are managed in accordance with standard airspace regulations.

requires them to interact with, and hence to have due regard for, civilian manned aircraft.

While this clearly makes good common sense, from a legal perspective it means that, if they are to operate in uncontrolled airspace, the UAVs must effectively comply with the regulations approved by the national and regional administrative bodies for airspace and operations. Unfortunately, the precise meaning of the term “have due regard for” – and the commensurate civil regulations and procedures – still needs to be codified by most of the regional bodies and ICAO⁸⁹.

At present, most authorities restrict the use of UAVs to special airspace where civil and commercial aircraft access is strictly controlled. They also prevent them from flying over populated areas. UAV operations within uncontrolled airspace usually require significant notice to national air traffic administrators and near-perfect weather conditions. Military operations also usually manage the use of airspace for UAV operations on a temporally or spatially segregated basis.

The key starting point in regard to operations in uncontrolled airspace is that the UAVs pose no greater risk to persons or property on the ground or in the air than that presented by a manned aircraft.⁹⁰ That is, the UAV must provide an equivalent level of systems performance to that of a manned aircraft and be as reliable. These issues of performance and reliability need to be addressed from the perspective of the structural integrity of the UAV, the systems and sub-systems performance, the stability and control of the UAV and the procedural and regulatory regime. Additionally, most regional authorities recognise that the air traffic management procedures for UAVs should mirror those applicable to manned aircraft and the provision of air traffic services to UAVs should be transparent to air traffic controllers (i.e. UAVs must fit in with other airspace users rather than requiring the existing users to adjust to accommodate UAVs).

Fundamentally, there are three aspects of safety legislation that must be addressed [103]:

- *Airworthiness*: The design of the aircraft must be approved; the aircraft must be manufactured in accordance with this design; and, the aircraft must be maintained in accordance with appropriate maintenance and configuration control procedures.

⁸⁹ There are a number of ongoing initiatives and are taking place at the national and regional level (e.g. JAA/EUROCONTROL, ASTRAEA, NATO FINAS WG) and several that have concluded (e.g. ERAST and ACCESS 5).

⁹⁰ At present, UAV's cannot reliably detect other aircraft and conflict situations so they are unable to share air or runway space with them. To enable this, they will need a fault-tolerant, multi-function ‘pilot’ capable of operating day or night and in all-weather conditions, with the capacity to at least replicate a human's ability to sense and avoid problems. In other words, the UAV must be able to sense other aircraft in its operating environment, monitor the health of its component systems (e.g. sense loss or corruption of communications, sense structural, systems, and other onboard failures), and take appropriate action on the basis of the situation. Since not all aircraft carry transponders the UAVs will need to use onboard sensors to detect other aircraft, and ideally coordinate this information with other available information [103].

- *Flight rules:* The responsibilities and authority of the ‘pilot’ must be defined, as must operating rules for different classes of airspace, weather conditions, etc and any equipment that may be required onboard the aircraft.
- *Operator qualifications:* The licensing and training regimes for any pilots or crew need to be defined, together with any periodic activities required to maintain the currency of these qualifications.

Within any ‘responsibilities and authority of the pilot’ section of the legislation we must also address the degree of autonomy that we consider acceptable for any command and control procedures that might involve a UAV. In other words, we must establish procedural approaches for autonomously, semi-autonomously or directly controlling a UAV in any given phase of its mission or section of airspace. At present, many national bodies require that UAVs operating outside special restricted zones have either certified pilots at the controls or that the UAVs be commanded and controlled in a particular way (e.g. under semi-autonomous or direct ground control for take-off).

Another component of the command and control procedures concerns the use ofUCAVs. Like UCMVs, for at least half of each missionUCAVs will carry weapons, which significantly increases the potential threat to other airspace users as well as those on the ground. Consequently, specific procedures and legal arrangements will need to be developed regarding emergency and command and control procedures forUCAVs [161].

Small UAVs – comparable in size to model aircraft – seem to be regarded by most authorities as equating to model aircraft and unlikely to require integration with most civil airspace users. Indeed, similar constraints to those imposed on the operators of model aircraft are often applicable to small UAVs (e.g. height, line-of-sight, proximity to airports, etc). In effect, therefore, while there are many ways of categorising UAVs (e.g. weight, height, endurance, role, type, etc), and each is valid in context, from the perspective of operating in uncontrolled airspace the discriminator seems to be flight rules, as these are the most relevant to UAVs and manned aircraft. In other words, although it may be (say) kinetic energy, size, or endurance that defines whether a UAV is regulated in a particular way, from a legal perspective the categories will probably be defined on the basis of the applicable flight rules because they govern the regulations [103].

The implications of insisting on an equivalent level of safety for UAV airworthiness to that of manned aircraft is that civil certificates of airworthiness (or their military equivalents) may be required if the UAVs are to fly in uncontrolled airspace. Moreover, the UAV would need to be manufactured in appropriately certified facilities to agreed standards using approved materials and subject to rigorous inspection, maintenance and flight test processes [97]. Ground control, launch and recovery mechanisms, which are critical to UAV operations, would also need to be included in the above procedures.

Similarly, stringent security standards may need to be developed to protect the UAV command and control links so that they are resistant to jammers and

spoofers. As many UAVs will fly beyond the line of sight of their ground stations, these communications will need to accommodate satellite or other ‘multi-hop’ technologies.

The result is that many UAV facilities may also need to invest in an appropriate level of physical security and the cost of many UAVs will increase and quite possibly force many of today’s manufacturers out of the market; or make UAVs less cost-effective options to their manned counterparts.

6.1.6 Ground Vehicles

Even though the technology was immature, unmanned ground vehicles have been around since the late 1930’s, 1940’s, or possibly early 1970’s, depending on your point of view. More recently, and in particular since the 1990’s, there have been a series of successful UGV research programs and three DARPA Grand Challenges, which have fast-tracked the rate at which the technology has developed.

However, unlike maritime and air vehicles there are not a set of internationally agreed conventions or principles governing the interaction of ground vehicles with one another and/or other users of their environment. There are laws and conventions that transcend single states (e.g. Europe and the United States), but to all intents and purposes the principles by which ground vehicles are governed are developed and executed at a national or state level. As a result, Australian law is used here and readers can extrapolate for their own countries.

At present, under Australian law [120], the definition of a vehicle is very broad and not exhaustive. It includes, for example, motorised wheelchairs that can travel over 10km/hour, but does not include trains, wheeled recreational devices or wheeled toys. Riders of these are treated as pedestrians. Consequently, like UMVs, the legal status of UGVs must first be established. Moreover, the status and obligations of UGVs may vary according to their function and size, just as they vary for commercial and domestic vehicles, vehicles carrying passengers, vehicles of a certain size, and so on.

It is probably reasonable to assume that an autonomous UGV will be classified as a vehicle, although unlike tele-operated UGVs, which are really just an extension of their operator, autonomous vehicles may be considered separate legal entities. Moreover, most laws applying to the movement of vehicles apply to drivers or riders,⁹¹ as they are assumed to be in control of the vehicle. Consequently, attribution of control, and hence responsibility for any breach of the relevant legislation, is harder to ascertain unless we ascribe these laws to the IDT.

To this end, it is probable that we will need to establish the specific role of either the functional components of the UGV that result in its mobility or at a more holistic or system level, both of which may have far-reaching consequences. We will discuss these later in the section on UVS and Tort Law. However, an alternative to ascribing responsibility for the vehicle’s operation might be to declare that the vehicle is ‘not under control’ in the manner that we might classify a run-away vehicle. In some sense this seems appropriate because there is no human operator ‘in charge’ of the vehicle. However, from a practical perspective,

⁹¹ The distinction being when the vehicle is horse-drawn or a motorcycle, etc.

the run-away status would seem to imply that the vehicle is temporarily out of control, which is clearly inaccurate.

For UGV's the 'rules of the road' will not be the only laws that require modification. For instance, what if an autonomous vehicle was to trespass and/or cause property damage? Additionally, as autonomous UVS must have sensors for navigation purposes there may well be issues of privacy raised, just as with Google Maps and Street View [125]. Clearly, the existing law of tort would probably be used, but against who; the owner, the user, the developer/integrator, the component manufacturer? Clearly, as is compulsory for commercial and domestic ground vehicles, insurance might be purchased or insisted upon so that liability may be assumed vicariously by underwriters. However, this really just shifts and quarantines the liability it does not introduce novelty into the problem with respect to solving it.

6.2 UVS and Tort Law

Regardless of whether or not UVS carry weapons they have many safety implications as they are typically large objects that travel at speed and are constructed of sub-systems and materials that are inherently hazardous. Furthermore, most of these cannot be avoided as they are needed for system operation. For a weaponised UVS the list includes:

- Inadvertent firing of weapons;
- Engagement of unintended targets;
- System fratricide from weapon firing;
- Unintended system operation injures personnel;
- Exposure of personnel to hazardous chemicals;
- Exposure of personnel to damaging levels of radiation;
- Exposure of personnel to fatal electrical voltage;
- Damage to the environment; and,
- Collision with other people, vehicles or property.

Clearly, if a UVS is involved in an accident, the issue of how it was used is of interest. However, it may be extremely difficult to say whether this was inappropriate as often the best that can be concluded is that the operator used the UVS or its HMI in a certain way and certain consequences followed. This may be particularly true for complex environments, where the only data gathered and usable in regard to identifying the proximate cause may come from the UVS itself, which may of course be purposefully designed not to allow information to be disclosed that may 'incriminate' the developers, programmers, etc,⁹² or admit liability on behalf of any insurance underwriter.

⁹² If, as is suggested in later sections, UVS are provided with separate legal identity then intentionally not providing "self-incriminating" evidence is considered legally sound practice in many jurisdictions.

The key tort⁹³ law concept is *negligence*, which implies that the developer, owner, or user failed to do something that was required. Legal culpability for negligence depends upon a *failure to warn* or a *failure to take proper care* [24]. A *failure to warn* occurs when (say) the manufacturer was knowingly aware of a risk or danger but failed to notify customers of this risk. *Failure to take proper care* is more difficult to prove because it usually involves cases where the risk or danger cannot be shown to have been known to the manufacturer but, it is argued, was in some sense obvious or easily foreseeable, even if the manufacturer failed to recognise it.

The usual legal defence against charges of *failure to warn* or *failure to take proper care* is that the plaintiff acted in accordance with the ‘industry standard’ [17]. In other words, (say) the manufacturer followed ‘world’s best practice’, as measured against their peers, regardless of whether these standards are explicitly stated or are simply implied. Clearly, a first step for militaries around the world would be to establish a set of standards and principles for the development, operation and force-integration of their UVS, although these may need to be characterised in terms of parameters that are fiducially referenced. At present, however, most UVS are simply demonstrated in some representative military environment.

Another aspect of legal liability is that it can be differentially apportioned, even for a single event: that is one party might be 10% responsible, the other 90% responsible and analysis of causal links in product liability cases is common and will be an important consideration when considering UVS-related mishaps. For instance, a badly designed object recognition algorithm may be responsible for the prosecution of the wrong target, but so might a weak battery, a bad sensor, or a malfunctioning actuator [68]. Moreover, the environment in which the UVS is used or the training regime may also be the principal or contributing factors to the failure.

Fortunately, the law already has a highly developed set of principles and precedents that apply to product liability, so we can apply them to the use and production of UVS as if they were unexceptional technological objects. Moreover, it is believed that many of the concerns frequently touted as being possible dangers that UVS might pose will ultimately fall under this interpretation [24].

Furthermore, as indicated previously, insurance might be purchased for the UVS such that liability is assumed vicariously by an underwriter against a range of responsibilities and duties (e.g. the duty to exercise due care when operating in the presence of pedestrians or other vehicles). In fact, as is often the case, the technology may operate more correctly and provide greater reliability than its human counterpart and it may turn out that insurance for unmanned systems is less than that for the manned systems. The idea of insuring artificial agents is usually

⁹³ Tort (or civil) law deals primarily with property rights and infringements and seeks justice by compelling wrong-doers to compensate those who were harmed for their loss. The criminal law, which in our case includes the Law Of Armed Conflict (LOAC) seeks justice by punishing the wrong-doers. We will consider the LOAC and criminal law implications later.

attributed to Solum [255] and one of the principal objections used against it is the lack of any novelty that insurers provide. In other words, in the final analysis, someone has to pay the insurance premiums and it will probably be the owners, users, manufacturers, or programmers, who are ultimately therefore bearing all the risk anyway, albeit amortised over time.

An important principle in compensation is that of causation analysis [145], which seeks to attempt to fix liability on those responsible for the legal cause of an injury (the drunk driver was the cause of the car accident, not the brewery that made the beer). This is a necessary element of any civil lawsuit, although a number of other elements are also necessary (e.g. duty to the injured party and damage caused by the failure to execute that duty). Moreover, it is frequently difficult to pick the single (or few) legally responsible causes from the variety of possible ones. For example, most accidents result from multiple causes and it can be difficult to untangle the various contributing factors (a UVS that crashes into a house may be the fault of the manufacturers, the developers of the IDT that suffered a ‘glitch’, the ‘pilot’, the mission planners who allowed it to pass by or over the house, the air traffic controller (in the case of a UAV) who failed to notify the pilot of the thunder clouds, and so on).

Consequently, before any compensation could be paid a court would have to decide out of the many possible factors that were necessary for the accident to have occurred, which one(s) had actually caused the damage. To do this it would first seek Causation in Fact, which connects the injury suffered, through a chain of circumstances, to the act of an accused, no matter how many other contributing factors may be present [24]. This enquiry is a fundamental and intractable element of proof of a tort and is roughly equivalent to establishing the ‘but for ...’ causes of the accident. Next, in order to establish liability, Proximate Cause is determined, which selects from all of the causes in fact the entity (or entities) which will be held responsible for the injury.

Given the distributed, complex and polymorphic nature of the IDT likely to be responsible for controlling the UVS, to say nothing of possible interaction with its human supervisor, the courts and civil parties may find it extremely difficult to specify the proximate cause of the damage. In fact, they may not be able to even identify the UVS as the cause. Moreover, if we assume that automation is sophisticated (i.e. it is as competent at driving as a human) we are presented with a series of other questions. When should a human elect to use the IDT’s judgement and when his own? Should these decisions be based on a formal framework of metrics, the relative performances of users and the IDT, or less rational factors? What would the consequences be of inappropriately selecting one or other of these options if the same choice had previously resulted in a beneficial outcome? What might any mitigating circumstances be?

One solution might simply be to identify an errant behavioural trait or the presence of a fault within the UVS using a truncated form of proximate cause. That is, not take the detailed causal relationship into account, but accord

culpability to the UVS at the system or possibly sub-system level. Although this may seem unreasonable (and legally questionable), economic pressures, driven through either consumer choice or insurance policies may well determine the cost and hence the viability of key UVS components, thus providing a ‘market forces’ solution to any ‘fairness’ test.⁹⁴ It is understood, however, that as this solution may hold ‘innocents’ liable for actions for which they were not responsible, albeit vicariously through an insurance premium and this may not be politically or socially acceptable.

However, given the typically adversarial nature of law; that no general rules exist on how to identify the causes of UVS accidents and infringements; that in the abstract the law does not provide good answers to questions; that accidents will not usually be a function of the human operator, but of the system inaccurately or ineffectively facilitating user understanding of system functions; and that human error is not removed by simply increasing the level of automation, it may be that ascribing liability in this way is preferable to the alternatives.

The most practical way for this to become a reality is for UVS or their IDT to assume some sort of separate legal personality. While it is recognised that some people will dismiss the notion of a machine holding a separate legal identity as fantastical and not something deserving of serious consideration, this concept does not imply that the UVS is sentient: they are real-world machines controlled by algorithmic processes and are tools, not people. Furthermore, the practical instantiation of UVS holding separate legal status is unlikely to differ substantially from the vicarious liability assumed by an underwriter.

Before discussing the prospect of a silicon-based architecture rather than a carbon-based one having a legal personality, however, issues of accountability and liability relative to the Laws of Armed Conflict (LOAC) are discussed. In this domain inappropriate actions may be criminal, any redress may need to be non-monetary, and insurance is more than likely unsuitable.

6.3 UVS and the Law of Armed Conflict

To date the IDT onboard a UVS has largely been used for navigation purposes. However, it can also be used for the autonomous identification and engagement of targets. This means we can now provide UVS with the potential to calculate which objects within a sensor’s field of view constitute targets, whether or not to engage them, and by which route or means they should be engaged. If the payload includes a weapon, then a UVS has the potential to make the final determination on whether to prosecute a target with lethal force, without operator intervention. If permitted, this would mark a sea-change in the role of technology in warfare as the human could be removed from the decision-making loop, with potentially lethal consequences. As soon as such weapons are justified by one nation, others will likely follow. The use, development or acquisition of such weapons may then

⁹⁴ In this regard, it is also assumed that there would be threshold tests relating to health and safety applied to avoid a technological ‘free-for-all’.

trigger our obligations under Article 36 of Additional Protocol 1 to the Geneva Conventions.⁹⁵

Even though a significant number of influential writers [47] [56] [92] [159] [195] have argued that future IDT will possess intelligence equal to – and probably in excess of – humans, particularly in stressful and high workload environments, at present most would agree that there are advantages of having real ‘eyes on the target’. Moreover, many commentators have suggested that “Humans will always make the final decision” [246] [256] [297] as, under the Laws of Armed Conflict (LOAC), there is a requirement to assess ambiguous situations with rational human judgement such that the response is not excessive to the anticipated military advantage. For example, in the words of Air Chief Marshall, Sir Brian Burridge [62]:

“When we go into combat, we have got to be sure what we are doing is legal and moral. I do not believe that, in future, even though technology will allow it, we will be allowed to indulge in robotic warfare. I simply do not see the international community regarding that as an appropriate way to fight. The notion of using UCAVs controlled ten time zones away to prosecute a battle is not something international law of the future will regard as acceptable. I think the notion of a person in the loop, the notion of positive ID, the notion of someone feeling the texture of what is going on in the battlespace, is going to be more and more prevalent. ... Overall, I think robotic warfare drives you away from what I term as emotional connectivity with the battlespace. My view is that winning the hearts and minds battle with the indigenous population requires this emotional connectivity.”

Unfortunately, retaining humans in this way is likely to provide its own problems as the issue is not so much “Will there always be a human in the loop?” but “Where will the human be in the loop?” Will it be at the level of each and every target engagement, or at the level “use of lethal force in this mission is granted?” Moreover, there is only modest evidence to suggest that autonomous or UVS-based weapons will be considered any differently from other weapons systems.

For example, IDT that permit lethal response options may be introduced as systems that act if human input/veto has not occurred within a certain period. As IDT improves and is introduced onto the battlefield in greater abundance, critical survival decisions will start to depend upon them. As a result, operational tempo

⁹⁵ Protocol Additional to the Geneva Conventions 12 August 1949, and relating to the Protection of Victims of International Armed Conflicts (Protocol 1) - adopted on 8 June 1977 by the Diplomatic Conference on the Reaffirmation and Development of International Humanitarian Law applicable in Armed Conflicts. **Article 36, New Weapons:** In the study, development, acquisition or adoption of a new weapon, means or method of warfare, a High Contracting Party is under an obligation to determine whether its employment would, in some or all circumstances, be prohibited by this Protocol or by any other rule of international law applicable to the High Contracting Party [217].

may well be diminished if the systems rely upon human intervention, particularly for complex or networked environments. Similarly, as information is increasingly processed and manipulated by these systems prior to decisions being made or confirmed by humans the link between *real* decision-making and automated response will become blurred.

Furthermore, even though active vision and intelligent sensing techniques are not yet as sophisticated at interpreting a situation as humans, automation is almost universally welcomed by operators because, if introduced correctly, it allows reliability and the overall system effectiveness to be improved, allowing humans to concentrate on other tasks. For instance, an automated vision system has the advantage of being able to ‘stare’ equally effectively in all areas of the sensor’s field of view. The human fovea is only about 0.5 deg, degrading to one tenth of the visual acuity at 2 deg [28]. This provides the human with a remarkably narrow angular focal view. Moreover, the human suffers from fatigue and is unable to focus on (or even detect) multiple targets within a scene. An automated system can alert the user to the presence of multiple potential targets and overlay them on the video – perhaps in a head up display to the operator. The user can then use his superior interpretative skills to verify or dismiss the targets.

Prior to the development of technologies that are able to meet the implied challenge of full autonomy, therefore, it is likely that users will retain responsibility for high-level tasks such as target designation, target verification, and the lethal application of force. Similarly, the intelligence onboard the UVS is likely to assist in the process by identifying and tracking potential targets and controlling and aiming the weapon system. Maintenance of broader situational awareness will probably be a shared and less controversial responsibility. Moreover, as the human reaction time is around 400ms, whereas a machine can respond in a few nanoseconds, the first use of fully autonomous UVS weapons will probably be in defensive roles, where critical decisions of survivability must be made quickly, or in the prosecution of fleeting targets, where the timeline is also necessarily short. Hence, it is very likely that the target verification processes will be compressed and highly dependent upon IDT, perhaps to the point where the only human input is oversight or veto.

Unfortunately, in order to achieve the appropriate level of human oversight, we will still be reliant upon technology for the users and officers superior to those engaged in combat operations to remotely watch (and potentially interfere with) a tactical engagement. This raises a number of issues:

- There is the question as to whether the senior commanders are legally obligated to interfere with ongoing operations if they realise that a violation of LOAC is about to occur. Schmitt [246] states that application of Article 28 of the Statute of the International Criminal Court indicates that senior commanders will progressively become more responsible for their subordinates’ conduct.
- A potential ‘hiatus’ in tempo may arise when those on the frontlines start to rely on the intervention of their superiors rather than using their own

discretion before proceeding with a particular course of action. It seems likely that, to avoid such interruptions in tempo, there will be pressure for IDT to be introduced in support of the decision-making process, for instance through the filtering or fusing of information, the generation of response option, plans, schedules, and even the control of weapons.

- It is an open question as to whether senior commanders have access to technology that allows them to observe the relevant aspects of a remote operation to the degree necessary for them to usefully intervene. Given the amount of potential information that such a commander is likely to have to deal with (and require for any lethal target prosecution), and the compressed timeline against which such decisions are likely to need to be made (otherwise he would be using a more conventional weapon), it would seem probable that the only way to meaningfully monitor and respond to any such workload would be through the use of IDT.
- Just as corporations strive to exploit research into human factors and human systems integration to develop interfaces and artificial agents that garner trust in their users, so manufacturers of IDT and UVS will attempt to do the same. In other words, IDT will progressively undertake the difficult work (data fusion, visualisation, etc) so that the decisions made by humans become simpler. The net effect of this ‘sleight of hand’ will be to diminish the possibility of users making truly informed choices.
- Once this occurs, it is a small step to allow UVS to make decisions relating to the use of lethal force may in response to particular defensive circumstances. Clearly, the human commander has made the decision to use the UVS in the first place, which might be considered the crucial discretionary judgement. However, if the circumstances become more complex than initially anticipated or the UVS acts autonomously for some other reason (it perceives an attack to be taking place) there may be an ambiguity in regard to who actually made the decision to engage.

This effectively brings us back where we started, which may mean that the decision on whether or when to use lethal force is handed to IDT sooner rather than later.

6.3.1 The Law of Armed Conflict

The Law of Armed Conflict (LOAC) governs the protection of non-combatants and aims to limit the effect of war on those not directly involved. It sets out a legal framework for commanders in the field and defines the method and means by which warfare can legitimately be employed by state actors. The body of law that underpins the LOAC has been extensively codified in a series of international treaties dating from the 1860’s to the present day. The most significant of these treaties are the 1907 Hague series of treaties, which governs the conduct of operations, and the 1949 Geneva Conventions [217], which governs the protection of people and property.

To some extent, the two branches have now merged under Additional Protocol I [217], which was added in to the Geneva Conventions in 1977. As signatories to the Geneva Conventions and its Additional Protocols, countries are obliged to honour their obligations in good faith. Furthermore, many countries also incorporate the provisions of these treaties into their domestic law.⁹⁶ Clearly, countries must understand their obligations under these laws to ensure that they do not develop, acquire, or deploy systems that contravene the spirit or the letter of their responsibilities.

The theoretic framework for debate about the morality of specific choices and actions in war is usually attributed to Walzer [288] who established a set of principles that effectively capture general moral principles. The work draws a key distinction between just reasons for going to war (*jus ad bellum*) and just acts in fighting war (*jus in bello*). In this work we deal almost exclusively with the latter, although brief reference to the potential impact of UVS on the principle of *jus ad bellum* is made in the next section. A brief description of Walzer's principles, drawn from [29] now follows.

The **Principle of Military Necessity** states that "a combatant is justified in using those measures, not forbidden by international law, which are indispensable for securing complete submission of an enemy at the soonest moment" [288]. Military necessity requires combat forces to engage in only those acts necessary to accomplish a legitimate military objective. It permits the killing of enemy combatants and other persons whose death is unavoidable and it permits the destruction of property if that destruction is imperatively demanded by the necessities of war. However, destruction of property as an end in itself is a violation of international law as there must be a reasonable connection between the destruction of property and the overcoming of enemy forces.

The **Principle of Humanity**⁹⁷ forbids the "use of means or methods of warfare which are calculated to cause suffering which is excessive to the circumstances" [29]. It has also been expressed as averting the infliction of suffering, injury or destruction not actually necessary for the accomplishment of legitimate military objectives.

The **Principle of Proportionality** provides a link between the principles of Military Necessity and Humanity. In simple terms, the principle generally relates to the reduction of incidental injuries caused by military operations and requires that the losses and damage resulting from military action should be proportionate (i.e. not be excessive) in relation to the anticipated military advantage. The proportionality principle, together with the principle of humanity, dictates that "civilians should not be made the object of attack and that while civilian casualties may be an inevitable consequence of an attack, every effort must be made to spare them, and other parties who are non-combatants, from becoming adversely affected" [288]. The principle of proportionality not only requires that an attacker must assess what feasible precautions must be taken to minimize incidental loss,

⁹⁶ In the case of Australia these are the Geneva Conventions Acts of 1957 and 1991 [29].

⁹⁷ This principle is codified under Article 35(2) of AP-I: It is prohibited to employ weapons, projectiles and material and methods of warfare of a nature to cause superfluous injury or unnecessary suffering.

but must also make a comparison between different methods or axes of attack so as to be able to choose the least excessively destructive method or axis compatible with military success. When making that assessment, the attacker should “naturally take into account likely friendly casualties” [29].

The related **Principle of Distinction**⁹⁸ seeks to ensure that only legitimate military objects are attacked. This principle has two components: the first, relating to personnel, seeks to maintain the distinction between combatants and non-combatants or civilian and military personnel; the second distinguishes between legitimate military targets and civilian objects. Military operations must only be conducted against military objectives, including combatants. Non-combatants and civilian objects are protected from attack, that is, they are not legitimate objects of attack. LOAC therefore requires that “belligerents maintain the clear distinction between armed forces and civilians taking no direct part in hostilities; that is, between combatants and non-combatants, and between objects that might legitimately be attacked and those protected from attack” [29].

Together these principles seek to temper the violence and range of war. The **Principle of Responsibility** seeks to tie in the actions of warfighters to morality more generally, by seeking to ensure that the agents of war are held accountable for their actions. For instance, when soldiers turn their weapons against non-combatants, or pursue their enemy beyond what is reasonable, they are no longer committing legitimate acts of war but acts of murder. Other issues that arise from this principle include the morality of obeying orders that are known to be immoral and the status of ignorance (i.e. not knowing the effects of one’s actions). In other words, there is at least an aspiration to identify those responsible for deaths in war, even if practical circumstances do not always permit this.

While there does not appear to be a universally agreed definition of “armed conflict” (the point at which LOAC applies) and [246] argues that, “given the means and methods of modern warfare it is no longer sufficient to use a threshold test of actor-based physical confrontation,” (almost) all nation-states now accept the basic Principles underpinning the LOAC and that we are legally bound to consider these principles together.

There is not a uniform interpretation of all aspects of LOAC.⁹⁹ Moreover, some principles are only codified in Additional Protocol I, which has not yet been ratified by a number of major nations (notably the US) [297]. Consequently, in order to deploy a system that relies upon an IDT designed and/or manufactured by another nation, a country which has ratified the protocol (e.g. Australia) may have to understand the manner in which the vendor nation interprets the LOAC and whether this accords with its own interpretation. This is a further complication and

⁹⁸ This principle is codified under Article 48 of AP-I and is complemented by a number of other provisions within the Protocol that relate to the consequences of distinction, including Article 51(2), which prohibits attacks on civilians and Article 52, which defines military objectives.

⁹⁹ For instance, a civilian loses the protection of the LOAC “If, and for such time as, he takes a direct part in hostilities” (Article 51 (3) Additional Protocol I). The interpretation of the phrase “Takes a direct part in hostilities” is controversial, with some nations employing a higher threshold test than others.

of particular significance to countries that acquire significant military capability from nations such as the US, as the Statute of the International Criminal Court (ICC), which (for example) Australia has signed and (for example) the US has not, provides that the responsibility for providing the legal means and methods of warfare falls upon the military using the technology [297]. In other words, if the IDT violates the LOAC the forces of the purchasing nation may be held criminally liable not the vendors of the IDT or their national government.

The key message from Principles of Necessity, Humanity and Proportionality is that any weapon developed, acquired, or deployed must not employ any degree of force that is not necessary for the purposes of war. Somewhat obviously, therefore, we may not employ an unlawful weapon on a UVS and abrogate our responsibilities in this regard by implying that the UVS ‘made the decision’ to use the weapon. Similarly, the Principle of Distinction requires that any deployed system must:

- Distinguish between civilians and combatants;
- Distinguish between civilian and military objectives;
- Direct operations only against military objectives; and,
- Not cause excessive incidental loss to civilians.

It remains an open question as to whether or not IDT will ever be able to successfully identify legitimate military targets all of the time; it challenges humans. Nevertheless, given the rate of technological progression and the level of investment made by militaries around the world into research in this area it seems inevitable that IDT will progress to the point where they could be put to work on aiding and/or making decisions about what objects constitute targets and by what means and manner these targets can be engaged. To this end, let us examine whether UVS might ever be sufficiently competent to undertake such duties.

In order to be of use the division of roles between the human and the UVS will primarily be driven by a combination of human and technological factors, including the cognitive load on any human operator, the degree of trustworthy “intelligence” and automation achievable within the UVS and its weapon, and the potential impact of latencies on the control and actuation of any weapon systems. Based on our deductions from the Principle of Distinction above, these may be expressed as a series of technical requirements for the UVS. It must:

- Positively and accurately identify objects of interest and intended targets;
- Determine and minimise any effects on non-combatants and infrastructure;
- Accurately control the delivery of the weapon wrt to aiming at any targets;
- Maintain a suitable record of action for the purposes of evidentiary hearings;
- Contain no known bugs or malfunctions that might cause ‘reckless’ behaviour.

In addition to the basic principles of LOAC, which determines the actions that are lawful and therefore permitted, governments then place further limitations on military forces in respect to the application of force (for operational, political, diplomatic, or legal reasons). These conditions are known as Rules of Engagement (ROE). ROE provide authoritative guidance for the application of force in operations and take two forms: actions a soldier may take without consulting a higher authority, unless explicitly forbidden, and actions that may only be taken if explicitly ordered by a higher authority [29].

To comply with any ROE the UVS must also accept additional restrictions in regard to the circumstances under which it may use discretionary lethal force. For instance, it may be necessary to vary the type of targets allowed, the nature or trajectory of a UVS mission, the degree of hostility used to trigger a response and the nature of any response that might be permitted to achieve the desired goal (e.g. a warning shot may be required). Alternatively, we must provide our commanders with the requisite level of understanding of IDT capabilities relative to these functions so that they may take the appropriate precautions prior to and during the execution of their duties.

If for the moment we set aside the complexities of interpreting specific ROE, with the exception of “positively and accurately identifying objects of interest and targets,”¹⁰⁰ technology is currently available to undertake each of the other functions. Consequently, it is not difficult to imagine a system that combines these functions in order to automate the tasks currently undertaken by a human (possibly in conjunction with an automated system), particularly for simple scenarios where complex judgements and ROE do not need to be taken into account (i.e. scenarios that give a UVS very little discretion). To this end, let us now try to imagine how the introduction of such a capability might evolve.

Stage One - A decision-aid is developed that assists the warfighter, allowing him to manage a larger number of tasks or targets. Such a program might allow the human to make goal-setting decisions with respect to UVS missions and navigation, high-level decisions regarding specific target designation and verification and compliance with the ROE. Other tasks, such as target acquisition and tracking, weapon aiming and control, record keeping, HUMS and weapons impact calculations might also be undertaken autonomously. In other words, the UVS would effectively act as humans do today, but would not take the key decisions on target discrimination, designation and verification – humans would.

Stage Two - A UVS is developed that outperforms a human: ROE are then issued to the effect that – aside from exceptional circumstances – warfighters must follow the advice of the IDT. Perhaps experience has shown that under duress or in high workload environments humans make more targeting errors. Regardless, there is now little or no reason for the humans to check the program for compliance. As a result, the role of the human in the decision-making process diminishes and the number of UVS and targets and tasks he can manage become considerable; he

¹⁰⁰ And under certain circumstances this function can also be accomplished using current technology.

devotes little or no time to any particular target. There are still times when a human is called upon to make a decision. If such events occur with regularity, however, the developers of the system will be tasked with developing a routine for handling these. As the capabilities of such a system grow, the need for the human to make decisions diminishes.

Stage Three - This begins when the decision is made to do away with the human altogether; perhaps the human input/over-ride causes delays relative to an adversary's capabilities and these have a tendency to result in casualties.

The Principle of Responsibility then requires us to determine who is accountable for any erroneous or 'criminal' actions. This will be the focus of our discussion from here on. First, however, let us consider the potential impact that UVS might have on *jus ad bello* (just reasons for going to war) as UVS might also influence a nation's entry into war.

6.3.2 UVS and *Jus ad Bellum*

There are a few ways in which UVS might affect *jus ad bello* [23][25][250]:

- They could directly threaten the sovereignty of a nation;
- They may introduce ambiguity into who made the decision to attack;
- They might make it easier for leaders who wish to start a war to actually start one;
- They may alter perceptions and ultimately make it harder to win a war;
- They may be used to deliberately carry out war crimes.

Challenge to Sovereignty – Autonomous UVS operating near political hotspots offer considerable potential for starting wars accidentally – or by more nefarious means (e.g. as a result of human manipulation). For instance, aUCAV on a surveillance mission might respond lethally to an Early Warning radar that is legitimately (and passively) protecting its host nation's airspace. Although not strictly a threat to either nation's sovereignty, this may be interpreted as an 'act of war' that leads to a more extensive conflict. Clearly, wars that start as a result of an accident do not fit the principle of *jus ad bellum* [25].

Ambiguity in Attack Decisions – Autonomous UVS with the capacity to make decisions relating to the use of lethal force may have been granted authority to respond to particular circumstances (see the above scenario). Clearly, the human commander has made the decision to use the UVS in the first place, which might be considered the crucial discretionary judgement, and which implies that all subsequent discretionary decisions derive from the principal authority. However, if the circumstances are more complex than initially anticipated or the UVS acts autonomously for some other reason (accidents are covered above) there may be an ambiguity in regard to who made the decision to engage in the first acts of war.

Once again, the commencement of such wars would not fit the principle of *jus ad bellum*, particularly if no formal declaration has been made.

War Crimes – Because the autonomous UVS cannot themselves be held accountable for their actions, some nations or commanders may elect to use them to deliberately conduct activities that constitute war crimes, thereby potentially absolving it/themself of any responsibility.

Lowering the Barrier to War – A central element in the decision-making process for most democratic nations is the estimation of the cost of a war in terms of lives lost, even if these are the lives of soldiers who have volunteered [25]. Consequently, in these nations a political strategy has evolved to limit the number of casualties such that, where possible, military involvement is kept to a ‘safe’ form. Moreover, significant investment is now made in a range of technologies that lower the risk to and increase the lethality of warfighters (e.g. armour, longer range weapons). This sometimes translates into civilian casualties. To this end, it is clear that the introduction of technologies such as UVS, which are specifically aimed at extending the reach of the warfighter to provide him with greater lethality and more protection, could lower the political barriers of entry into war by offering to ‘take the blood out of war’¹⁰¹ [252].

Perceptions – Autonomous UVS may be perceived as providing those militaries using them as having an ‘unfair’ advantage and, while not illegal under the LOAC, this perception might be used against (presumably Western) military intervention or even to incite resistance among the local population. Moreover, the media is a powerful weapon in today’s warfare and if UVS were to kill innocents this might lead to widespread revulsion. Finally, as coalition forces have discovered recently in the Middle East, there is no substitute for soldiers on the ground that have the capacity to interact with the local civilian population. Ultimately, therefore, wars waged with autonomous UVS may make it harder to win and establish a lasting peace.¹⁰²

Another point made eloquently by [252] is that the vicarious engagement of warfighters in combat through UVS controlled from thousands of miles away effectively combines the previous two points. The net effect on our potential adversaries is to make them believe that we are cowards, thereby raising the likelihood that they will engage us in or attempt to prolong war.

A final note on this score; if we assume that only nations that fight just wars were to develop UVS technologies none of the above would be a concern. Unfortunately, history shows us that “all wars involve at least one unjust (or badly mistaken) nation” [25] so the prospect that UVS will be used to enable such future injustices is a legitimate concern.

¹⁰¹ Naturally, this is not meant to imply that we should increase the risk to our forces and reduce their lethality in order to improve the Just War considerations. The point simply notes that this is one potential impact of UVS on *jus ad bellum*.

¹⁰² This is perhaps more pertinent to *jus post bellum* (justice after war).

6.3.3 *Autonomous Weapons*

A full account of what is meant by autonomy and whether UVS can actually possess it requires answers to questions that philosophers have wrestled with for centuries. Regardless of its level of autonomy, however, to some greater or lesser extent a UVS is reliant upon a combination of external human-supervisory commands and onboard sensors and processing to capture, represent, and interpret environmental cues that are then autonomously combined and manipulated (and possibly presented to the user) such that the result is a series of mobility and payload response commands representing mission priorities.

In current parlance, the phrase “autonomous weapon” typically means that the weapon is a ‘fire-and-forget’ system capable of acting independently of human control, but only to the extent that it can determine mid-course and final trajectories and acquire targets on the basis of certain limited stimuli. There are several such weapons in existence today, some of which use of over-the-horizon and beyond-visual-range systems that include Automatic Target Recognition (ATR) capabilities that allow the weapons to select targets when they enter a designated area of interest. These features preclude full human control.

However, these weapons are really only semi-autonomous, in that they wait for their programming to be satisfied. That is, the determination of location, value and risks are undertaken by humans who ultimately control the weapon. For instance, some weapons must have their target sets pre-programmed before they are launched, although they can be re-programmed in-flight if real time intelligence indicates that the target has already been destroyed or if the situation has changed (e.g. civilian traffic has entered the area). Similarly, HARPY¹⁰³ can loiter for several hours before detecting, locking onto, and then destroying enemy radars.

Moreover, the SGR-A1¹⁰⁴ can detect and identify targets within a 4km radius (or 2km at night) and provide either a lethal or non-lethal response. Although the manufacturer indicates that the ultimate decision “should be made by a human not a robot” the system has an automatic mode. Similarly, Israel also has robotically

¹⁰³ HARPY is a fire-and-forget, UAV-based weapon system that can be launched from a ground vehicle or from a ship. It is designed to loiter for extended periods during which time it can detect and then destroy radar emitters. No operator input is required, although target verification can be insisted upon: the target radar “lights up,” the UAV’s onboard sensors acquire and compare the radar signal to a library of hostile emitters, prioritise the threat and (if verification is required/given) the drone enters its attack mode. If the radar is turned off before HARPY strikes, the UAV can abort its attack and continue loitering. If no radar is spotted during the mission, the UAV is programmed to self destruct over a designated area. The latest versions allow visual ID and attack of targets, even after they turn off their emitters.

¹⁰⁴ The SGR-A1 “robot” was jointly developed by the Korea University and Samsung Techwin Corporation. It was deployed along the South Korean border of the Demilitarised Zone (DMZ) in 2007 and uses a combination of electro-optic and infrared sensors to detect and track mobile targets and to autonomously control, aim and fire a machine gun. The SGR-A1 is also reputed to provide verbal commands to its targets (in Korean) to surrender; apparently recognising the action of holding one’s arms high in the air as the act of surrender, which prevents or suspends the automatic firing response.

controlled fifty calibre machine guns along its border with Gaza. The current policy is that “in the initial phases of deployment” there will be a man in the loop; leaving open the prospect of more autonomous operations in the future. There are a wide range of other examples of semi-autonomous weapons and weaponised UVS that include the Phalanx anti-aircraft/missile weapon¹⁰⁵, the Tomahawk anti-ship cruise missiles¹⁰⁶, the anti-submarine Captor mine¹⁰⁷ the TALON SWORDS, GATERS, Fire Ant and MAARS UGVs¹⁰⁸, and the Patriot III missile batteries¹⁰⁹ [23] [64].

At present, it is also possible for many of these systems to be re-tasked by someone from a different command chain to the person who originally initiated its use. Moreover, the person who first used the weapon may not be aware of the new mission or of the rationale behind the change. Nevertheless, under the LOAC responsibility for the weapon lies with the person who actually uses it: if he has doubts over the outcome he should not launch it. In reality, therefore, past a

¹⁰⁵ The MK-15 Phalanx Close-In Weapons System (CIWS, pronounced “sea whiz”) is an automatic, fast reaction, rapid fire 20mm machine gun that provides US Navy ships with terminal defence against attacks that have penetrated other fleet defences. It is designed to engage anti-ship cruise missiles and fixed wing aircraft at short range without assistance from other shipboard systems. It automatically engages incoming anti-ship missiles and high-speed, low-level aircraft that have penetrated the ship’s primary defense envelope. It performs search, detection, tracking, threat evaluation, firing, and kill-assessments of targets autonomously, while providing for manual override.

¹⁰⁶ Tomahawk cruise missiles are launched in the general direction of their intended target. At some distance from the anticipated target location the missile enters a serpentine trajectory (several other patterns are also possible) to search for the target using passive and active radar techniques. Once the target is detected and the onboard algorithms satisfied, the missile locks onto the target and automatically enters its attack phase.

¹⁰⁷ The MK-60 CAPTOR mine is often referred to as the “mousetrap that chases the mouse.” It uses onboard acoustic sensors to detect and classify hostile submarines, while ignoring surface ships and friendly submarines. Upon acquisition of a suitable target it launches a modified Mark 46 torpedo, which searches and homes in on its intended target using a circular search strategy.

¹⁰⁸ The TALON SWORDS UGV, developed by Foster-Miller & Qinetiq, is a tracked, semi-portable vehicle that carries an M240, M249 machine gun or a Barrett 0.5 calibre rifle. The system has been operationally deployed. The Modular Advanced Autonomous Robot System (MAARS) is a replacement for the TALON SWORDS and can carry a 40mm grenade launcher or an M240B machine gun. The GATERS (Ground-Air Tele-Robotic System) is a somewhat larger UGV from the 1980’s that carried hellfire missiles. Fire Ant (manufactured at the Sandia National laboratory) is also of 1980’s vintage and while tele-operated from a navigation stand-point it had an autonomous weapon firing mechanism; the GATERS sensor package was aimed by an operator, but the image processing algorithm autonomously fired the weapon in response to target motion (tanks).

¹⁰⁹ The Patriot missile system uses ground-based radar to detect, identify and track its targets. An incoming missile could be 80km away when the Patriot’s radar locks onto it and consequently not visible to a human being, much less identifiable. The Patriot can also operate in a completely autonomous mode with no human intervention as an incoming missile flying at Mach 5 (i.e. traveling at about one mile/sec) does not leave much time for a human to respond, making automatic detection and launch an important feature.

certain point these weapons are no more controlled by humans than artillery or mortar barrages¹¹⁰ and once the targeting package onboard the weapon is satisfied the target is destroyed. Consequently, at one level, these semi-autonomous weapons would appear similar to autonomous UVS with weapons, except that the IDT onboard the UVS is re-tasking the weapon ‘on the fly’ instead of a human from a different command chain.

However, the use of fire-and-forget weapons is only considered lawful if they are “equipped with onboard sensors or have access to external sources of targeting data sufficient to ensure effective target discrimination” [71]. Moreover, commanders do not launch these weapons on a purely speculative basis. That is, while there may not be a priori knowledge of the precise location or immediate accessibility of a given target prior to launch, there is usually some intelligence about the target (e.g. it is believed to exist, to be within a specified Area Of Interest (AOI), to be of a particular type, pose a particular threat, etc.). Additionally, the weapons may be safely destroyed if suitable targets are not acquired.

UVS, on the other hand, are often used in intelligence-gathering, surveillance and reconnaissance (ISR) roles, where there is no a priori knowledge of the existence or nature of targets. Given our lack of knowledge regarding the potential diversity, nature and location of any targets that might ‘pop up’ it is also unlikely that we will have a good understanding of the surroundings and hence the potential effect of weapons on non-combatants and infrastructure, except in specific circumstances (e.g. uncluttered environments such as air-air or underwater engagements).

Another way to describe the difference between the current and future technologies is that the extant weapons represent ‘conduit warfare’ whereas fully autonomous UVS and their weapons represent ‘intermediary warfare.’ That is, fire-and-forget weapons represent technology conduits through which warfare is conducted by extending the reach of the protagonists. However, they do not fundamentally alter the terms or conditions of engagement. Newer technologies, on the other hand, promise to remove some of the key elements of the decision-making process in the target prosecution process from the protagonists. Clearly, regardless of the degree of sophistication underlying the weapon’s autonomy, there is a need to safeguard against unauthorised persons gaining control of the UVS or its weapons. In fact, surety of such control will likely be one of the underlying tenets of any such fielded system.

6.4 Accountability and Liability

In an earlier section we briefly discussed some of the issues of liability and accountability from the perspective of tort law. Unfortunately, if UVS are to be used operationally and the Principle of Responsibility is to be upheld, who is responsible for breaches of the LOAC involving an autonomous weaponised UVS? Although somewhat far-fetched, the following (based on [256]) will illustrate the matter.

¹¹⁰ Which cannot be surrendered to and for which there are few moral or legal concerns.

Let us assume that a weaponised UGV has been tasked with some high priority goal. En-route to this objective it encounters adversaries, who surrender. The UGV recognises their surrender and stops to guard them, but also determines that the imposed delay will prevent it from achieving its primary objective. The UGV also determines that if it were to stop guarding the prisoners they would in all probability interfere with (and therefore prevent) it achieving its primary and higher priority goal: it kills the surrendered adversaries and proceeds to its higher objective.

This is clearly a war crime, but who is responsible? Moreover, as the Principle of Responsibility requires agents of war be held accountable for their actions, the argument may reasonably be raised that this alone precludes UVS from making such decisions or taking part in such scenarios. To some even the intuitive implausibility of having to hold a UVS to account will be sufficient to preclude their ever being allowed to make the final decision regarding the lethal use of force. There are a number of issues.

Providing UVS with separate legal personality and holding it ‘responsible for its actions’ does not imply that it has sentience, intentionality or will. Furthermore, we must consider reasons why humans commit war crimes: hatred, bigotry, racism, malevolence, and so on; all reasons that are characteristically intentional and not something that we expect to ascribe to a UVS. Unintentional reasons (i.e. a “failure to take proper care”) really relate to levels of technical competence in the UVS, the users, manufacturers, etc and these are covered in tort law arguments. Interestingly, if a UVS had a moral sense the law would have something to act upon. Even so, infants have little or no moral sense, but are accorded legal personality and hence an implied level of responsibility. Similarly intelligent animals (e.g. dogs) are assumed to know their master’s bidding [68] and may be punished for their disobedience.

On this matter, [256] helps us conceptualise the key problems of attributing responsibility to artificial agents by comparing them to child soldiers. Clearly, children have a degree of autonomy and are capable of a wide range of decisions and actions (significantly more so than any existing UVS, in fact). Like UVS, the children will *probably* kill the right person but may not. Regardless, it is widely recognised that, as the children do not fully comprehend the ethical dimensions of what they do, they are not the appropriate objects of punishment. That said, our moral repugnance at the use of child soldiers is perhaps more at the use of children per se rather than at the use of agents of war that have potentially unreliable or unpredictable target discrimination capabilities and a diminished sense of moral responsibility.

As indicated in the section on tort law the legal system ascribes liability by tracing the vector of causation back to the human agency where mistakes were made. The sins of omission and commission are then scrutinised with respect to negligence, recklessness, intentional malfeasance and so on [145]. To understand these issues in greater detail let us follow the approach of [256] and expand upon the collective ideas put forward by [104] [181] [246] [252] [256] and [297] and others and try to attribute responsibility for any accidents or infringements.

An Unfortunate Mistake

The allocation of responsibility is important for the principle of *jus in bello*. Sometimes accidents will happen, but if it is typically impossible to hold individuals responsible for the casualties of war then this violates the Principle of Responsibility. To this end, if the UVS does not over time demonstrate sufficient and predictable capacity to discriminate targets correctly, issues of culpability will almost certainly be raised. Consequently, while it is acceptable to ‘apologise and move on’ for the accidental destruction of a civilian target, some nations will probably use such an approach for target misidentifications made by the automatic target recognition algorithms (say, when a man with an axe is mistaken for a man with a rifle). However, this strategy is not likely to be acceptable when the violations become indiscriminate or numerous, or when the breaches of LOAC are more complex. Other candidates will be sought.

The User/Commander

The Statute of the International Criminal Court (ICC) provides that the responsibility for providing the legal means and methods of warfare falls upon the military using the technology [246]. To this end, the officer who ordered the deployment of the weapon is held responsible for the consequences of its use. In other words, he accepts the risk that the weapon may go awry when he makes the decision to send it into action, even if the consequences were not intended. Clearly, suitable precautions must be taken by the users prior to any mission.

However, as [256] points out, if we adopt this precedent as the norm for fully autonomous UVS, we neglect that current weapon systems do not autonomously select their own targets, except within the framework of a fire-and-forget weapon. Fully autonomous UVS and weapons, which attack targets that have not been chosen for them, have an inherent ‘unpredictability’ that must relieve the user from his responsibility – at least to some extent. In other words, the autonomy of the UVS implies that the orders issued by the supervisor influence, but do not prescribe, its actions as the decision-making programs are operating in the real world and make decisions unforeseen by the human operators at time of deployment.

The Acquisition Organisation

Clearly, prior to any acquisition and in-service deployment of an IDT-controlled weapon extensive testing and evaluation will be needed. For example, in order to accept liability the agency must quantify the anticipated performance and reliability of the IDT. To do this it will need a series of metrics, although historically, most of the results associated with artificial intelligence, robotics, UVS, and autonomous systems have been in the form of demonstrations. A trail of documentation would also need to be carefully and accurately recorded as there may be a need to record decisions made in the field for the purposes of evidentiary hearings, court proceedings, improvements, upgrades, etc. This is the case with all modern weapons systems, most of which now crucially depend upon information coming from multiple sensors, filtered and fused using IDT and provided to a human-machine interface (HMI) for a user to make a decision.

The acceptance of liability by governments for sophisticated equipment that can behave in a ‘somewhat unpredictable’ manner is not a novel concept. However, the IDT for an autonomous UVS may be so complex as to make reproducible testing very difficult, alternatively it may have the capacity to learn from its mistakes. Either way, the more intelligent the decision-making process, the harder it is likely to be to functionally establish or fully test the response of the system to repeatable or verifiable system stimuli. How then can an acquisition agency be held responsible for the actions of an IDT that it cannot fully test? Similarly, how can this agency accept liability on behalf of its nation or the likely users?

*The Developer*¹¹¹

Given the complications articulated above it is tempting to consider holding those who designed, manufactured, or programmed the IDT responsible for its actions. However, inherent in the nature of an autonomous UVS is that its IDT will make choices other than those directly programmed into it. For instance, when immersed in the real environment, it is highly unlikely that the IDT’s decision-making will be entirely predictable. This is not to say that it will be random, but that the internal states of the machine are likely to be so complex that the actions of the machines may be indeterminable [256]. Clearly, the designers of a system should not be held responsible if they cannot predict its actions or the environment in which it will be immersed.

Secondly, it seems likely that the developers will know about the possibility that the IDT could misidentify targets (and hence engage them in error). Similarly, if appropriate to the sophistication of the weapon, it is likely that they will understand that the IDT could potentially develop unresolved or contradictory states that result in other ‘unpredictable’ behaviour like that described in the UGV scenario at the start of this section; the systems will not be released untested. As companies employ lawyers to advise them on their responsibilities in producing, advertising and selling their products and UVS companies are unlikely to be any different in this regard it seems highly likely that the UVS developers will adequately fulfil their legal obligations in regard to any potential issues of product liability. Thereafter, assuming that there was not a faulty implementation of an otherwise acceptable design, blaming the developers would also seem to be both unfair and impractical.

Thirdly, in order to ensure culpability falls upon the guilty developer we must account for the complex and polymorphic nature an autonomous UVS. For instance, many UVS will comprise architectures where the decision-making elements are distributed across a number of programs and processors with results that are derived from data input from a wide range of sensors, a number of concurrently interacting components, machine and humans, none able to make the decision in its own right. There will, therefore, be interaction between multiple programs developed by multiple programmers (possibly from different software houses) on a variety of processors, operating systems, and architectures (perhaps unknown to each other in advance) and possibly distributed across a network of

¹¹¹ This includes the manufacturers, systems integrators, systems designers, soft/hardware engineers, etc.

UVS platforms. Moreover, the sensed data can probably be interpreted from a number of perspectives, each commensurate with different aspects of the integration or mission. Additionally, some military UVS may be designed to be unpredictable so as to inject a degree of flexibility or creativity into the system. Establishing a causal link would seem to be a task of Herculean proportions.

Finally, in all likelihood, the academics who devised the algorithms used in the decision-making process may have had no intention of it being used in this way. Therefore, to allocate any responsibility to them – even if their research was funded by a military organisation – would seem to stretch the bounds of legal or common sense.

Regardless of where responsibility lies, we must try to identify and apply a clear legal doctrine that simultaneously conforms to the LOAC (and any rules of the road), allows UVS development, and progresses its integration into capability such that it can be used with or without weapons. If we do not, those less responsible or less scrupulous, who also wish a capability edge, will likely simply develop and integrate less restrained versions into their militaries.

6.4.1 *Legal Personality for UVS*

Given the difficulties in attributing responsibility to the humans or their agencies, why not ascribe it to the UVS? To address this we must consider the legal concept of *Agency*. Agency is a highly specialised field of law but may be summarised as: an agent is empowered by their principal to negotiate and make various arrangements on their behalf; thereafter, the principals are bound by the contracts that their agents sign as if they themselves had signed, unless it is possible to prove misconduct. Moreover, the agent's actual authority extends to cases of apparent authority, where the agent has no actual authority but where the principal permits him to believe that he has authority [24]. To this end, as UVS become more sophisticated and are able to perform a range of complex actions on behalf of their human supervisors it is attractive to think of them as agents of their users.

The usual philosophical and legal debate over whether or not legal personality¹¹² can be ascribed to artificial agents centres around the list of necessary and sufficient conditions that must be met by the artificial agent in order for it to be recognised as an equivalent to a human. However, while UVS may ultimately achieve this status, this is not likely to occur any time soon. Consequently, we may choose to afford UVS *quasi-legal* status and allow them responsibility for a limited set of decisions and actions.¹¹³ In this way, we might consider UVS to hold diminished responsibilities such that the liability for certain decisions could then be transferred to the UVS or differentially apportioned between the UVS and humans or their host organisations.

¹¹² Typically, a legal persona has the capacity to sue and be sued and to hold property in its own name [68].

¹¹³ Some legal entities (e.g. children, corporations and the mentally impaired) frequently act through agents. Moreover, not all legal entities share the same rights and obligations; some (e.g. marriage) depend upon age whereas others (e.g. voting and imprisonment) are restricted to humans.

Clearly, in order to consider UVS as agents, we must imagine some form of ‘contractual’ obligation to exist between the user and the UVS. The mechanism through which electronic agents are able to create contracts is typically known as “attribution law” [149], which acknowledges the ability of electronic agents to conclude contracts independently of human review and alter the rights and obligations of their principals in their absence. In essence this law stipulates that a person’s actions include those taken by human agents and those taken by electronic ones. In other words, any transaction entered into by an electronic agent is attributed to the person using it, the requisite intention flowing from the programming and use of the IDT. However, legal opinion seems to be divided upon the question as to whether this attribution accords with the general principles of property and contract law: different European courts have produced contradictory rulings [270]. Nevertheless, regardless of the current statutes, the legal principles remain valid.

Another legal concept that is useful is that of Copyright Law¹¹⁴. That is, it may be possible to think of the UVS and its human supervisor as holding joint rights to the decisions, in the same way as joint authors are considered co-owners of a single body of work. The human’s contribution to the decision would be the input directing the UVS output, which might vary from significant to trivial depending upon the level of human supervision required. The key is that the decisions made by both parties are merged in an inseparable and interdependent manner such that they result in a unitary outcome [33]. There are however problems with this concept. For instance, what if the UVS adaptively learns from its environment and takes decisions independent of the supervisor’s or programmer’s original input or intention? Under most jurisdictions, the UVS would then become the ‘owner’ of the decision, because in copyright terms this would be considered a derivative of the original work?

There appear to be five options that may lead to a resolution of the problem; three that require us to apply or slightly modify our interpretation of current ‘legal doctrine’ (and possibly apply some ‘fine tuning’ to the existing law) and two that adopt more radical approaches, one requiring us to treat UVS as legal agents of their users, and another where the UVS are accorded separate legal personality.

6.4.2 UVS as Tools

Given the difficulty of ascribing intention to a UVS it is perhaps easier to acknowledge that they are just tools of their supervisors such that they allow lethal engagements to take place. This provides the most straight forward approach to

¹¹⁴ The author of a work is the initial owner of the copyright in it, and may exploit the work or transfer some or all of their rights in that work to others. The author is generally the person who conceives the copyrightable expression and fixes it, or causes it to be fixed, in a tangible form [32]. However, there are precedents for those who conceive and fix the work in a tangible medium not holding the copyright. Those who pay for the creation of the work, rather than the employee who conceives the work, hold the rights. To this end, it may be useful to think of the human supervisor as a commissioner or employer of the UVS,¹¹⁴ who in turn is the writer or executor of the work.

the problem, which is to insist on a rigorous interpretation of the current LOAC doctrine. In effect, it disregards the decision-making involvement of the UVS in the engagement entirely, thereby assuming the UVS never plays an active cognitive role in the target prosecution, regardless of its autonomy. The actions of the UVS are then ascribed to its user, regardless of whether or not they are intended, predicted, or mistaken.

While simple in its application, as the sophistication of UVS increases, it will be less and less realistic to assume that they are simply tools of their principals. Hence the limitations will become either unjust or inefficient and will almost certainly stifle the introduction of UVS onto the battlefield, which in turn may lead to nations ignoring the protocols.

6.4.3 *Human Intervention*

Another straight forward approach is to simply consider the lethal prosecution of a target illegal unless human intention can be identified at all (appropriate) stages of the process. This is based on similar logic to that used in the preceding section and may in fact simply be a practical outcome of applying such a doctrine.

The likely impact of such a proposal, however, is that operational tempo will be diminished if the UVS must rely upon human intervention, particularly for complex, networked or high workload environments. The commensurate effect will then either be to impede the introduction of UVS into capability or to promote a disregard for certain aspects of the LOAC. As soon as such interpretations are justified by one nation, others will probably follow.

6.4.4 *Technical Equivalence*

Yet another straight forward strategy would be to accept that technological progress is largely monotonic and that, once UVS have demonstrated their capacity to discriminate between targets to a given level of competence, entry into war accepts that, just as humans can make mistakes, any UVS flexible enough to be delegated our judgements can also be expected to err. This is reinforced by the considerations of software and hardware reliability outlined in the earlier sections on UVS. In the absence of other criteria we could assume that human equivalence is sufficient to use as the threshold test for the level of automation that we are willing to accept as being within the LOAC. However, this again poses problems.

First, with further improvements in sensor technology, processing and fusion, a UVS will be provided with information that is equivalent to or better than that provided to a human. Consequently, in at least some circumstances the UVS will be capable of making superior assessments in regard to the application of lethal force. At this time, the push to insert such systems into military service will increase.

Second, there is a growing perception that aside from any improvements in sensor quality, technological advancement (e.g. agents that can filter and fuse data) can significantly affect the quantity and quality of information available to a commander prosecuting a target. This will enable the human to more precisely

determine whether (for instance) non-combatants are to be placed at risk. This has two consequences: the first is to legitimise precision engagement and criminalise collateral damage, such that human equivalence may be inadequate [159]; and, the second is that the commander may increasingly rely upon indirect, filtered information, which may not contain information relevant to his obligations under the LOAC [177]. This issue may become even more acute in networked multinational coalitions when the information supplied by or to a force may be filtered for/from commanders of different nationalities.

Aside from the challenge of defining the various criteria by which we measure human equivalence, the other limitation of this strategy is that while the LOAC accept that civilian casualties are sometimes inevitable, it is also a requirement that the taking of life not be indiscriminate. That is, if we accept that it is legitimate for technology to err in terms of discrimination, some nations may use this approach to introduce indiscriminate targeting or breaches of LOAC in other areas of technology. In other words, while some would probably argue to the contrary, this solution could fundamentally undermine a key principle of the LOAC.

On the other hand, technology will continue to improve and it is reasonable to assume that the agencies who bear responsibility for developing UVS will continue to be charged with devising strategies for rectifying any technical problems with target discrimination techniques. These improvements will present a reference by which the quality of less discriminatory targeting UVS are judged in a way that is little different to other weapons technologies.

6.4.5 UVS as Agents

The final two solutions involve taking the artificial agent theory seriously by treating the UVS either as the legal agents of their users or as independent legal personalities.

The fourth option requires us to treat the matter as if the engagement were akin to a 'contract' between the two opposing parties, with the UVS acting as an agent on the supervisor's behalf. Then, in the same way that a person can be bound by signing an unread contract – or even a contract generated by a computer over which the principal has no direct control – we assume that the supervisor may be bound by the decisions of a UVS acting in accordance with his general intentions even though he was not cognisant of the detail of its actions in terms of any specific engagement. This option, however, requires us to distinguish between autonomy and unpredictability.

In the law of agency, a party's assent is not necessary to form a contract. It is sufficient, judged according to 'standards of reasonableness,' for one party to *believe* that the other party intended to agree; the real but unexpressed state of the other party's mind is irrelevant. In other words, at the point of use a commander will have certain expectations of the way in which a UVS will behave, deploy its weapons, etc. In our case this would mean that we adopt the doctrine that, without ever 'knowing the mind' of the IDT, a 'contract' between the user and the UVS is

assumed. If the UVS behaves autonomously, but not erratically or unpredictably, responsibility for its actions would then be in accordance with the current interpretation of the LOAC. Alternatively, if the behaviour of the UVS is deemed erratic or unpredictable then it would be acknowledged that liability for its actions do not fall upon the users.

While attractive, the practical limitations of this approach are significant. For instance, the distinction between autonomous (expected) and erratic or unpredictable behaviour for UVS must be determined. That is, standards of reasonable autonomous behaviour must be established and defined. Aside from the significant technical hurdles of such an undertaking, most militaries are unlikely to want to expose the capabilities of their potentially leading edge technologies. Nevertheless, it may be possible to describe these in broad, but acceptable, qualitative or descriptive terms.

Even if the world's militaries were willing to share the relevant data, the UVS may be so complex as to make reproducible testing impossible, or they may have the capacity to adaptively learn from their experiences and environment. In fact, at present, due to the stochastic nature of most systems, the best that is likely to be achieved is a Monte Carlo simulation to determine performance averages and bounds. This is likely to be cost prohibitive for real systems and the problem will be compounded for multiple or networked systems. Additionally, to achieve our goal of establishing standards of reasonable and erratic behaviour it is not enough to simply determine statistical norms, we must also compare the stochastic performance of the UVS against a series of previously agreed metrics that represent (say) human or other 'normal' behaviour; and most would agree that these do not currently exist. Furthermore, we must then establish which criteria define human equivalence, how many of these are essential as opposed to merely useful, and what we do if the thresholds for several criteria are exceeded, but not others.

6.4.6 Separate Legal Personality

We have already indicated that lethal operations conducted by UVS do not appear to fall within any of the existing exceptions to Walzer's Principles; so why not create a new one? To do this we need to focus on the *fact* of the engagement rather than the *process* of the engagement. Specifically, we would need to accept that human intention need not underlie the use of lethal force in an engagement. In other words, we assume that the user's generalised and indirect intention is sufficient to render the lethal prosecution of a target legal. While reliant upon the technical competence of the UVS to achieve a certain practical standard (e.g. human equivalence) this would, in effect, extend the current interpretation of the Principle of Responsibility to incorporate the UVS, which would need to hold separate legal personality.

Holding separate legal status does not imply that the UVS is considered a person.¹¹⁵ It simply recognises that being human is not a necessary condition of

¹¹⁵ Although the technical legal meaning of *person* is "subject of legal rights and duties" [1].

being accorded legal personality.¹¹⁶ Furthermore, there are well-established precedents for dealing with abstract and inanimate objects that enjoy legal rights and duties (e.g. nation states, the church, gods, shrines, corporations, ships, dead people, trees, and animals). There are also modern (and evolving) theories that define the legal criteria for agency and identity, not on the basis ontological properties (mind, soul, empathy, and the reflexive capacities) possessed by an entity, but on the basis of its ability to communicate [270].

Ultimately the legal status of a UVS will probably be decided as a matter of pragmatism rather than on the basis of “computational substrate or internal architecture” [255]. That is, it is unlikely that such legal personality will come to the courts ready-formed. Furthermore, unless there is demonstrable economic, doctrinal or capability benefit for such a change there is unlikely to be strong internal pressure from the military or legal communities to accord it. It is clear, however, that the vagaries of the ‘rules of the road’ and LOAC relative to the autonomy of UVS may provide impetus, although this should not be overstated. Quite possibly, the greatest considerations will be the political and financial cost-benefit analyses, as few modern laws are proposed without debate surrounding the relative merits of such packages.

6.4.7 Impediments to Artificial Agency

Artificial agents are now ubiquitous and include avatars (graphical icons that represent a real person), vReps (virtual representatives that are used to humanise online relationships and provide a single point of contact for customer enquiries through natural dialogue), digital buddies (software programs that automate chat with users), and shopping bots (programs that collect and compare information online and recommend products) [149]. These agents are essentially self-contained code units that interact with humans, one another and/or their environment. They often fulfil multiple roles, including information provision (e.g. filtering or fusing information), decision-making (e.g. the generation of goals, plans, and schedules), and action execution (e.g. closure of contracts). Moreover, their functions are similar – and in some cases superior – to the humans they replace. In the civil environment, however, an agent’s actions are to a large extent reversible: the contract can be declared void, the goods returned, etc. In the military domain, the consequences of an IDT’s actions may not be reversible as they may result in loss of life, destruction of cultural heritage, etc [246].

Some will simply dismiss the notion of agency or separate legal personality for UVS as fantastical and not something deserving serious consideration. Moreover, there are obviously a number of objections to non-humans holding legal status and rights. Several commentators (e.g. [12] [24] [68] [149] [255] [270]) have discussed these from the perspective of artificial intelligence, contract law, and the

¹¹⁶ Similarly, being human has historically not been considered a sufficient condition to be recognised as a legal persona: in England, prior to the middle of the 19th century, a married woman was not considered to have separate legal status to her husband. Similarly, in the US slaves were considered “non-persons” until the Emancipation Proclamation of 1863 [68].

agents holding rights under various Articles of the US Constitution. The philosophical objections usually include:

- It is only right that natural persons be given legal rights;
- UVS lack soul, consciousness, intentionality, free will, etc; and,
- As human creations UVS can never be more than human property.

Readers are referred to the works above and those referenced therein for a more detailed discussion. Here we simply note that there are clearly contrary intuitions and arguments, and that each depends on values or assumptions that are not necessarily universally shared (i.e. is a child or a human clone a human creation?).

The commentators' various legal objections usually include:

- It is unnecessary as other solutions are adequate;
- The inability to sue an errant agent is not in practice a significant loss;
- The agent has no assets and so any judgement against the agent is meaningless¹¹⁷;
- In multi-agent systems, which comprise multiple copies of the same code and are in communication with one another, it is not clear whether the entity should be recognised as singular or plural in a legal sense.

Other objections to treating artificial agents as true agents may be interpreted as:

- UVS would lack the legal power to give consent because they are not persons. However, as [68] points out, Roman slaves were not considered persons, but still had the legal capacity to enter into contracts on behalf of their masters.
- The agent onboard the UVS would not have the intellectual capacity or the ability to exchange or represent 'promises' of a nature that a user could depend upon. Here we note that in the civil domain web and software agents are ubiquitous and we interact with them every time we shop online.
- Some legal systems require contracts to exist between agents and their principals: artificial agents are not persons therefore they cannot enter into contracts in their own name. However, as we have already pointed out, in Anglo-American law it is only necessary for the principal to be willing for the agent to bind him as regards third party contracts.
- The final objection, which is raised by [149], relates to the need for the agent to have the mental capacity to comprehend the nature of the act being performed. Given that UVS are not likely to have human-like intelligence at this level of cognition for some time to come, the current legal doctrine would probably need to be modified before artificial agency could be allowed.

¹¹⁷ Perhaps not the case if the UVS holds legal personality and has insurance purchased on its behalf.

6.4.8 *Technical Impediments to Agency*

In regard to an IDT onboard the UVS, we may make the following assumptions [145]:

- It will make decisions
- It will be employed for its creative intelligence and judgement
- It will be distributed across an ensemble of programs and processors
- It will be polymorphic and immersed in a changing environment
- Consequences may be unpredictable, surprising and unintended

Unfortunately, as we have seen from the earlier section on UVS, many of the decision-making elements of the UVS are distributed across a number of programs and processors such that no one processor, program, programmer or user will know the full context of certain decisions. Moreover, if the UVS forms a component of an NCW environment, it may be that agents originating at another node in the network are executed within the UVS (or vice versa), such that any users are unaware either when or where the agents are executing or from where they originated. The agents simply cooperate with one another across platforms & operating systems autonomously.

Furthermore, in complex environments, it is not necessarily possible to know (a priori) what concepts (legal or technical) are relevant to any given situation. For instance, an IDT may interpret data from a number of different perspectives and manipulate the information depending upon the nature of the mission or problem. Additionally, some defence UVS controllers may be designed to be unpredictable so as to inject a degree of flexibility or creativity into the system, as predictable systems are not necessarily optimal for military operations.

Developers will naturally strive to achieve 'best practice' by implementing basic rules of thumb, keeping their designs simple, providing suitable documentation and creating initially stable designs. Empirical evidence¹¹⁸ suggests, however, that it is practically impossible to fully debug a sophisticated program on an unblemished processor. Moreover, we are considering the interaction between multiple programs on a variety of processors, operating systems, and architectures across a range of platforms (perhaps unknown to each other in advance) all operating in environments with which they are inherently unfamiliar. To this end, we can be reasonably sure that the UVS will malfunction at some level and that we are unlikely to be able to predict the specific nature or timing of the failures. This will present great challenges to traditional legal analysis in regard to the attribution of responsibility even at the processor level.

The upshot is that the operation of autonomous UVS entails considerable fundamental uncertainty, particularly in complex environments. The failure of a

¹¹⁸ There are inherent problems with software reliability as, despite the best endeavours of mathematicians, programs cannot yet be verified as correct over an arbitrary set of inputs [145]. And the problem is not limited to software. The processors on which the software runs can also be thought of as programs encased in silicon and they too usually contain bugs. Consequently, from time-to-time, software fails.

UVS may then not always be due to human negligence during their creation, acquisition or operation, particularly if the layout of the card was devised by an automated system or a reinforcement learning algorithm was used to adaptively modify the code or coefficients in some algorithm. As a result it may be beneficial to avoid trying to identify a causal relationship between any injuries caused and the functional component level of the UVS. That is, we may conclude that there are a class of circumstances that allow identification of faults within or by the UVS at the systems level and that not do not require detailed investigation into what, who, when, or why the injury occurred at a component, card, processor or algorithmic level.

Although this may seem unreasonable and it is recognised that the law may hesitate to assign responsibility to a UVS or its IDT when it could more reasonably have been ascribed to someone who *might* have been able to control the outcome it may be preferable to the alternatives, particularly if the best that can be achieved through causal analysis is indeterminate or ‘unpredictable.’

6.4.9 Is the UVS the Real Decision-Maker?

Perhaps the most potent objection to an IDT ‘taking responsibility’ is the question “Is the UVS the real decision-maker?” In order to take the option of separate legal personality seriously, therefore, we must establish whether, given its limited capabilities, the UVS is in fact be the real decision-maker. The key argument in this regard centres on whether the natural decision-maker (the human commander) making the decision to use the UVS in the first place constitutes the crucial discretionary judgement as the power to make these discretionary decisions essentially identifies the principal authority.

However, if we assume that the UVS will probably assume a role in warfighting on the basis that it improves warfighting efficiency and effectiveness – that is, for instance, it allows the prosecution of more targets per person – how can we reasonably argue that the real decision-maker is its supervisor, who may only be making a few discretionary interventions on behalf of the UVS, leaving the technology to cope with the majority of decisions, particularly if the human has no direct contact with the targeting process – i.e. the UVS has undertaken the ‘mechanical’ tasks required to identify and engage the targets. In these cases, it may be necessary to have a formal hand-over procedure so that we could say “The UVS was the decision-maker until 03:15 UTC on 12th June 2008, whereupon the human took over.”

Another reason for objecting to the IDT being held as the real decision-maker is based on the belief that it cannot follow a system of rules (however complex) sufficient to enable it to make appropriate judgements and exercise discretion. Here there are a number of considerations:

Changes of Circumstance: It is not possible to design a UVS that can anticipate all possible targeting decisions and circumstances. How then can we expect it to make good decisions in certain situations? Conversely, if we cannot codify such eventualities, how can we prepare Rules of Engagement (ROE) for humans to

make judgements consistent with the LOAC under the same circumstances? For a UVS (as for a human commander), the ROE can therefore be designed to minimise such possibilities. Secondly, the ROE can provide for changes of circumstance by specifying that, if the UVS finds itself unable to carry out its relevant functions, a change in decision-maker (i.e. from the UVS to the human) is permitted.

Decision Outcome: From time-to-time humans wish to improve decision outcomes not simply by recognising novel or more complex approaches within the frame of an existing problem but by re-framing the basis of the problem itself. In this way it might be considered that they have superior decision-making capacity to a UVS. As a minimum, therefore, we must design a UVS with the capacity to recognise and deal effectively with complex novelty and to maintain information regarding any decisions it has made in a form readily interpretable by humans. In this way, we reduce objections to a level commensurate with our (a priori) capacity to codify such eventualities for humans.

Moral Judgements: How would an IDT exercise a moral judgement to implement the LOAC against a simple test of fairness – as opposed to a literal interpretation of what is strictly permissible? This is a particularly significant issue and is discussed extensively by [23] and briefly in the next section on the Ethical Control of UVS.

Legal Choice: In order to reduce the complexities of allocating responsibility, it has been suggested that a UVS can act as a quasi-agent and may be considered legal persona (that might hold insurance). Given this, the UVS may need legal representation in any dispute over liability (as a legal persona it will also have certain legal rights). How could the IDT exercise judgement and discretion regarding over matters pertaining to (say) the settlement of a complex lawsuit? The most straight forward response is that we have already determined that the IDT must maintain a suitable record of action for the purposes of evidentiary hearings, which would seem to satisfy the most basic requirement of this objection. Thereafter a lawyer could be engaged on behalf of the UVS, who could act as ‘trustee’ for the UVS in regard to any decisions regarding litigation. This would seem to be similar to the representation of minors in acrimonious child custody cases under various jurisdictions.

Before leaving the discussion on legal issues, two final thoughts:

- First, it is often stated that “Humans will always make the final decision” because they have ‘mastered logic’ (or some derivative of that phrase). The implication is that they can apply rational human judgement to special circumstances in a manner that IDT cannot. While it is true that humans are generally good at pattern recognition – and therefore recognising ‘abnormal’ circumstances – it should be noted that only *some* humans have mastered logic to a high degree; others have not. On the other hand, IDT now have the edge over most humans when it comes to

assimilating, ranking and storing information from a diversity of sources, processing it using mathematical formulae and then interpreting and applying the results of these calculations as actionable outcomes.

- Second, if the UVS holds separate legal personality and has obligations, it will presumably also have rights. Consequently, it might be reasonably argued that it holds some form of legal title over the hardware and software that enables it to operate [254]. At the very least, it might seem reasonable that the insurance company might apply for some form of title to protect its financial interests.

6.5 Ethical Control of UVS

Autonomous UVS will be integrated into military service when they are technically capable of undertaking their defined roles, it makes sound economic and military sense to do this, and there is an appropriate legal framework to accommodate them. There may also be other advantages. For instance, what if we could ensure that UVS make morally superior decisions to a human? Among other things, this might reduce the diversity of interpretations of LOAC, which would not then be made against a backdrop of mainly Western, but nevertheless varying, ethical systems and international cultures, but would be superseded by a more unified and global interpretation of them that would evolve from the specific implementations of what is technically achievable.¹¹⁹

Such challenges have held fascination for humans since Isaac Asimov published his three laws in “I, Robot” [27]¹²⁰ and risks sounding as though we are now crossing the line into the realms of science fiction or fantasy. Nevertheless, if we consider that many automated systems are now in the position of making split-second decisions that have life-or-death consequences it is perhaps not so far-fetched a notion. Furthermore, if we also interpret the challenge from the perspective of what it would take to design an IDT capable of making decisions from the perspective of what the UVS *ought* to do rather than simply allowing it to do what it *may* do,¹²¹ the discussion sounds more plausible.

Let us start by considering an industrial robot. It has limited autonomy and must ‘choose’ between courses of action based on appropriately ‘moral’ judgements, albeit that they are effectively framed a priori by safety standards (e.g. [210] [257]). These systems tend to operate in very limited frames of

¹¹⁹ We make extensive reference to the work of Arkin and other’s in this short section. The aim is to highlight some of the complexities and opportunities for instantiating the ROE and LOAC-based control of autonomous weapons. For a more extensive treatment of the topic the reader is referred to [19] – [23].

¹²⁰ These are: (i) a robot may not injure a human being or, through inaction, allow a human being to come to harm; (ii) a robot must obey orders given to it by human beings, except where such orders would conflict with the First Law; and, (iii) a robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

¹²¹ That is, how might we ensure that the UVS behaviour is in line with what we expect of it in accordance with the LOAC so that it “shows compassion” when the opportunity arises.

reference and it is reasonably safe to assume that the engineers and programmers who design and build the systems may be able to conceive the variety of options the system will be presented with. Consequently, they are able to devise and program appropriate responses. As a result, these systems do not need to specifically evaluate the consequences of their actions on the basis of some moral code or ethical stance; this has been done for them a priori on the basis of (say) the occupational health and safety laws and the common sense of the design team. On the other hand, there are a number of factors that suggest that military UVS will need explicit ethical governance in order to comply with the LOAC [22].

- The UVS may encounter situations that the designers had not anticipated, either because the UVS is used in a way for which it was not designed or because the designers had not anticipated some complex aspect of the environment;
- The increasing autonomy of the UVS will allow it to make increasingly complex decisions. This has an impact in two regards:
 - The variety of context that the UVS could find itself in will preclude a priori interpretation of each and every situation
 - The number of potential decisions or degrees of freedom available to the UVS will preclude a priori knowledge of each and every potential action
- The polymorphic nature of the system and the sheer complexity of the systems engineering precludes accurate prediction and testing of UVS behaviour

Despite this, we can articulate some guiding principals for implementing ‘ethical control’ in UVS by following these basic attributes [16]:

- Consistency – any contradictions in the informing theory must be avoided (i.e. an action must not be simultaneously right & wrong in a given set of circumstances)
- Completeness – the system must ‘know’ how to act in any ethical dilemma
- Practicality – the technical implementation must be feasible and we must be able to follow any action recommended
- Agreement with Intuition – the actions must agree with expert ethicist intuition

Furthermore, it may be that we encode the governance of the UVS relative to the LOAC in line with the ‘decomposition’ strategy outlined in the section on Verification and Validation. That is, rather than attempting to consider the autonomous UVS holistically or at a system level, we devise a strategy that allows us to divide it into its constituent functions and verify that each of these components or functions behave ethically. In this regard [37] presents an interesting methodology, tools and techniques for modelling the heterogeneous real time components of a system as the superposition of three layers: a

behavioural layer (specified as a set of transitions), interactions between the behaviours, and priorities (used to choose between the possible interactions).

The central idea – and the one most pertinent to the instantiation of ethical governance or control – is that while complex systems are built through the assembly of components, the main problem is that these components must be integrated in a way that ensures their correct interaction. Aside from obviously describing the key integration challenge for the functional elements of a UVS (and therefore possibly of considerable use in terms of V&V), this methodology also describes the overarching problem for the ethical control of a UVS within the framework of ROE and the LOAC. That is, the guiding principals may be considered consistent, complete and practical when interpreted in human terms, but there is no mechanism for provably demonstrating the composition, integration and software coding of these ‘abstract’ principles have ‘mechanical’ integrity. Consequently, the enactment of any physical functionality within the UVS against these ‘rules’ in a way that preserves the intended properties of the LOAC framework in real time is unprovable. However, the BIP (Behaviour-Interaction-Priority) [37] framework provides a basis for the study of property-preserving transformations or transformations between sub-classes of systems or principles, which in turn allows such component interactions to be verified.

6.5.1 Could UVS Ethically Out-Perform Humans

Humans have an impressive ability to rapidly, spontaneously, and effortlessly recognise and identify a large variety of objects even under unusual conditions. On the other hand, a human’s capacity to detect, locate, and track certain objects, particularly in cluttered or attenuated environments, at long ranges, in highly stressful situations, or for protracted periods is inferior to those of sensors combined with intelligent processing. Moreover, the human eye has limited spectral sensitivity, does not see well at night, and its processing can be tricked rather easily [130]. With improvements in sensor, sensor processing, fusion, and communications technology it therefore seems likely that an IDT will be provided with the capacity to acquire, select, represent and retain certain types of information with a capability approaching or exceeding that of a human.

This inevitably leads us to the conclusion that in some circumstances the UVS may be capable of making superior assessments in regard to the target identity and hence the application of lethal force. For instance, during periods of high stress a soldier may feel the need to act in self-defence and may respond lethally to all targets within a given sector, which may inadvertently cause civilian casualties.¹²² As IDT-controlled UVS need make no appeal to self-preservation they can value civilian lives above that of their own. This does not negate the likelihood that civilian lives will be lost, but it does indicate a way in which the IDT might ethically out-perform a human. Moreover, before proceeding (or dismissing as fantastical the notion that autonomous ethical governance could out-perform a human), it is instructive to consider the findings of [262] and which were first

¹²² Self-defence is a common justification for the exculpation of responsibility for civilian casualties [16].

noted in this regard by [22]. The findings were taken from a sample of 1,320 US soldiers and 447 US marines.

- Only 47% of soldiers and 38% of marines agreed that non-combatants should be treated with dignity and respect;
- 44% marines and 41% soldiers agreed that torture should be permitted to save the life of a fellow marine/soldier and 39% marines and 36% soldiers agreed that torture should be permitted to obtain important intelligence about the enemy;
- 45% of soldiers and 60% of marines stated that they would not report a fellow soldier/marine if he had injured or killed an innocent non-combatant;
- 33% marines and 27% soldiers did not agree that their NCOs and officers made it clear that they were not to mistreat non-combatants;
- Even though they reported receiving ethical training, about 30% of those surveyed reported facing ethical situations in which they did not know how to respond
- Combat experience – particularly the loss of a colleague – was related to an increase in ethical violations.

In addition to the negative aspects highlighted above, there are also a number of ways in which UVS might out-perform humans on the battlefield [19]:

- They can be designed without emotion, which might cloud their judgement;
- The absence of a sense of self-preservation means they can be used sacrificially;
- They are unlikely to be affected by decision bias – the interpretation of (usually contradictory) incoming information such that, particularly in stressful situations, it is fitted to pre-existing scenarios or belief sets of the decision-makers;
- UVS can potentially integrate more information from more sources and more sensors more quickly than humans can (without getting fatigued or distracted);
- The sensors onboard the UVS might be used to report or police human behaviour (independently and objectively – unlike their human counterparts), which might also lead to a reduction in unethical battlefield behaviour over time;¹²³

Before any responsible military could proceed with the introduction of such a capability, however, a series of metrics aimed at determining whether or not the

¹²³ It should be noted, however, that while mobile sensors on the battlefield might be used as policing devices this may have a negative “big brother is watching” effect on some of the troops in combat. This might in turn impact negatively on intra-unit trust and cohesion; although this was not the experience of police forces that introduced video taping into the interview process.

UVS actually delivered greater ethical behaviour on the battlefield would need to be established against the principles of the LOAC. Such a set of criteria might include [228]:

- Does the prospect of lower risk to their own forces result in the more frequent use of military force by governments
- Are autonomous UVS able to achieve greater levels of target discrimination
 - Are there fewer non-combatant casualties
 - Is there a reduction in collateral damage
- Do the IDT improve proportionality
- Are fewer of the forces that use the UVS killed or injured
- Could a captured UVS be used malevolently

6.5.2 *Implementation Issues*

Regarding the instantiation of ethical governance, there are two primary issues:

- What is considered ethically acceptable; and,
- What can technically be achieved?

In a study conducted by [198], where the researchers have canvassed opinion on the use of lethality by autonomous systems, initial responses to the first of these questions are presented. The survey demographic spans the general public, researchers in the field of robotics, policy-makers, and military personnel. The results of this survey are not discussed here further, but are reported in [23]. It should, however, be noted that ethics are relative: an act carried out by a UVS may be morally acceptable to a programmer, but perceived as immoral by indigenous cultures that witness its actions in theatre [22].

In regard to the second question, and relative to ‘everyday’ ethics, a number of authors [23] [46] [55] [180] [194] [198] [228] have acknowledged that the LOAC provide us with a ‘rigid’ framework that reduces the complexity of implementation somewhat. In other words, we neither want nor need the UVS to derive its own code of beliefs from first principles regarding the moral implications of the use of lethal force. Rather we wish it to apply those that have previously been derived by humanity and coded in the LOAC. All the same, we must be able to certify that our UVS behaves in accordance with these coded laws and any ROE, which (fortunately) are designed to be self-consistent, even if not universally agreed upon in terms of specific definition.

Implementing an artificial ‘moral agent’ that has the capacity to govern the actions of a UVS will involve a broad range of engineering, ethical and legal considerations and a full understanding will require a dialogue between philosophers, legal theorists, engineers, computer scientists, developmental psychologists and other social scientists. Moreover, it would be naïve to assume that such an agent would solve all issues of responsibility; this exercises even the International Criminal Court. Similarly, it is not anticipated that such a system

would have the capacity to adequately interpret all situations to the satisfaction of the local commander or user. Under these circumstances it is presumed that they could assume responsibility for over-riding the agent.

In addition to the difficulties outlined in the section on Contextual Decision-Making, therefore, the practical instantiation of an ethical moderator will tax us in a number of ways:

- The creation of such a ‘technical conscience’ will challenge our understanding of morality to its extreme as, in order to code the conscience, we must formally express our knowledge of the moral framework in a manner that engineers and computer scientists can both understand and express in software/hardware [290];
- Ethical reasoning is based on abstract principles, which often conflict with each other in specific situations. If more than one law, code or principle applies it is often not clear how to resolve the conflict as the favourite tools of logicians and mathematicians (e.g. first order logic) are not usually available [178];
- The LOAC and ROE are intended to be self-consistent and complete. Moreover, they provide a broadly agreed framework upon which we can agree a basis for the development of this agent. However, the premises, beliefs and principles that humans use to make ethical decisions are varied and often intimately linked to religious beliefs [16] (i.e. there are a variety of accepted ethical theories upon which to base our computations).
- Additionally, the interpretations of the LOAC are not yet agreed between even the democratic nations of the world [16];
- If such an agent is able to prevent an IDT from acting in some unethical fashion it must also have the capacity to explain to its supervisor (at least to some degree) the underlying reasons for its logic [14];
- The conditions, premises or clauses are not precise, are subject to interpretation and may have different meanings in different contexts [13]; and
- The actions and conclusions that we obtain from such a system are likely to be abstract, so even if a rule is known to apply the most appropriate action may be difficult to execute due to its vagueness [22].

There are a number of groups working on the development of techniques that represent artificial ethical behaviour (e.g. [7] [14] [22] [55] [180] [194] [198] [228]), most proposing to use a variety of ethical theories (e.g. Kantian, Utilitarianism, Social Contract, Cultural relativism, etc). Several authors have pointed out that a distillation of ethical directions from the LOAC – a ‘top-down’ approach, if you like – might require us to instantiate a more rigid structure than we would wish; and thereby risk overlooking events or circumstances which then cause the UVS to behave inappropriately because it is bound by a set of rules. Alternatively, we might seek to use an ‘optimisation’ or ‘bottom-up’ approach that relies upon adaptive programming or machine learning. Constraints (e.g. the LOAC) could then be applied and the UVS placed in a variety of circumstances

(training cases) to make it learn using trial and error. However, this too may have its difficulties.

First, because it does not have a complete ‘set of rules’ to guide its overarching behaviour the UVS may effectively assemble its own, and, if those are inadequate (or the training data is incomplete), the UVS may once again perform inappropriately.

Second, the UVS will be trained against data selected by their owners and operators. Just as education can enhance cognitive skills it can also create fanatics, dogmatics, cynical manipulators, as well as prejudiced, confused, and selfishly calculating minds [47]. Consequently, while the developer of a reinforced learning algorithm that is aimed at enhancing the capabilities of a UVS would presumably expect it to be immersed into environments and then trained or used in a way that leads to improvements in the IDT’s response options relative to the LOAC, there may be those less scrupulous who seek to exploit (and possibly identify) aspects of an algorithm that result in the UVS-equivalent of a fanatic. That is, a system that over-rides some pre-ordained taboo, and otherwise high-priority, action on the basis of some learnt behaviour. This obviously introduces another level of complexity into the problem.

Arkin [22] approaches the topic from the pragmatic perspective of first applying it to limited and specific contexts and, as they learn, applying it more widely to new contexts and applications. The work approaches the design of the weapon control not simply as a problem in ethics, but from the perspective of safety. It uses a hybrid architecture that strictly adheres to the rights of non-combatants by regarding discrimination paramount (i.e. using deontological reasoning) while simultaneously considering proportionality on the basis of military necessity (i.e. using utilitarian reasoning). It effectively starts from an extension of the medical premise “First - do no harm”¹²⁴ (i.e. do not engage an enemy until obligated to do so). Based on the derived situational awareness picture, a set of circumstances, C, are defined which are parameterised and characterised on the basis of interpretable stimuli, S. C then comprises vectors to represent the building blocks of situational awareness [96]:

- The perception of the elements in the environment;
- The comprehension of the current situation; and,
- The projection of the future status.

The development of each of these blocks represents increasing levels of complexity. Initially, the interpretable stimuli, S, might consist of a vector expressed in terms of a perceptual class¹²⁵, a threshold value for each perceptual class, and so on. Complex, statistical analysis of (say) the geographic distribution and disposition of these classes would then allow us to parameterise and define circumstances that can then be passed to our ‘ethical moderator’ and acted upon in

¹²⁴ As opposed to the “Shoot and ask questions later” or “Apologise and move on” philosophies.

¹²⁵ To represent people (non-combatant/combatant), buildings, walls, vehicles, weapons, gestures, etc.

conjunction with the UGV's purpose, P, which is essentially the same as the (parameterised) high level mission or goals of the UVS set by the high-level interaction between the UVS and its supervisor.

A range of actions, A, is then defined for the UVS (e.g. its motion and payload response) that can be factored as two orthogonal components: strength and orientation. Strength would denote the magnitude of the response (e.g. speed or force), orientation the direction of action (e.g. retreat from target, move towards target, warning shot, engage target, etc). The instantaneous response of the IDT would then be expressed as a vector representing each of the degrees of freedom (DOF) for the UVS where weapons targeting and firing are considered within these DOF. We may now write a set of maxims [194]

If (C and P) \rightarrow A	A is Obligatory for the IDT
If (C and P) $\rightarrow \neg A$	A is Forbidden for the IDT
If $\neg((C \text{ and } P) \rightarrow A)$ and $\neg((C \text{ and } P) \rightarrow \neg A)$	A is Permissible for the IDT

Actions can then fall into one of three classes: obligatory (the UVS must act in this way, based on moral grounds), permissible (the act is morally acceptable, but not required) or forbidden (the act is morally unacceptable). Only acts that are deemed to be obligatory may be carried out. In other words, if an action is permissible it may only be carried out if it is considered obligatory in the context of the mission.

Using such a technique, we may now make sense of scenarios such as "UGV under attack and motivated by the need to protect Blue Force (Circumstance) shall try to reduce threat (Purpose) by responding with lethal force (Action)". As previously mentioned a record of action (in this case the UGV's ethical deliberations) would also be required; not just for after-action review, but also for the application of reinforcement learning techniques.

Once again, this does not imply the UVS has sentience or intention as we have effectively just articulated a way to deliver 'artificial morality'. That is, getting a UVS to act as if it were a moral agent. The first step in this process was to understand – and to provide the UVS with cognisance of – the possible harmful consequences of its actions, such that it can select from its repertoire accordingly.

The point of this commentary is to provide an indication of what might currently be possible and to highlight that development of such moral reasoning capability presents us with an interesting conundrum [19].

If a UVS is given an order such that $(C \text{ and } P) \rightarrow \neg A$ (i.e. the UVS effectively recognises that it has been told to carry out an unethical order), given that the UVS may hold separate legal identity, should it refuse to carry out the order? If it does not who is to blame?

In practical terms, this actually returns us to an earlier discussion regarding location of the crucial discretionary decisions. Clearly, if the UVS is to assume a role in warfighting on the basis that it improves warfighting efficiency and

effectiveness – and that warfighting remains for the foreseeable future a human endeavour – it will once again be necessary to have a formal hand-over procedure so that we can say “The UVS was the decision-maker until ..., whereupon the human took over.”

Interestingly, if we record the complex decisions being computed by the IDT, which may be necessary data for evidentiary hearings (and to develop and test any ethical governors) it may be possible to combine this recorded data with other ethical reasoning agents to improve and better understand complex matters of liability and reasoning in stressful and cluttered environments. Similarly, we may be able to apply adaptive learning techniques to the ethical reasoning components of the IDT and use a range of techniques that model moral and emotional perceptions such as guilt. This may then in effect allow us to use moral motivation for UVS behaviour, which somewhat ironically, may provide us with the capacity to ‘punish’ the UVS for errant behaviour – or provide the UVS with the incentive to commit war crimes.

6.5.3 *Architectural Considerations*

According to [184] a good architecture describes the following:

- The overall system organisation
- Ways in which the architecture supports likely changes
- Components that can be re-used from other systems
- Design approaches to allow standardisation of functional components
- How the architecture addresses each system requirement

There are at least four stages (i.e. architectural possibilities) that would allow moral control of an autonomous weapon, each representing a more sophisticated level of ethical governance or control. These are described by [19] as:

Responsibility Advisor - While it may seem fanciful to assume that machines will take over the role of ethical decision making in war, it is more reasonable to assume that, as technology allows, they may be used in some advisory role. Thereafter, when human decision-makers have been put in the position of passively receiving interpretations generated (for instance) by data fusion and hypothesis generation-aid machines and they become less able to recognise emergent problems or be able to interpret the complex ones, they may progressively defer to their autonomous advisors. It is a small step from here to accepting their autonomous governance [20].

This would form a good first step towards any eventual goal of developing an autonomous UVS that acts ethically. Essentially, the advisor would probably form part of the HMI used for pre-mission planning and managing operator over-rides. It could potentially advise (either in advance or during the mission) the UVS operators and commanders of their ethical responsibilities should the UVS be deployed in specific or complex situations. It would then require their explicit authorisation to use lethal force. One could start by designing an advisor that

provided guidance on the interpretation of the LOAC and ROE to selected personnel in specific circumstances, thereby retaining the notion that the human supervisors are making the decisions as it is they who must decide whether or not to follow the machine's recommendations. Another benefit of such a system would be that the necessary data required to make such decisions would start to be collected [20]. Additionally, as the interpretation of the LOAC has not yet converged to the point where there is universal agreement, difficult decisions could be flagged and passed to human specialists for later or more detailed analysis.

The ensuing architectures are likely to be instantiated when the performance of the ethical advisor has been shown to outperform human interpretation of the LOAC and ROE. Commanders and supervisors would eventually be compelled to follow the machine's advice or take responsibility for the actions of the machine thereafter. Over time, one would expect the role of the human in the ethical decision-making process to diminish and the number of dilemmas he handles to become considerable; perhaps he devotes little or no time to any particular one. There will still be times when a human is called upon to make a decision, however, and if such events occur with regularity, the developers of the system would need to be tasked with developing improved architectures for handling these, described by [19] as.

Ethical Governor – This is an extension of the previous design that physically transforms or suppresses UVS-generated lethal action to permissible action; either converting it to non-lethal action or permissible lethal action.

Ethical Behavioural Control – This approach constrains all individual control behaviours to be only capable of producing lethal responses that fall within acceptable bounds.

Ethical Adaptor – This architecture would provide the UVS with the ability to update the agent's constraint set and ethically related behavioural parameters, but probably only in a restrictive manner.¹²⁶

To this end, we need to understand the requirements of such architecture. For example, if weapons are ever placed on a UVS, the equivalent of a Weapons Safety Board (WSB) will need to sign off on the system's architecture with respect to the safety of the overall system. To this end, we would need to start by agreeing with the WSB on a level of automation with which they are happy and then work with them to arrive at an approved architecture. We will then have to determine and articulate the instantiation of this architecture against the state-of-the-art technologies, such that this instantiation is also signed off. To do this, we will need to define the sensor, technology, and 'intelligence' capabilities (most of which are software) that allow the requisite level of target

¹²⁶ Arkin [19] offers the view that such a component would be based upon an 'after-action, reflective review of the system's performance or by using a set of affective functions (e.g. guilt remorse, grief, etc.)' produced by a violation of the LOAC.

detection, discrimination, location, tracking, etc intrinsic within the agreed systems architecture.

Implementation of such a system is likely to be modular with specific code units run on physically separated hardware. It is also likely to have a variety of quarantined time-outs, queries, and messages to shift the weapon control unit back to a safe-operating mode, if required. Distributed systems are likely to introduce another level of complexity. Given the embryonic status of the development process, another desirable property is that the architecture supports growth so that extensions can be added incrementally. For instance, initially we are likely to want to test the systems against a representational yet small sub-set of forbidden and obligatory constraints; as opposed to trying to encode the entire LOAC and interpreting complex ethical situations at the outset [22]. In this way, the basic research required prior to the deployment of any operational systems can be conducted away from the complexities and heat of the battlefield.

Although this certification process is likely to be a lengthy one, we can postulate acceptable generic architectures. For example, it is likely that the governor will make decisions by projecting action into the future on the basis of a model of the UVS, its current and potential behaviour and any environmental dynamics. It will then evaluate the outcomes according to a function or some other selected criteria. This evaluation function then represents the objectives and ethical constraints through the return of high values for plans that meet mission goals without violating the performance envelope of the UVS. Typically, this will involve some form of search through a set of potential plans until an acceptable or feasible plan is found. Consequently, the key is to apply pruning techniques so that only successful outcomes are generated. As with code and system verification and validation there are several major risks:

- The moderation outcome makes inefficient use of resources;
- The moderation outcome could not be generated;
- The moderation outcome generated is not feasible;
- The moderation places non-combatants or civilian infrastructure at risk;
- The moderation places users, supervisors (or possibly the UVS) at risk.

It can be argued that we need to worry about the last one the most. However, to assure ourselves that this (and that the governor is working correctly) is not an issue we do not need to verify the entire system, only its evaluation function. If the evaluation function is correct then the UVS and/or users cannot be placed at risk. Furthermore, it is likely that human oversight will (initially at least) ‘double-check’ the results.

We may now imagine an architecture or algorithmic sequence, which combines the major Principles of the LOAC (i.e. responsibility, military necessity, target discrimination and proportionality) such as that suggested by [23].

Responsibility – Permission for the UVS to use lethal force autonomously or in specific situations (e.g. prior to or at certain points in the mission) is granted by the human supervisor. If relevant, the type or nature of weapon selection is also authorised.

Military Necessity – The UVS makes a determination, based on its observation, comprehension and projection of the circumstances and its assigned mission, as to the need to use lethal force at any instant. It may determine (for example) that it needs to move to obtain a different perspective, fuse data from other on or off-board sensors, or allow the use of lethal force.

Discrimination – Targets and other objects of interest (combatants, non-combatants, buildings, and so on) are accurately identified in a manner that allows positive discrimination between civilians and combatants and civilian and military objectives.

Proportionality – Based on the disposition of the objects of interest and the likely impact of the use of lethal force (weapon selection, firing pattern, etc) collateral damage calculations are performed.

Obligation – In order to satisfy the goal that the UVS is doing ‘what it ought’ rather than ‘what it can’ in order to use lethal force the parametric criteria must satisfy the ethical moderator not just that the proposed action is permissible (i.e. that any and all forbidden constraints are upheld), but also that the action is obligatory (i.e. at least one obligating direction must be upheld).

Weapon Control – The UVS then attempts to accurately control delivery of the weapon with respect to aiming at any targets and maintain a suitable record of action for the purposes of evidentiary hearings and adaptive improvement.

Given unlimited computational capacity this problem could then be treated as one of dynamic, constrained optimisation set in a time-varying environment. The governor could continuously re-compute and execute policies over some multi-objective cost function, adaptively learning from its experience. Unfortunately, there are limits to the processing capability available, tasks relating to the autonomous application of lethal force are likely to be highly time-constrained, the constraints and the solutions to any cost functions will only provide good approximations to the specific parameters under consideration, and the optimisation will depend upon the real world variables being optimised. As a result, the best we will likely achieve will be an approximate solution to a precise set of principles.

Furthermore, given the modular architecture proposed, each element of the system would need to detect that it has failed and inform any other components. These requirements are strongly linked to the need for the system to determine conditions under which the prescribed directives are unachievable, either within a required time frame, the broader capability framework of the sensors, weapons, etc, or for some programming shortfall in the ethical moderator. By having this

level of self-awareness, and notifying users of such limitations, the human can then over-ride the moderator, as appropriate.

Under these circumstances, or if the ethical moderator has declined to use lethal force for any other reason, the option for operator over-ride should result in responsibility for any actions once again resting with the human. The complication here is that, unless the reasons are appropriately explained to the user, they may not be apparent. However, if the algorithm uses the sequence above this might provide a framework that could help the machine explain its reasoning to the human in terms they can understand. Whether this can be done in sufficiently timely manner for the human to make an informed judgement under the pressures of battle, however, is less clear.

It is clearly a fundamental requirement that the moderator be able to provide permissible and obligatory solutions. To do this it may attempt to recognise those that are not permissible, based on its interpretation of the ROE, its own capabilities, etc. As a result, it is possible that the moderator could work in this domain, defining solutions that it cannot achieve in order to determine those that it can.

However, thus far, we have cast the solution as a binary problem for which the actions are either permissible or not. Relative to the capabilities of a specific UVS there may be areas of grey where solutions are difficult rather than impossible, and that further human judgement needs to be applied. One solution to such challenges is to attempt to explicitly compute the ethical cost functions that are defined in UVS response-behaviour space. These plans may then be treated as input to reinforced learning techniques that can then learn by physically interacting with the ethical deliberations. However, even though the early instantiation of such behaviour bounds are likely to be relatively straightforward and well-understood, the abstract principles of ethical reasoning and the absence of first order logic will likely lead to unknown or un-modelled factors. Consequently, the evaluation function is unlikely to be 'crisp' and the application of any reinforced learning techniques complex.

A decision must also be made as to how to structure the ethical governance algorithms. As with UVS planning systems, the computational advantages of using partial solutions are attractive, but present difficulties. For example, if the process is interrupted a partial solution algorithm may not provide a feasible strategy, whereas complete solution approaches should always be able to provide at least one feasible plan; and it is likely to be imperative that a timely solution be available at all times.

As with mission-planners, once operating most ethical moderators cannot simply stop to compute a new plan every time the circumstances demand it. Consequently, at a most basic level, it is also essential that the moderator prescribe solutions that have 'lives' longer than the situational awareness predictions of the UVS. If this is achievable the UVS will at least operate ethically within its environment while other plans are computed. In order to achieve this, however, moderation cycles must be performed concurrently with normal system operation and it is impossible to know a priori what update rates will be required for such systems. As a result, if possible, we should incorporate a degree of adaptive or

reinforced learning into the prediction component of the adaptive algorithms to allow them to determine their own update requirements. That is, we should allow them to learn from changes in the environment and their own deliberations in order to accommodate a better sampling frequency of sensor inputs and moderation outputs.

However, as previously mentioned, when applying any reinforced learning techniques to ethical moderators, caution will need to be exercised as the UVS will be trained against data selected by their owners and operators. As a result, while clearly aimed at enhancing the capabilities of a UVS, the algorithm could be immersed into environments that train it to over-ride some pre-ordained taboo on the basis of some learnt behaviour.

Concluding Remarks

Both now and historically – and in particular since the industrial revolution – there has been a great deal of interest in building autonomous military vehicles. This has met with mixed success, both programmatically and technologically. Nevertheless, militaries have persisted for the self-evident reason that UVS offer increased levels of force protection, reduced workload, and extended reach and access for their warfighters into areas of the battlefield intentionally denied them.

In the past few decades in particular, however, there has been an explosion in the miniaturisation, maturation, diversity and commercial availability of components and systems engineering techniques required to successfully automate a UVS. To date these advances have come largely in the areas of communications, sensing, signal processing, data and information fusion, systems engineering and integration, launch and recovery, human factors, platform, aero/hydro-dynamics, mobility, collision avoidance, mission planning and re-planning, propulsion, size, and energy storage. However, more recent advances now offer the prospect of levels of autonomy and functionality which for the first time could bring about agile, versatile, persistent, reliable and lethal autonomy with levels of robustness and survivability that could cope with some of the rigours of modern warfare.

It is acknowledged that progress in several scientific areas is still required before we achieve the generic requirements of such systems, which include persistence, low cost, stealth, ready deploy and retrieve-ability; the capacity to detect, locate, track, identify and engage targets autonomously; the ability to gather, disseminate and act on several types of information; the capacity to be networked together and to higher-value, manned assets; individual platform and sensor elements that can self-organise; and, systems that do not impose significant risk or burden upon the operators.

In particular, some of the higher-order challenges, such as systems and software VV&A, contextual decision-making, autonomous perception, dynamic planning (and plan repair) in complex and cluttered environments, and multi-UVS systems collaboration and scalability will require considerably more research. Additionally, even when these ‘technological’ matters are resolved there are many other force integration issues that will still need to be determined, such as how these systems can function reliably within an effects-based and network centric battlespace or with complex, multi-levelled, geographically dispersed, and multi-faceted human interaction within tactically changing command and control environments.

Clearly, therefore, UVS will not provide ‘bloodless battlefields’ or ‘press-button wars’ any time soon; nor will they provide solutions to every capability

challenge faced by today's defence force planners. They will, however, probably have a transformational impact on military forces and could play a key role in many future force structures. As a result, many tasks have already been identified that UVS could undertake more readily than humans, particularly in regard to in their capacity to carry out dangerous, repetitive and mechanically-oriented tasks, which will free the warfighters for other missions. Furthermore, in the short term, UVS are particularly well-suited to well-structured and uncluttered environments and tasks where system or mission failure has little or no impact on humans.

Even if the requisite levels of sophistication and functionality are achieved, however, it is not clear that the value proposition for autonomous UVS has yet been made in quantifiable terms that executive decision-makers and users can feel comfortable with; or that autonomous UVS are not exploitable due to their primitive levels of creativity and innovation; or that we understand the optimal human-UVS mix as a function of role, systems, and mission complexity or scale to counter these potential vulnerabilities.

Nevertheless, given the rate of technology progression, interest and investment, serious issues are raised in relation to the development and use of such systems, particularly in terms of future capability planning, organisational structure, procurement, training, development of doctrine and policy. In particular, it seems likely that the defence procurement and planning cycles will continue to struggle to capitalise on the promising technological opportunities and novel concepts of operation that are emerging as a result of this developmental tempo, and that this will continue to tax the symbiotic relationship that exists between the technology, its assessment and capability exploitation, and the requirements for a stable and appropriate legal framework.

Within a timeframe commensurate with the 'soft end' of the planning cycles of most defence forces, it does not seem an unreasonable prediction that within a decade we shall see a fully autonomous unmanned combat vehicle in service. As then occurred with cruise missiles, tanks, aircraft, radar and the Global Positioning System, where within a decade of their first use such systems became vital components of military consideration, it seems likely that UVS will then quickly play a dominant role in defence force planning.

Furthermore, as military UVS are only one of many fields accelerated by progress in this area and there is a synergistic relationship between the science of robotics more generally and the development of military UVS, once established fully autonomous UVS are likely to advance very rapidly because as soon as primitive techniques are developed and stable, they will be used in a wide variety of environments and circumstances. They may then be copied and run on smaller, cheaper processors with any lessons derived from (say) adaptive learning techniques instantly transportable between UVS and across domains, without the need for these lessons to be re-learned. Using Moore's law as a guide we can then assume that there will soon be a much larger number of autonomous UVS capable of processing and interpreting data many times faster than the current set of operationally deployed and relatively primitive machines. Once this level of sophistication is achieved, however, UVS may become independent agents that

are capable of initiating and making their own plans. At which point they are no longer really just tools of the warfighter.

However, before we get too carried away, let us conclude by framing the context of our aspirations with the following poignant anecdote. When programmed to present on the Australian Army's future requirement for robots Brigadier Steve Quinn, CSC, the then Director General Land Development,¹²⁷ gave a particularly short presentation.

“The long term Australian Army capability requirement for robots is relatively simple to articulate: they should be able to ‘Seek out and close with the enemy, kill or capture him, seize and hold ground, repel attack by day or by night, regardless of season, weather or terrain.’ When your robots can do that, come and see me!”

He had simply quoted the Role of Infantry and felt that no further elaboration was necessary or warranted. That is the challenge for autonomous military vehicles.

¹²⁷ He is currently the Chief, Land Operations Division at the Australian Department of Defence's Defence Science and Technology Organisation (DSTO).

References

- [1] Al-Majeed, W.: Electronic Agents & Legal Personality. In: BILETA Annual Conference, Hertfordshire (2007)
- [2] Alberts, D., Garstka, J.J., Stein, F.P.: Network Centric Warfare: Developing & Leveraging Information Superiority, C4ISR Cooperative Research Program (CCRP), <http://www.dodccrp.org/>
- [3] Alberts, D.S., Hayes, R.E.: Code of Best Practice for Experimentation, Washington, DC, Dept. of Defence, CCRP (2002)
- [4] Albus, J.S.: Brains, Behaviour and Robotics. BYTE/McGraw-Hill, New York (1981)
- [5] Albus, J.S.: Outline for a Theory of Intelligence. IEEE Transactions on Systems, Man and Cybernetics 21(3) (May/June 1991)
- [6] Albus, J.S., Meystel, A.M.: Engineering of the Mind: An Introduction to the Science of Intelligent Systems. Series on Intelligent Systems. John Wiley & Sons, NY (2001)
- [7] Allen, C., Wallach, W.: Why Machine Ethics. IEEE Journal of Intelligent Systems, 12–17 (July/August 2006)
- [8] Allen, S.C., Jones, C.R., Guerlain, S.: The Effect of Command & Control Team Structure on Ability to Quickly & Accurately Retarget UAVs. In: Proc. 2nd Conference on Autonomous Intelligent Networked Systems, Arlington, VA (2004)
- [9] Allen, T., Hill, A., Underwood, J., Scheduling, S.: Dynamic Path Planning with Multi-Agent Data Fusion – The Hierarchical Re-planner. In: Proc. International Conference on Robotics & Automation (ICRA), Kobe, Japan, May 12-17 (2009)
- [10] Allen, T., Underwood, J., Scheduling, S.: A path planning system for autonomous ground vehicles operating in unstructured dynamic environments. In: Australasian Conference on Robotics and Automation (2007)
- [11] Allen, T., Hill, A., Underwood, J., Scheduling, S.: Dynamic path planning with multi-agent data fusion - the parallel hierarchical re-planner system. Submitted to International Conference on Robotics and Automation (2009)
- [12] Allen, T., Widdison, R.: Can Computers Make Contracts. Harvard Journal of Law 9(1), 25–52 (1996)
- [13] Allison, S.: Options for DSTO Research in Automation of the Battlespace, Internal DSTO Discussion Paper (May 2003)
- [14] Anderson, M., et al.: Towards Machine Ethics. In: AAAI 2004 Workshop on Agent Organisations: Theory and Practice, San Jose, CA (July 2004)
- [15] Anderson, M., Anderson, S., Armen, C.: Towards Machine Ethics: Implementing Two Action-based Ethical Theories. In: 2005 AAAI Fall Symposium on Machine Ethics, AAAI Technical Report FS-05-06, pp. 1–7 (2005)

- [16] Anderson, M., et al.: An Approach to Computer Ethics. *IEEE Journal of Intelligent Systems*, 56–63 (July/August 2006)
- [17] Anderson, S.: *We are Our Values, Questioning Matters*, Mayfield (1999)
- [18] Arbib, M., Fellous, J.-M.: Emotions from brain to Robot. *TRENDS in Cognitive Sciences* 8(12) (December 2004)
- [19] Arkin, R.: *Behavior Based Robotics*. MIT Press, Cambridge (1998)
- [20] Arkin, R., Collins, T.R., Endo, T.: Tactical Mobile Robot Mission Specification and Execution. In: *Mobile Robots XIV*, Boston MA (September 1999)
- [21] Arkin, R., Moshkina, L.: Governing Lethal Behaviour: An Ethical Stance. In: *Proc. IEEE Int. Symposium on technology & Society*, Las Vegas, NV (June 2007)
- [22] Arkin, R.: *Governing Lethal Behaviour*, Technical Report GIT-GVU-07-11, College of Computing, Georgia Institute of Technology (2007)
- [23] Arkin, R.: *Governing lethal behaviour in autonomous robots*. Chapman & Hall, Boca Raton (2009)
- [24] Asaro, P.: Robots and Responsibility from a Legal Perspective. In: *Proc. IEEE Conf. Robotics & Automation, Workshop on Robot Ethics*, Rome (April 2007)
- [25] Asaro, P.: *How Just Could a Robot War Be?*, Umeå University, Umeå, Sweden, May 18 (2007)
- [26] Ashleigh, M.J., Nadhakumar, J.: Trust and Technologies: Implications for IT Supported Work Practices. In: *Proc. ECIS Conf.*, Gdansk, Poland, June 6-8 (2002)
- [27] Asimov, I.: *I, Robot*. Gnome Press (1950)
- [28] Australian Transportation Safety Board (ATSB), *Limitations of the See-and-Avoid Principle*, Research Report (1991) ISBN 0 642 16089 9
- [29] Australian Defence Doctrine Publication 06.4: *Law of Armed Conflict*, May 11 (2006), <http://defweb.cbr.defence.gov.au/adfwc>
- [30] Bailey, T., Nebot, E.: Localisation in Large Unstructured Environments. *Int. J. Robotics Autonomous Systems* 23(4-5), 449–472 (2004)
- [31] Balch, T., Arkin, R.: Behavior-based Formation Control for Multi-Robot Teams. *IEEE Trans. on Robotics & Automation* 14(6) (December 1998)
- [32] Baker, C., Finn, R.A.: *Operational Impact of Jamming GPS Guided Munitions*, DSTO Research Report, DSTO-RR-1129 (1998)
- [33] Barfield, W.: *Intellectual Property Rights in Virtual Environments: Considering the Rights of Owners, Programmers and Virtual Avatars*, BePress Legal Series, Working Paper 861 (2005), <http://law.bepress.com/expresso/eps/861>
- [34] Bar-Yam, Y.: *Complexity of Military Conflict: Multi-Scale Complex Systems Analysis of Littoral Warfare*, Report for Contract F30602-02-C-0158 (April 2003)
- [35] Bar-Yam, Y.: *Dynamics of Complex Systems*, Perseus (1997)
- [36] Bar-Yam, Y.: *Unifying Principles in Complex Systems*. In: Roco, M., Bainbridge, W. (eds.) *Converging Technology for Improving Human Performance* (2003)
- [37] Basu, A., Bozga, M., Sifakis, J.: Modeling Heterogeneous Real-time Components in BIP. In: *Proceedings of the Fourth IEEE International Conference on Software Engineering and Formal Methods, SEFM 2006* (2006)
- [38] Bay, H., et al.: Speeded-Up Robust Features (SURF). *Computer Vision and Image Understanding* 110(3), 346–359 (2008)
- [39] Benjamin, M.R., Curio, J.A.: COLREGS-Based Navigation of Autonomous Marine Vehicles. In: *IEEE/OES, AUV 2004, Sebasco Harbour, Maine* (2004)

- [40] Benjamin, M.R., Curcio, J.A., Leonard, J., Newman, P.: Navigation of Unmanned Marine Vehicles in Accordance with the Rules of the Road. In: International Conference on Robotics and Automation (ICRA), Orlando, FL (May 2006)
- [41] Benjamin, M.R., Curcio, J.A., Leonard, J., Newman, P.: Protocol-Based COLREGS Collision Avoidance Navigation Between Unmanned Marine Surface Craft. *Journal of Field Robotics* 23(5) (May 2006)
- [42] Berthold, K.P.H., Brian, G.S.: Determining Optical Flow. Massachusetts Institute of Technology (1980)
- [43] Bertsekas, D.P.: Auction Algorithms, *Encyclopaedia Optimization*. Kluwer, Dordrecht (2001)
- [44] Blackburn, M.R., Laird, R.T., Everett, H.R.: Unmanned Ground Vehicle (UGV) Lessons Learned, SPAWAR Technical Report 1869 (November 2001)
- [45] Boff, K.R., Kaufman, L., Thomas, J.P.: Handbook of perception and Human Performance. John Wiley & Sons, NY (1988)
- [46] Borenstein, J.: The Ethics of Autonomous Military Robots. *Studies in Ethics, Technology & Law* 2(1) (2008)
- [47] Bostrom, N.: How Long before Super Intelligence? *Linguistic & Philosophical Investigations* (March 2006)
- [48] Bourgault, F., Goktogan, A., Furukawa, T., Durrant-Whyte, H.: Coordinated search of a lost target in a Bayesian world. *J. Advanced Robotics*, 187–195 (2004)
- [49] Bowen, G., MacKenzie, S.: Autonomous Collaborative Unmanned vehicles: Technical Drivers & Constraints, Defence R&D Canada, DRDC CR-2003-003 (2003)
- [50] Bowen, J., Stavridou, V.: Safety-Critical Systems, Formal Methods and Standards. *Software Engineering Journal* (December 1992)
- [51] Bowyer, R., Bogner, R.: Cooperative Behaviour in Multi-Sensor Systems. In: Proc. of 6th Int. Conf. on Neural Information Processing. IEEE, Perth (1999)
- [52] Boyd, J.: Organic design for Command & Control (May 1987), <http://www.d-n-i.net>
- [53] Braitenberg, V.: *Vehicles: Experiments in Synthetic Psychology*. MIT Press, Cambridge (1984)
- [54] Bruemmer, D.J., et al.: Shared Understanding of Collaborative Control. *IEEE Transactions on Systems, Man & Cybernetics, - Part A Systems & Humans* 35(4) (July 2005)
- [55] Brjngsjord, S., Arkoudas, K., Bello, P.: Ethical Robots: The Future Can Heed Us. *IEEE Journal of Intelligent Systems*, 12–17 (July/August 2006)
- [56] Broadbent, D.E.: *Perception & Communication*. Pergamon Press, Oxford (1958)
- [57] Brooks, R.A.: A Robust Layered Control System for Mobile Robots. *IEEE Journal of Robotics and Automation* (April 1986)
- [58] Brooks, R.A.: *Cambrian Intelligence: The Early History of the New AI*. MIT Press, Cambridge (1999)
- [59] Brooks, R.A.: *The Future of Flesh & Machines: How Robots will Change Us*. Pantheon Books, Random House (2002)
- [60] Brown, E.D., Gaskell, N.J.J.: Report on the Law relating to Autonomous Underwater Vehicles. Report commissioned by Southampton Oceanography Centre, vol. 1 & 2. Published by Society for Underwater Technology (1999)
- [61] Brown, K., et al.: A Demonstration of Distributed Sensing & Situational Awareness. In: Land Warfare Conference, Brisbane (October 2006)

- [62] Bruemmer, D., et al.: How to Trust Robots Further than we can Throw Them. In: Proc. CHI Conference on Human Factors in Computing Systems, Vienna (April 2004)
- [63] Burridge, B.: Post-Modern Warfighting with Unmanned Vehicle Systems: Esoteric Chimera or essential Capability. *RUSI Journal* 150(5) (October 2005)
- [64] Canning, J.: Safety Implications of Unmanned Systems. In: Proc. AUVSI Conference, Orlando, FL, August 29-31 (2006)
- [65] Canning, J.: A Concept of Operations for Autonomous Weapons. In: NDIA Conference on Disruptive Technologies, Washington, D.C., September 6-7 (2006)
- [66] Chieh-Chih, W., et al.: Simultaneous Localization, Mapping and Moving Object Tracking. *Int. J. Rob. Res.* 26(9), 889–916 (2007)
- [67] Chieh-Chih, W., Thorpe, C., Thrun, S.: Online Simultaneous Localisation and mapping with detection & Tracking of Moving Objects: Theory & results from a Vehicle in Crowded Urban Areas. In: Proc IEEE Conference on Robotics & Automation (ICRA) (September 2003)
- [68] Chopra, S., White, L.: Artificial Agents – Personhood in Law and Philosophy. In: European Conference on Artificial Intelligence, pp. 635–639. IOS Press, Amsterdam (2004)
- [69] Cohen, M., Parasuranman, R., Freeman, J.: Trust in decision aids: a model and its training implications, Tech. Report USAATCOM TR-97-D-4 (1999)
- [70] Collins, T.R., et al.: Field results for tactical Mobile Robot Missions. In: Unmanned Systems 2000, Orlando, FL (July 2000)
- [71] Commander's Handbook on the Law of Naval Operations, NWP 1-14M
- [72] Craig, W.R.: Flocks, herds and schools: A distributed behavioral model. In: Proc. 14th Annual Conference on Computer Graphics and Interactive Techniques (1987)
- [73] Cummings, M.L.: Supervisory Control Challenges in NCW, University Virginia, Doctoral Dissertation (2003)
- [74] Cummings, M.L.: Human Supervisory Control of Swarming Networks. In: Proc. 2nd Conference Autonomous Intelligent Networked Systems, Arlington, VA (2004)
- [75] Cummings, M.L.: Human Supervisory Control Challenges in Network Centric Operations. In: Proceedings Unmanned Vehicle Systems, Canada (2005)
- [76] Cummings, M.L., et al.: Predicting Operator Capacity for Supervisory Control of Multiple UAVs. In: Innovations in Intelligent Machines. Springer, Heidelberg (2009)
- [77] David, G.L.: Distinctive Image Features from Scale-Invariant Keypoints. *Int. J. Comput. Vision* 60(2), 91–110 (2004)
- [78] Davidson, E.J.: Evaluation Methodology: The Nuts & Bolts of Sound Evaluation. Sage Publications, Thousand Oaks (2005)
- [79] Davison, A.J., et al.: MonoSLAM: Real-Time Single Camera SLAM. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29(6), 1052–1067 (2007)
- [80] Davis, P.K., Blumenthal, D.: The Base of Sand, A White paper on the State of Military Combat Modelling, RAND Note N-3148-OSD/DARPA (1991)
- [81] Deans, M.: Bearings-only localization and mapping. Robotics institute Carnegie Mellon University, Pittsburgh, PA (2005)
- [82] Delude, C.M.: MIT Neuroscientists See design flaws in Computer Vision Tests, McGovern Institute (January 2008)
- [83] Dempster, A.G.: Dilution of precision in angle-of-arrival positioning systems. *Electronics Letters* 42(5), 291–292 (2006)

- [84] Dias, M.B., Stenz, A.: A Free Market Architecture for distributed Control of a Multi-Robot System. In: 6th International Conference on Intelligent Autonomous Systems, Venice, Italy, July 2000, pp. 115–122 (2000)
- [85] Dissanayake, M., et al.: A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Transactions on Robotics and Automation* 17(3), 229–241 (2001)
- [86] Dixon, S.R., Wickens, C.D.: Imperfect Automation in UAV Flight Control, Institute of Aviation, AFHD Technical Report, AFHD-03-17/MAAD-03-2 (2003)
- [87] Dixon, S.R., Wickens, C.D.: Control of Multiple UAV's - A Workload Analysis. In: Proceedings 12th International Symposium Aviation Psychology, Dayton, OH (2003)
- [88] Drake, S.P., Brown, K.L., Finn, R.A.: The Autonomous Control of Uninhabited Aerial Vehicles for the Passive Location of Radars, DSTO Research Report (2005)
- [89] Draper, J.V., Kaber, D.B.: Human-Robot Interaction. In: Mital, A., et al. (eds.) *Industrial & Occupational Ergonomics: Users' Encyclopaedia* (1999)
- [90] Durrant-Whyte, H.F., Clark, S., Dissanayake, M.W.M.G., Newman, P., Csorba, M.: An experimental and theoretical investigation into simultaneous localisation and map-building (SLAM). In: Proc. 6th Sym. on Experimental Robotics, Sydney (1999)
- [91] Durrant-Whyte, H.: Autonomous Land Vehicles, ACFR Internal Report (2004)
- [92] Dyson, G.: *Darwin Amongst the Machines: The Evolution of Global Intelligence*. Addison-Wesley Publishing Corp., Reading (1997)
- [93] Ebert, C., Jones, C.: *Embedded Software: Facts, Figures & Future*, vol. 42(4). IEEE Computer Society, Los Alamitos (2009)
- [94] Elliott, E.D., Eagleson, R.: Web-based teleoperated Systems using EAI. In: Proc. IEEE Conf. on Systems, Man and Cybernetics, October 12-15, vol. 1, pp. 749–754 (1997)
- [95] Endsley, M.R., Kaber, D.B.: Level of Automation Effects on Performance, Situational Awareness, and Workload in a Dynamic Control Task. *J. Ergonomics* 42(3), 462–492 (1999)
- [96] Endsley, M.R., Jones, W.M.: A Model of Inter- & Intra-Team Situational Awareness. In: Endsley, M.R. (ed.) *New Trends in Cooperative Activities*. Human Factors & Ergonomics Society (2001)
- [97] EUROCONTROL, Specifications for the Use of UAVs as Operational Air Traffic Outside Segregated Airspace, EUROCONTROL-SPEC-0102 (July 2007)
- [98] Fikes, R., Nilsson, N.: STRIPS – A New Approach to the Application of Theorem Proving to Problem Solving. *J. Artificial Intelligence* 2(3-4), 189–208 (1971)
- [99] Finn, R.A.: Automation of the Battlespace Initiative, Internal DSTO Paper (2004)
- [100] Finn, R.A., Chalmers, G.: Operations Analysis Methodology for Force Level Electronic Warfare, DSTO-RR-0298, DSTO Research Report (2001)
- [101] Finn, R.A.: Spectral Dominance – A Way Ahead, DSTO Research Report (2000)
- [102] Finn, R.A., Brown, K., Lindsay, A.: Miniature UAVs & Future Electronic Warfare. In: Proc. Land Warfare Conference, Brisbane (October 2002)
- [103] Finn, R.A., et al.: The Australian UAV Development Project. In: Proceedings of the Australian International Aerospace Congress, Melbourne (2005)
- [104] Finn, R.A.: Legal Considerations for the Weaponisation of UGVs. *J. Intelligent Defence Systems* 1 (September 2008)
- [105] Finn, R.A., Pincombe, Chalmers, G.: Force Level Electronic Warfare for the ADF. In: Proc. of the SPIE Conference, Orlando, USA (April 2001)

- [106] Finn, R.A., et al.: Automation of the Battlespace: An Integrated Air-Ground, Multi-Vehicle, Multi-Sensor, Autonomous UVS Demonstration. In: Proc. AUVSI Conference, Anaheim (2004)
- [107] Finn, R.A.: Stereoscopic See-and-Avoid for Small UAV's, TTCP Research Paper, TTCP-AER-TP6 (April 2008)
- [108] Finn, R.A., Drake, S.: Intelligent Control of Multi-UAV Cooperatives. *J. Intelligent Decision-Making Technology* 1(1) (2008)
- [109] Franklin, S., Finn, A., Pattison, J., Jain, L.: Towards Scene Understanding Using a Cooperative of Mobile Robots. *J. Intelligent Systems* 20(issue TBD) (accepted, 2009)
- [110] Freedman, S.T., Adams, J.A.: Inherent Components of UV Situation Awareness. Dept. Electrical Engineering, Vanderbilt University, Nashville, TN, 1-4244-0991-8/07 (2007), <http://www.ieeexplore.ieee.org/iel5/4413560/4413561/04414129.pdf>
- [111] Freedy, A., DeVisser, E., Weltman, G., Coeyman, N.: Measurement of Trust in Human-Robot Collaboration. In: Proceedings International Conference on Collaborative Technologies and Systems. IEEE, Los Alamitos (2007)
- [112] Frost, Sullivan: Unmanned Systems Markets, Report N384-16 (April 2008)
- [113] GAO Report, Evaluation of the Air Campaign, GAO NSIAD 97-134 (1997)
- [114] Gaudin, S.: IBM's Roadrunner smashes 4-minute mile of supercomputing. *Computerworld* (June 2008)
- [115] Gat, E.: Artificial intelligence and mobile robots: case studies of successful robot systems. MIT Press, Cambridge (1998)
- [116] Gerkey, B., Mataric, M.J.: A Formal Framework for the Study of Task-Allocation in Multi-Robot Systems. *Int. Journal Robotics Research* 14 (2003)
- [117] Gibbins, D.: Ladar 3-D scene analysis - Classification update and real time considerations, University of Adelaide Research Report for DSTO (2009)
- [118] Goodrich, M., Olsen, D.: Seven Principles of Efficient Human-Robot Interaction. In: Proc. IEEE Conf. Systems, man & Cybernetics, pp. 3943-3948 (2003)
- [119] Gothard, Cory, B.M.P., Peterman, P.: Demo III processing architecture trades and preliminary design. In: Proc. SPIE, vol. 3525 (1999)
- [120] Govt. South Australia Attorney General's Department, Australian Traffic Act (SA) (December 1999), http://www.transport.sa.gov.au/publications/road_rules/index.asp
- [121] Granda, Kirkpatrick, T.M., Julien, T., Peterson, L.: The Evolutionary Role of Humans in the Human-Robot System. In: Proc. Human factors Society 34th Annual Meeting, pp. 664-668 (1990)
- [122] Grocholsky, B.: Information-Theoretic Control of Multiple Sensor Platforms, PhD Thesis, University of Sydney (2002)
- [123] Gulam, H., Lee, S.: Uninhabited Combat Aerial vehicles and the Law of Armed Conflict. *Australian Army Journal* 3(2) (Winter 2006)
- [124] Hart, S., Staveland, L.: Development of the NASA Task Load Index: Research in Empirical and Theoretical Results. In: Hancock, P., Meshkati, N. (eds.) *Human Mental Workload*. North Holland Press, Amsterdam (1988)
- [125] Helft, M.: Google Zooms In Too Close for Some, *New York Times* (June 2007), <http://www.nytimes.com/2007/06/01/technology/01private.html>

- [126] Henderson, A.: The Legal Status of Unmanned Undersea Vehicles. Naval Law Review, <http://www.jag.navy.mil/documents/NJSCourses/NLRVolume53.pdf>
- [127] Henderson, J., Overholt, J.: Robotics Strategy White Paper, Army Capabilities Integration Centre & Tank-Automotive Research & Development Engineering Centre, ACIC-TARDEC (March 2009)
- [128] Hendler, J.: Computers Play Chess, Humans Play Go. IEEE Journal of Intelligent Systems, 2–3 (July/August 2006)
- [129] Hespana, J., Prandini, M., Sastry, S.: Probabilistic Pursuit Evasion Games. In: Proc. 39th IEEE Conf. on Decision & Control, pp. 2272–2277 (2001)
- [130] Hinds, P.J., Roberts, T.L., Jones, H.: Whose Job Is It Anyway? A Study of Human-Robot Interaction. J. Human Computer Interaction 19 (2004)
- [131] Hmam, H.: Scan-Based Emitter Passive Localization. IEEE Trans. on AES 43(1) (2007)
- [132] Hoffman, G., Breazeal, C.: Anticipatory Perceptual Simulation for HRI Joint Practice. In: Proceedings of 23rd AAAI Conference on AI (2008), <http://www.aaai.org>
- [133] Hogendoorm, H.: The State of the Art in Visual Object Recognition, Master's Thesis, Helmholtz Institute, Experimental Psychology Division, Universiteit Utrecht, Utrecht, The Netherlands (August 2006)
- [134] Huang, H.M.: Autonomy Levels for Unmanned Systems (ALFUS) Framework. In: Proceedings 2006 Performance Metrics for Intelligent Systems (PerMIS) Workshop, Gaithersburg, MD, August 21-23 (2006)
- [135] Hylton, Todd (Program Manager), DARPA Nano UAV Challenge, <http://www.darpa.mil>
- [136] ICAO, Convention on International Civil Aviation (Doc 7300), Signed at Chicago on 7 December 1944 and amended by the ICAO Assembly, 9th edition, 116 p. (2006), http://www.icao.int/icao/en/m_publications.html
- [137] Intel Corporation, Intel details upcoming pew processor generations, Intel Press Release 11th (August 2008)
- [138] Jacoff, A., Messina, E., Evans, J.: Performance Evaluation of Autonomous Mobile Robots. Industrial Robot: An International Journal 29(3) (2002)
- [139] Jacoff, et al.: Robocup Rescue – Rules (2009), <http://www.robocupresuce.org>
- [140] Johnson, C.L.: Inverting the Control Ratio: Human Control of Large Autonomous Teams. In: Proc. Workshop Humans & Multi-Agent Systems, Melbourne (July 2003)
- [141] Joint Robotics Master Plan, Office of the Under Secretary of Defense (AT&L), Land Warfare & Munitions (2004), <http://www.jointrobotics.com>
- [142] Kaku, M.: The Physics of the Impossible. Anchor Books, NY (2008)
- [143] Karumanchi, S., Allen, T., Bailey, T., Scheduling, S.: Non-parametric Learning to Aid Path Planning over Slopes. In: Proc. of Robotics Science & Systems Conference, University of Washington (June/July 2009)
- [144] Karumanchi, S., Allen, T., Bailey, T., S.: Scheduling, Learning to Interpret Sensed Environmental Conditions Using Proprioceptive Feedback. Submitted to J. Autonomous Robots (2009)
- [145] Karnow, C.: Liability for Distributed Artificial Intelligence. Berkley Technical Law Journal 11, 149–174 (1996)

- [146] Kaupp, T., Makarenko, A.: Measuring Human-Robot Team Effectiveness to Determine an Appropriate Autonomy Level. In: Proc. IEEE Conf. on Robotics & Automation, Pasadena, CA, USA, May 19-23 (2008)
- [147] Kellerstrass, E.J.: Drone RPV and Aerospace Power, Air University review (September-October 1973), <http://www.airpower.maxwell.af.mil>
- [148] de Kerchove, R.: International Maritime Dictionary, 2nd edn. French & European Publications Inc. (1961)
- [149] Kerr, I.: Babes and the Californication of Commerce. University of Ottawa Law & Technology Journal, 285-324 (2004), <http://www.uoltj.ca>
- [150] Khatib, O., et al.: Vehicle/arm coordination and multiple mobile manipulator decentralized cooperation. In: Proc. of the 1996 IEEE/RSJ International Conference on Intelligent Robots and Systems 1996, IROS (1996)
- [151] Kilgore, R., Harper, K.: Mission Planning & Monitoring for Heterogenous Unmanned Vehicle teams: A Human-Centred Approach. In: Proc. AIAA Infotech@Aerospace Conference (2007)
- [152] Kim, J.R., Vidal, H., Shim, O.: A hierarchical approach to probabilistic pursuit-evasion games with unmanned ground and aerial vehicles. IEEE Transactions on Robotics & Automation 18(5) (October 2002)
- [153] Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., Osawa, E., Matsubara, H.: Robocup: A Challenging Problem for AI and Robotics. In: Kitano, H. (ed.) RoboCup 1997. LNCS (LNAI), vol. 1395. Springer, Heidelberg (1998)
- [154] Klein, G., Moon, B., Hoffman, R.: Making Sense of Sensemaking. IEEE Journal of Intelligent Systems, 18-21 (July/August 2006)
- [155] Kogut, G.T., Birchmore, F., Pacis, E.B., Ahuja, G., Sights, B., Everett, B.: Using advanced computer vision algorithms on small mobile robots. In: SPIE Proc. 6230: Unmanned Systems Technology VIII, Defense Security Symposium, Orlando, FL, April 17-20 (2006)
- [156] Kolling, A., Carpin, S.: Multirobot cooperation for surveillance of multiple moving targets - a new behavioral approach. In: Proceedings 2006 IEEE International Conference on Robotics and Automation, ICRA (2006)
- [157] Kuchar, J.K., Yang, L.C.: A Review of Conflict Detection and Resolution Modelling Methods. IEEE Transactions on Intelligent Transportation Systems 1(4) (2000)
- [158] Kuc, T.-Y., Lee, I.-J., Baek, S.-M.: Perception, reasoning and learning of multiple agent systems for robot soccer. In: IEEE Conf. on Systems, Man & Cybernetics (1999)
- [159] Kurzweil, R.: The Age of Spiritual Machines: When Computers Exceed Human Intelligence, Sydney, Allen & Unwin (1999)
- [160] Lambert, D., Scholz, J.: Ubiquitous Command and Control (UC2). In: Proceedings Information, Decision and Control Conference, Adelaide, South Australia (1999)
- [161] Lazarski, A.J.: Legal Implications of UCAVs. J. Air & Space Power (March 2001)
- [162] Lee, J., Moray, N.: Trust, self-confidence, and operators' adaptation to automation. Journal of Human-Computer Studies 40, 153-184 (1994)
- [163] Lee, J., See, K.: Trust in Automation: Designing for Appropriate Reliance. J. Human Factors 46(1) (Spring 2004)
- [164] Leonard, J., Durrant-Whyte, H.: Simultaneous Map Building and Localization for an Autonomous Mobile Robot. In: Proc. IROS 1991, IEEE/RSJ International Workshop on Intelligent Robots and Systems (1991)
- [165] Lewis, M., Polvichai, J., Sycara, K., Scerri, P.: Scaling Up Human Control for Large UAV Teams. In: Proc. HF & Ergonomics Soc. 48th Annual Meeting (2004)

- [166] Liao, E., Scerri, P., Sycara, K.: A framework for very large teams. Proc. Autonomous Agents & Multi-Agent Systems (AAMAS) Conference (2004)
- [167] Lindsey, A., Pecheur, C.: Simulation-Based Verification of Livingstone Applications. In: Proc. Workshop on Model-Checking for Dependable Software-Intensive Systems, San Francisco (June 2003)
- [168] Lopez, I., Sanz, R., Hernandez, C.: Architectural Factors for Intelligence in Autonomous Systems. In: Proc. AAAI Workshop on Evaluating Architectures for Intelligence (August 2007)
- [169] Lowe, C.: Invisible Boomerang Bot (September 2006), <http://www.defencetech.org>
- [170] Luck, M., McBurney, P., Preist, C.: Agent Technology: Enabling Next Generation Computing (2003), <http://users.ecs.soton.ac.uk/~mml/papers/al2roadmap.pdf>
- [171] MacKenzie, D., Arkin, R., Cameron, J.: Multi Agent Mission Specification and Execution. *Autonomous Robots* 4(1) (January 1997)
- [172] Marsh, L., Calbert, G., Tu, J., Gossink, D., Kwok, H.: Multi-Agent UAV Path Planning. In: Proc. Modelling and Simulation Society of Australia & New Zealand, Melbourne (December 2005)
- [173] Marsh, L., Gossink, D., Drake, S., Calbert, G.: UAV Team Formation for Emitter Geolocation. In: Proc. IDC 2007, Adelaide (February 2007)
- [174] Mataric, M.J.: Issues and Approaches in the Design of Collective Autonomous Agents. *Robotics and Automation Systems* 16, 321–331 (1995)
- [175] Matsumura, J., Steeb, R., Herbert, T.J., John Gordon, I., Glenn, R.W., Steinberg, P.: Lightning over water: Sharpening America's light forces for rapid reaction missions, Tech. Rep. MR-1196-A/OSD, RAND, Arroyo Center (2000)
- [176] Matsumura, J., Steeb, R., Herbert, T.J., John Gordon, I., Rhodes, C., Glenn, R.W., Barbero, M., Gellert, F.J., Kantar, P., Halverson, G., Cochran, R., Steinberg, P.: Exploring advanced technologies for the future combat systems program, Tech. Rep. MR-8333-OSD, RAND, Arroyo Center (2002)
- [177] Matsumura, J., Steeb, R.: Unmanned but not un-tethered: Robots on the future battlefield, Tech. Rep. Summer Review, RAND, Arroyo Center (2005)
- [178] Markram, H.: The Blue Brain Project. *Nature Reviews of Neuroscience* 7, 153–160 (2006)
- [179] McCarley, J., Wickens, C.D.: Human Factors Concerns in UAV Flight, University of Illinois Human Factors Division Technical Report AHFD-05-05/FAA-05-1
- [180] McLaren, B.: Computational models of Ethical Reasoning. *IEEE Journal of Intelligent Systems*, 2–3 (July/August 2006)
- [181] McClelland, J.: The Review of Weapons in Accordance with Article 36 of Additional Protocol I. IRRC 85(850) (June 2003)
- [182] McCarley, J.S., Wickens, C.D.: Human Factors Concerns in UAV Flight, University of Illinois HF Division Technical Report AHFD-05-05/FAA-05-1
- [183] McCarley, J.S., Wickens, C.D.: Human Factors Implications of UAVs in the National Airspace, Aviation Human Factors Division, Institute of Aviation, Technical Report AHFD-05-05/FAA-05-01 (April 2005)
- [184] McConnell, S.: *The Software project Survival Guide*. Microsoft Press (1998)
- [185] Mason, K., Howard, C.: Knowledge Intensive Interpretation of Signal Data. In: Orchard, B., Yang, C., Ali, M. (eds.) IEA/AIE 2004. LNCS (LNAI), vol. 3029, pp. 170–175. Springer, Heidelberg (2004)

- [186] Menzies, T., Pecheur, C.: Verification and Validation and Artificial Intelligence. In: *Advances in Computers*, vol. 65. Elsevier, Amsterdam (2005)
- [187] Michalewicz, Z., Schmidt, M.: Planning & AI Techniques Applicable to the UAV Pre-Mission Planning Problem, Interim Report for DSTO
- [188] Michalewicz, Z., Schmidt, M., Michalewicz, M., Chiriac, C.: *Adaptive business Intelligence*. Springer, Heidelberg (2007)
- [189] Micire, M.: Analysis of the Robotic-Assisted Search & Rescue Response to the WTC Disaster, MSc Thesis, University of South Florida (July 2002)
- [190] de Miguel, M., Briones, J.F., Silva, J.P., Alonso, A.: Model Based Integration of Safety Analysis and Development. In: *Proc. 9th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing*, Korea (April 2006)
- [191] Milgram, S.: The small world problem. *Psychology Today* 1, 61–67 (1967)
- [192] Miller, A.: Modelling Human Workload Limitations on Multiple UAV Control. In: *48th Annual Meeting of Human Factors & Ergonomics Society*, Santa Monica (2004)
- [193] Mitchell, P.M., Cummings, M.L., Sheridan, T.B.: Mitigation of Human Supervisory Control Wait Times through Automation Strategies, Technical Report, MIT Humans & Automation Laboratory, MIT, MA (2005)
- [194] Moor, J.: The Nature, Importance, and Difficulty of Machine Ethics. *IEEE Journal of Intelligent Systems*, 18–21 (July/August 2006)
- [195] Moravec, H.: *Robot: Mere Machine to Transcendent Mind*. Oxford University Press, Oxford (1998)
- [196] Moravec, H.: When will computer hardware match the human brain. *Journal of Evolution and Technology* 1 (1998)
- [197] Moreland, M., Hanselmann, T., Moran, W., Sarunic, P.: Scheduling of Passive Sensors for target Tracking, DSTO Client Research Report (2008)
- [198] Moshkina, L., Arkin, R.: Lethality & Autonomous Systems: Survey and Design Results, Georgia Institute Technology, Technical Report GIT-GVU-07-16 (2007)
- [199] Moutarlier, P., Chatila, R.: Stochastic multi-sensory data fusion for mobile robot location and environment modelling. In: *5th Int. Sym. on Robotics*, Tokyo (1989)
- [200] Murphy, R.R., Rogers, E.: Human-Robot Interaction, DARPA/NSF Final Report (2001), <http://www.csc.calpoly.edu/~erogers/HRI/HRI-report-final.html>
- [201] Nass, C., Lee, K.: Does Computer-synthesised speech manifest personality? Experimental tests of recognition, similarity-attraction and consistency-attraction. *J. Experimental Psychology – Applied* 7 (2001)
- [202] National Research Council, *Autonomous Vehicles in Support of Naval Operations*. National Academic Press (2006) ISBN-0-309-09676
- [203] NATO, Code of Best Practice for Command and Control, C4ISR Cooperative Research Program, AC/243 (Panel 7) TR/8 (1998)
- [204] Navy Unmanned Undersea Vehicle Master Plan, Dept. of Navy (March 2001)
- [205] Navy Unmanned Undersea Vehicle Master Plan, Dept. of Navy (November 2004)
- [206] Navy Unmanned Surface Vehicle Master Plan, Dept. of Navy (July 2007)
- [207] Nettleton, Durrant-Whyte, H., Sukkarieh, S.: A Robust Architecture for Decentralised Data Fusion. In: *Proc. SPIE, Sensor Fusion: Architectures, Algorithms, & Applications*, Orlando, FL (2000)
- [208] Newell, A.: The Knowledge Level. *Artificial Intelligence* 18(1), 87–127 (1983)

- [209] Newell, A., Simon, H.: GPS: A Program that Simulates Human Thought. In: Feigenbaum, Feldman (eds.) *Computers and Thought*. McGraw-Hill, NY (1963)
- [210] New South Wales State Government, *The Work Cover Website* (2007), <http://www.workcover.nsw.gov.au/default.htm> (last accessed: 20-December-2007)
- [211] Nourbakhsh, I.R., et al.: *Human-Robot Teaming for Search & Rescue*. In: *Pervasive Computing*. IEEE Computer Society, Los Alamitos (January-March 2005)
- [212] Nilsson, N.J.: *Principles of Artificial Intelligence*. Tioga Publishing Co. (1980)
- [213] Nickerson, J.V., Reilly, R.R.: *A Model for Investigating the Effects of Machine Autonomy on Human Behaviour*. In: *Proc. 37th Int. Conf. on Sys. Sci., Hawaii* (April 2004)
- [214] O'Day, S., et al.: *Metrics for Intelligent Autonomy*. In: *Proc. Performance Metrics for Intelligent Systems, PerMIS* (2004)
- [215] Olsen, S.R., Wood, D.B.: *Fan Out: Measuring Human Control of Multiple Robots*. In: *Proc. CHI 2004, Vienna, Austria* (2004)
- [216] Olsen, S.R., Goodrich, M.A.: *Metrics for Evaluating Human-Robot Interactions*. In: *Proc. Performance Metrics for Intelligent Systems, Gaithersburg, MD* (2003)
- [217] Office of the UN High Commissioner for Human Rights, *Protocol Additional to the Geneva Conventions of 12 August 1949 and relating to the protection of Victims of International Armed Conflicts (Protocol 1), Adopted 8 June (1977)*, <http://www2.ohchr.org/english/law/protocol11.htm>
- [218] Office of Naval Research, *Technology Readiness Level*, <http://www.onr.navy.mil/>
- [219] Parasuraman, R., Riley, V.: *Humans & Automation: Use, Misuse, Disuse, Abuse*. *Journal of Human Factors* 39(2), 230–253 (1997)
- [220] Parasuraman, R., Barnes, M., Consenzo, K.: *Adaptive Automation for Human Robotic Teaming - The Future Command & Control Systems*. *International Journal of Command & Control* 1(2) (2008)
- [221] Parker, L.E.: *Current State of the Art in Distributed Autonomous Mobile Robotics*. In: *Distributed Autonomous Robotic Systems, Tokyo, vol. 4*. Springer, Heidelberg (2000)
- [222] Pecheur, C., Simmons, R.: *Formal Verification of Autonomy Models: From Livingstone to SMV*. In: Rouff, C., et al. (eds.) *Agent Technology from a Formal Perspective - NASA Monographs in Systems & Software Eng.* Springer, Heidelberg (2006)
- [223] Pecheur, C.: *Verification of Intelligent Control Software*. In: *Proceedings ASTRA 2006, Noordwijk, Netherlands* (November 2006)
- [224] Penrose, R.: *The Emperor's New Mind – Concerning Minds and the Laws of Physics*. Oxford University Press, Oxford (1989)
- [225] Perla, P.D.: *The Art of Wargaming*. Naval Institute Press (1990)
- [226] Peter, G., Peter, E., Markus, V.: *Simultaneous Motion and Structure Estimation by Fusion of Inertial and Vision Data*. *Int. J. Rob. Res.* 26(6), 591–605 (2007)
- [227] Picard, R.: *Affective Computing*, MIT Media Lab Perceptual Computing Section Technical Report No. 321 (1997), <http://www.media.mit.edu/~picard/>
- [228] Powers, T.: *Prospects for a Kantian Machine*. *IEEE Journal of Intelligent Systems*, 46–51 (July/August 2006)
- [229] Quintana, E.: *The Ethics and Legal Implications of Military Unmanned Vehicles*, Occasional Paper for the British Computer Society (February 2008)
- [230] RAND Corporation, *Swarming and the Future of Conflict* (2000)

- [231] Rantanen, E., Nunes, A.: Taxonomies of Measures in Air Traffic Control Research. In: Proc. International Symposium on Aviation Psychology (2003)
- [232] Ratches, et al.: Aided & Automatic Target Recognition Based upon Sensory Inputs from Image Forming Systems. IEEE Trans. on Pattern Analysis & Machine Intelligence 19(9) (1997)
- [233] Rehmann, A.J.: Guide to Human Performance Measurements and Crew Requirements, Dept. transport Technical Report, DOT/FAA/CT-TN95-49 (1995)
- [234] Report to Congress: Development & Utilization of Robotics and Unmanned Ground Vehicles, Office of the Under Secretary of Defense (Acquisition Technology & Logistics), Land Warfare & Munitions, Joint Ground Robotics Enterprise (2006)
- [235] Robertson, S.: Can Military UGV's Save Lives, and, If So, at What Cost, ACFR Internal Report (2006)
- [236] Rodriguez, G., Weisbin, C.: A New method to Evaluate human-robot systems performance. Autonomous Robots 14, 165–178 (2003)
- [237] Rose, M.F.: (chair) Technology Development for Army Unmanned Ground Vehicles, Committee on Army Unmanned Ground Vehicle Technology, Board on Army Science & Technology, National Research Council (2002)
- [238] Roumeliotis, S., Bekey, G.: Collective Localisation: a distributed Kalman Filter Approach. In: Proc. IEEE ICRA, pp. 2958–2965 (2000)
- [239] Rustad, M.L., Koenig, T.H.: Cybertorts and Legal Lag: An Empirical Aanalysis. Southern Californian Law Journal 13(1) (2004)
- [240] Rubel, R.: Using Wargames for Command & Control Experimentation, Research Report, Wargaming Department, Center for Naval Wargaming Studies (June 2003)
- [241] Sato, R.: Daily Galaxy, March 26 (2008), <http://www.dailygalaxy.com>
- [242] Scerri, P., Pynadath, D., Tambe, M.: Towards Adjustable Autonomy in the Real World. J. Artificial Intelligence 17, 171–228 (2002)
- [243] Scerri, P., Pynadath, D., Tambe, M.: Adjustable Autonomy in the Context of Coordination. In: Proc. AIAA 3rd Unmanned Unlimited Technical Conference (2004)
- [244] Scheduling, S., Durrant-Whyte, H., Robertson, S.: An Assessment of Defence Unmanned Ground Vehicle Systems and Technology. DSTO-CAS COE in Defence Autonomous & Unmanned Ground Vehicles Research Report (2008)
- [245] Scheduling, S.: Unmanned Ground Vehicles, DSTO-ACFR University of Sydney Centre of Expertise Report, DSTO Project 1034930 (2004)
- [246] Schmitt, M.: Impact of High & Low-Tech Warfare on the Principle of Distinction (November 2003), <http://www.michaelschmitt.org/images/Hi%20Tech%20Low%20Tech%20Schmitt%20formatted.doc>
- [247] Schneidewind, N.F.: Reliability Modelling for Safety-Critical Software. Proc. IEEE Trans. on Reliability 46(1) (March 1997)
- [248] Sciortino, J.C., et al.: The Role of AI & CAS in Network-Centric Warfare. In: 45th Joint EW Conference, Naval Postgraduate School, Monterey, CA (May 2000)
- [249] Shai, A., Amnon, S.: Trajectory Triangulation: 3D Reconstruction of Moving Points from a Monocular Image Sequence. IEEE Trans. Pattern Anal. Mach. Intell. 22(4), 348–357 (2000)
- [250] Sheridan, T.B., Verplank, W.: Human & Computer Control of Undersea Teleoperators, Cambridge, MA, Man-Machine Systems Laboratory, Dept. Mechanical Engineering, MIT (1978)

- [251] Sheridan, T.B.: *Tele-robotics, Automation, & Human Supervisory Control*. MIT Press, Cambridge (1992)
- [252] Singer, P.W.: *Wired for War: the Robotics Revolution of the Twenty First Century*. Penguin Press (2009)
- [253] Showalter, S.: The Legal status of Autonomous Underwater Vehicles. *Marine Technology Society Journal* 38(1) (Spring 2004)
- [254] Smith, R.C., Cheesman, P.: On the representation of spatial uncertainty. *International Journal of Robotics Research* 5(4), 56–68 (1987)
- [255] Solum, L.B.: *Legal Personhood for Artificial Intelligence*. North Carolina Law Review (April 1992)
- [256] Sparrow, R.: Killer Robots. *Journal of Applied Philosophy* 24(1) (2007)
- [257] Standards Australia, *Industrial robot systems – safe design and usage*, Technical Report 2939–1987, Standards Association of Australia (1987)
- [258] Steeb, R., Matsumura, J., Steinberg, P., Herbert, T., Kantar, P., Bogue, P.: Examining the army’s future warrior – force-on-force simulation of candidate technologies, Tech. Rep. MG-140, RAND, Arroyo Center (2004)
- [259] Steinfeld, A., et al.: Common Metrics for Human-Robot Interaction. In: *Proc. Human-Robot Interaction 2006*, Salt lake City, Utah (March 2006)
- [260] Stone, P., Veloso, M.: Multi-Agent Systems: A Survey from a machine learning perspective. *Autonomous Robots* 8(3), 345–383 (2000)
- [261] Sukkarieh, S., Durrant-Whyte, H.F.: Toward the development of SLAM for an unmanned air vehicle. In: *FSR 2001 Int. Conf. on Field & Service Robotics*, Helsinki (2001)
- [262] Surgeon General’s Office, Mental Health Advisory Team (MHAT) IV Operation Iraqi Freedom 02-07, Final report (November 2006)
- [263] Sujit, P.B., Sinha, A., Ghose, D.: Team, Game and Negotiation based Intelligent Autonomous UAV Task Allocation for Wide Area Applications. In: *Innovations in Intelligent Machines*. Springer, Heidelberg (2009)
- [264] Sutcliffe, J.H., Sukkarieh, S.: Decentralised tracking and slam applied to UAV sensing & operations. In: *Australian International Aerospace Congress*, Brisbane (2003)
- [265] Sutcliffe, J.H., Sukkarieh, S., Yelland, B.: Simultaneous localisation and map-building, Enabling the revolution in UAV sensing and operations. In: *UAV Asia-Pacific Conference*, Melbourne (February 2003)
- [266] Sutcliffe, J.H., Waldie, J.M.A.: Real-time decentralised data fusion: a framework for future tactical operations. In: *Land Warfare Conference*, Brisbane, pp. 249–257 (2002)
- [267] Tactical EW Division, NRL, Affordably Expendable UAVs, Paper prepared for US Secretary of Navy (April 1999)
- [268] *Technology Development for Army Unmanned Ground Vehicles*, NRC (2002)
- [269] Teo, R., Finn, R.A.: *Project CAREFUL – Scenarios & Measures of Effectiveness*, Joint DSO/DSTO Research Report (February 2005)
- [270] Teubner, G.: *Rights of Non-Humans*, Max Weber Lecture 2007/04, ISSN 1830-7736, European University Institute (January 2007)
- [271] Timofeev, A., Kolushev, F., Bogdanov, A.: Hybrid Algorithms, of multi-agent control for mobile robots. In: *Int. Joint Conference on Neural Networks*, vol. 6 (1999)
- [272] *TTCP-EWS-TP4, Next Generation Tactical ES Geolocation Technologies*. In: *Proceedings of Workshop held at NRL*, Washington (April 1999)

- [273] Vu, T.-D., Burlet, J., Aycard, O.: Grid-based Localization and Online Mapping with Moving Objects Detection and Tracking: New Results, Laboratoire d'Informatique de Grenoble. In: Proc. 2008 IEEE Intelligent Vehicles Symposium (IVS), Eindhoven University of Technology, Eindhoven, The Netherlands, June 4-6 (2008)
- [274] United Nations Convention on the Law of the Sea (UNCLOS), Division for Ocean Affairs and the Law of the Sea, Office of Legal Affairs, United Nations (1982), <http://www.un.org/Depts/los/index.htm>
- [275] Underwood, J., Hill, A., Peynot, T., Scheduling, S.: Error Modelling and Calibration of Exteroceptive Sensors for Mapping Applications. Submitted to J. Field Robotics (Special Issue on Mapping) (2009)
- [276] Unmanned Aerial Vehicles Roadmap 2002 – 2027, OSD (December 2002)
- [277] Unmanned Aerial Systems Roadmap 2005 – 2030, OSD (December 2005)
- [278] Unmanned Systems Roadmap 2007 – 2032, OSD (December 2007)
- [279] Unmanned Systems Integrated Roadmap 2009 – 2034, OSD (March 2009)
- [280] Unmanned Aerial Systems Roadmap – 2007 Edition, Australian Defence Force (2007)
- [281] Unmanned Ground Vehicles Lessons Learned, SPAWAR Systems Centre (2001)
- [282] Unmanned Aircraft Systems Flight Plan, US Air Force (May 2009)
- [283] Vidal, R., et al.: Probabilistic Pursuit-Evasion Games: Theory, Implementation and Experimental Evaluation. IEEE Transactions X(Y) (2002)
- [284] Vidal, R., Sastry, S.: Vision-based detection of autonomous vehicles for pursuit-evasion games. In: Proc. IFAC World Congress on Automatic Control (2002)
- [285] Vollmerhausen, R.H., Jacobs, E.: The Targeting Task Performance Metric: A New Model for Predicting Target Acquisition Performance, Technical Report, AMSEL-NV-TR-230, US Army CERDEC, Fort Belvoir, VA 22060 (2003)
- [286] Walrod, J., Nielsen, C., Hertel, K., Odell, D.: A Network-Centric UUV Acoustic Measurement Range System. In: Proc. IEEE Oceans 2003 Conference (August 2003)
- [287] Walter, W.G.: An Imitation of Life. Scientific American 182, 42–45 (1950)
- [288] Walzer, M.: Just & Unjust Wars: A Moral Argument with Historical Illustrations. Basic Books, NY (1977)
- [289] Wickens, C.: Automation in Air Traffic Control: The Human Performance Issues. In: Scerbo, M., Mouloua, M. (eds.) Automation Technology & Human research trends. Lawrence Erlbaum Assoc., Mahwah (1999)
- [290] Wiegel, V.: Building Blocks for Artificial Moral Agents. In: Proc. Tenth International Conference on Simulation and Synthesis of Living Systems, Bloomington, IN (2006)
- [291] Williams, B., Nayak, P.P.: A Model-Based Approach to reactive Self-Configuring Systems. In: Proc. 11th AAAI Conference (1996), <http://www.aaai.org>
- [292] Williams, Dissanake, S.G., Durrant-Whyte, H.: Towards Multi-Vehicle Simultaneous Location And Mapping. In: Proc. International Conference on robotics & Automation, Washington, DC (2002)
- [293] Yanco, H., Drury, J.: A Taxonomy for Human-Robot Interaction. In: Proceedings of the AAAI Fall Symposium on Human-Robot Interaction, pp. 111–119 (2002)
- [294] Yanco, H., Drury, J.: “Where am I?” Acquiring Situational Awareness using a Remote Robot Platform. In: Proceedings of the IEEE Conference on Systems, Man and Cybernetics (October 2004)

- [295] Yerkes, R.M., Dodson, J.D.: Relationship of Strength to Stimulus to Rapidity of habit Formation. *J. Comparative Neurology & Psychology* 18 (1908)
- [296] Zhen, Balasuriya, J.A., Challa, S.: Recent Developments in Vision-based Target Tracking for Autonomous Vehicle Navigation. In: *Proc. IEEE Intelligent Transportation Systems Conference (ITSC)*, September 2006, pp. 17–20 (2006)
- [297] Zwanenburg, M., et al.: Humans, Agents, and International Humanitarian Law, http://www.lea-online.net/publications/ZwanenburgBoddensWijngaards_LEA05_CR.pdf

Acronyms

ADF	Australian Defence Force
ATD	Automatic Target Detection
ATR	Automatic Target Recognition
BDA	Battle Damage Assessment
COTS	Commercial off the Shelf
COLREGS	Collision Avoidance Regulations
DOD	Department of Defence
DP	Dimensional Parameter
EOD	Explosive Ordnance Disposal
EW	Electronic Warfare
FOV	Field Of View
GAO	Government Accountability Office
GCS	Ground Control Station
HMI	Human-Machine Interface
HRT	Human-Robotic Teaming
ICC	International Criminal Court
IDT	Intelligent Decision-Making Technology
IED	Improvised Explosive Device
ISR	Intelligence, Surveillance & Reconnaissance
JAUS	Joint Architecture for Unmanned Systems
LOAC	Law of Armed Conflict
MOE	Measure of Effectiveness
MOFE	Measure of Force Effectiveness
MOP	Measure of Performance
MTI	Moving Target Indicator
NCW	Network Centric Warfare
OCU	Operational Control Unit
OSD	Office of the Secretary of Defence
OT&E	Operational Test & Evaluation
RC	Remote Control
REA	Rapid Environmental Assessment
ROE	Rules of Engagement
SAR	Synthetic Aperture Radar
SAS	Synthetic Aperture Sonar
STANAG	Standard NATO Agreement
TRL	Technology Readiness Level
UAV	Unmanned/Uninhabited Air Vehicle
UCAV	Unmanned/Uninhabited Combat Air Vehicle
UCUV	Unmanned/Uninhabited Combat Underwater Vehicle
UGV	Unmanned/Uninhabited Ground Vehicle
UMV	Unmanned/Uninhabited Maritime Vehicle
UNLCLOS	United Nations Convention on the Law of the Sea
USV	Unmanned/Uninhabited Surface Vehicle
UUV	Unmanned/Uninhabited Underwater Vehicle
WSB	Weapons Safety Board