Technical Evaluation Report

Glyn WYMAN
2 Palmer Gardens
Wivenhoe, CO7 9FL
UNITED KINGDOM

Glyn.wyman@gmail.com

ABSTRACT

This paper presents the findings of the Technical Evaluator of the Specialist Meeting held 16th and 17th October 2018 in Lisbon Portugal. A summary of each presentation is included as an annex. It concludes that autonomous agents are evolving and further research is required to obtain a robust and secure system. The desire to analyse the systems formally conflicts with the complexity of the software employed leading to lengthy time scales. A wide set of risk analysis tools are available but the capability to drill down and achieve a quantitative value requires extensive research.

1.0 INTRODUCTION

Autonomous agents have evolved over a number of years and the proliferation for commercial use is evident. Unmanned vehicles are readily available in a number of domains and are most useful in adverse and contested environments covering all domains including space. Activities where vehicles will be exposed to situations where safety of life has been identified is an obvious application. Tedious surveillance is a further example where the concentration of the observer may lapse or be compromised by external distractions. Energy consumption and consequential heat dissipation can be critical factors, in selecting the appropriate vehicle and associated sensors. Robustness and security become significant when commercial systems are adopted for military applications. The advances in driver-less road vehicles are significant and attracting considerable finance.

The IST-164 Study Group has been active in this area and sponsored the event. The meeting was subdivided into four sessions addressing the following topics:

Intelligence in Autonomy

Securing Autonomous Platforms

Risk Assessment for Platforms and Missions

Mission Concepts and Modelling

The chair of the Working Group also co-chaired the specialist meeting and had support from the members.

2.0 THEME

Topics for discussion raised in 'The Call for Papers' included:

Mission concepts integrating unmanned and autonomous systems

- Scenarios for current and future applications of unmanned and autonomous systems
- Mission planning and execution with increased levels of autonomy

- Modelling, simulation, and emulation of unmanned and autonomous systems
- Advances in man-machine interfaces and enabling technologies
- Training in the human-machine work force

Unmanned and autonomous system platforms

- Taxonomies and characterizations of unmanned and autonomous system platforms
- System, data architectures, and embedded system process control
- New concepts and development in unmanned and autonomous system vehicles
- Interoperability among unmanned and autonomous system vehicles

Risk management and risk assessment for autonomous systems

- Vulnerability and risk assessment tools for unmanned and autonomous systems
- Security and advanced threat assessment of unmanned and autonomous systems
- New challenges in risk management of autonomous cyber-physical systems

Validation and verification for autonomous vehicles and software

- Validation and verification of artificial and computational intelligence for autonomy
- Ensuring trustworthiness of sensor data
- Security assurance in embedded systems and systems with commercial components
- Secure and efficient implementation of crypto mechanisms
- Extended reach with cyber operations and compromised systems

Enabling technology for autonomy

- Blockchain architectures for distributed systems and logistics chain management
- Autonomous agent novel learning methods and algorithms
- Secure and robust machine learning
- Tamper protection
- Cryptography for distributed autonomous systems in heterogeneous environments
- State of health assessment of autonomous systems

3.0 PERCEIVED MILITARY ISSUES

The Military are required to operate efficiently in hostile environments which could be inaccessible to humans; unmanned or autonomous vehicles are a potential solution. Recent developments in the field of robotics and unmanned vehicles in commercial areas could be exploited for military applications, which coupled with extant vehicles and procedures can enhance capabilities. A strong desire to quantify the risks associated with particular options enables commanders to refine their selections, but in a large number of instances qualitative assessments will be the best available. The analysis of systems using formal methods is highly desirable but the complexity renders the formal analysis impractical. The transition to the so called ‘hyperwar’ is inevitable where the human reaction time is much longer than the reaction time demanded for action to be instigated. A further aspect is that robotics will offer/give a programmed reaction irrespective of any changes to the environment or psychological state of the human user. The ethical aspects need to be resolved but the benefits of employing autonomous agents is considerable particularly in a multi-national environment where agents can co-operate freely.

Verification and validation of function is a major concern where assurance cannot reach 100% and safety of life is involved. Refining missions to make best use of assets whilst adhering to specific constraints is a hard problem compounded when security parameters are introduced. Assistance using deep learning can be implemented and the wider use of AI is an issue under research elsewhere but closely aligned.

4.0 EVALUATION

4.1 FACILITIES

The room was well appointed with the delegates sat in raked rows with full visibility of the screen and adequate acoustics. WiFi access was provided giving access to the Science and Technology Office (STO) website and other internet facilities allowing the delegates to view the papers which had been uploaded by the secretariat prior to the symposium. I understand that some delegates had difficulty accessing this facility which may be a function of the application they were running but not resolved at the site, this forced those delegates to rely on the presentations. Coffee was served during the breaks with ample opportunity to engage with the presenters to elaborate on the research and to seek points of clarification, an essential aspect of a Specialists' Meeting. I observed very productive discussions in the breaks.

4.2 OVERVIEW

The use of unmanned vehicles in hostile environments coupled with the proliferation of deep learning algorithms associated with the situational awareness will increase. Much of the software was designed for commercial use and will exhibit vulnerabilities and associated risks without the rigour imposed on military projects. The themes identified in the call for papers did not attract the research papers from academia anticipated. The structure of the Specialists' Meeting mirrored the papers which had been submitted and was subdivided into four sessions. In addition three keynote speakers were invited to provide a background for the discussions. In reading the abstracts I was concerned about the quality but had little time to rectify the position because of the late submissions, I understand that the technical committee had already approached the authors to enhance their presentations. The concern remains about the quality of the publication, all authors should submit a full paper. The presentations lacked technical depth which was explored during the discussions and in the margins, with the consequential discussions gaining the necessary detail and insight into the submissions. Unfortunately the scheduled 'Panel Discussion' to conclude the meeting designed to highlight the contentious issues had a very limited audience when a significant number of the delegates left to follow a tour. It is strongly recommended that activities do not conflict whilst arranging the programme with the hosts. Without the additional networking achieved through the breaks the value of the Specialists' Meeting would have been limited.

4.2.1 Intelligence in Automation

A general overview of robots and autonomous systems was presented with videos demonstrating particular attributes. Guidance is generally regarded as relatively advanced, however, provision must be made to counter the situations in which autonomous systems are compromised through denial of resources, e.g. when communications or satellite navigation systems may not be available. Situational awareness has advanced with the evolution of sensors but limitations may be imposed through payload weight restrictions. Additional research in this area is advised, coupled with improvements in object identification in sophisticated models. A video demonstration of a robot grasping an object was shown but detail and the techniques applied were not exposed. Correction of gait provided evidence that feedback control has improved and the robot was able to recover from a fall. Robots are evolving rapidly and the military should benefit from the progress.

Intelligence in the form of object identification was shown to be subject to false identification if spurious pixels are introduced; not distinguishable by a human. A mathematical approach can be applied to identify a manifold which defines the state space which is most susceptible to this activity. A concern identified is that training data for neural nets will not cover the state space of military interest. Both the training data and the model need to be protected to avoid miss classification caused by actions of third parties. Formal analysis is theoretically possible but the complexity demands extensive processing over an extended period, the general consensus was that it would not be practical. The use of deep learning methods is becoming widespread but the topology is heuristic with currently no understanding as to how to implement improvements. Results are, suprisingly, far better than anticipated. The desire to achieve a more formal standing requires further research.

4.2.2 Securing Autonomous Platforms

A tutorial on Blockchain technology was offered which provoked discussion on the security of the architecture. A loose connection to autonomous agents was offered but aspects of the high demand on the infrastructure and the processing requirement to support blockchains may rule out the technology. The principle of the bitcoin architecture is that full visibility is maintained including tracing any transaction to the source which could be an advantage in autonomous systems. Other variants can protect the source which could also have applications in the military domain. The development of SingularityNET, which is open source, allows users to exploit extant code but also gives an advantage to nefarious users.

Autonomous vehicles can be designed in a modular fashion to increase the flexibility to which the vehicles can be employed. A presenter advocates a distributed integrated modular avionics system (DIMA) to achieve compliance with standards to ease certification. The design is based on a federated architecture to ensure all payloads remain robust. The protocol incorporates a crypto signature. Caution is raised about the scalability and that each configuration is non trivial to incorporate. The vehicles will be subject to the complete spectrum of attacks from kinetic through EMP to the disruption of neural nets, some protection mechanisms are well established but the same degree needs to be afforded to the machine learning elements. It is recognised that both structure and training data needs protection. A third party with knowledge of either can influence the classifier significantly. It has been shown that analysing the statistics of the identifier can prompt further analysis to confirm a false result. There are multiple avenues through which a third party could compromise the data and, if feasible, diversity of data should be introduced.

4.2.3 Risk Assessment for Platforms and Missions

A general appraisal and the use of UAVs was offered with the associated benefits including the greater flexibility available. The vehicles are generally regarded as a force multiplier particularly in situations where humans are prone to attention deficit or where a mission demands swarming behaviour. In adversarial environments communications can be strained creating an incomplete graph which may violate the assumption on which control is based. When all nodes are not fully informed the infrastructure becomes inefficient with increased overheads to assess where information is sparse. Payloads and control systems require power which must be dissipated, with small vehicles this may be a limiting factor on the uses. A trend towards biologically inspired systems is encouraged but caution is required to ensure overall security particularly when third parties can introduce alien nodes.

A series of risk assessment strategies and concepts were expressed with no one preferred method identified. Two schools became evident with divergent processes, those with a security perspective and those with a safety critical view. Risk assessment should encompass both disciplines to contribute to the commanders perspective. Differences were made public with the recommendation to encourage engineers to train in both disciplines. A wide range of generic risk assessment tools are currently available but specialist knowledge is required to generate consistent estimates. The wide distribution of estimates is of concern particularly when decisions as to whether a vehicle can fly is to be made. Much of the science is based on empirical data which may not map to the environment of interest. A specific example was given in which quantitative figures are generated from an empirically derived equation and further that it produced a wide variation in the results with minor interpretation of the answers given in the associated questionnaire. One paper looked at the vocabulary and offered the basis to have a common understanding, of particular interest is an attempt to have the various tiers for autonomy defined. A wide ranging concept was proposed in which the users could drill down through the options and generate a quantitative measure from elemental aspects on a logarithmic scale.

4.2.4 Mission Concepts and Modelling

An exercise was described in which unmanned vehicles and unattended sensors were assets used to defend a military establishment, details of the exercise will be available after the event. The planning and the potential were provided during the presentation. A desire to achieve interoperability was expressed starting with an agreed ontology. The history of UAVs was outlined which perhaps explains the lack of commonality as the vehicles evolved and had diverse environments in which to operate. An attempt to use assets in an efficient manner whilst operating within the mission constraints was described. The aim is to map the utility function, which could include security, onto the mission space which is dynamic. The model is based on a Markov chain and showed potential.

5.0 FEEDBACK FROM TER PRESENTATION TO IST PANEL

An initial appraisal of the meeting was given to the IST Panel in which the Specialists' Meeting was declared as satisfactory, some of the Panel Members were not party to the presentations and discussions giving a view based on the abstracts and papers available prior to the event which they thought was poor. I agree with their position but the effort by the technical committee to improve the content was successful and the quality is based primarily on the presentations and discussions. If the full papers are not available and the publication resorts to the extended abstracts that publication will be of limited quality. A move to improve the quality of the papers and enhance the reputation of the STO would be to introduce a more rigorous paper selection aligned with the IEEE process. This demands a longer lead time and will impact on the time spent on the preparation by the technical committee effectively introducing a peer review. We appear to be running under the cusp of high reputation and generating papers which are rarely cited. Greater technical depth could be achieved by giving stimulus to appropriate researchers, significantly improving the standing of STO events.

6.0 CONCLUSIONS

The Specialists' Meeting provided a forum to extend the knowledge of the delegates and exposed areas where more research is necessary. The vulnerabilities, risks and benefits of deploying and integrating autonomous systems equipped with complex learning systems was explored and discussed. In particular the implementation of convolutional neural nets (CNN) was identified, where a better understand of the operation of the classifiers is required which is currently selected by heuristics. Numerous risk assessment tools are available but none appear to be optimal. Attempts to quantify the risk associated with deploying autonomous systems and integrating them into a larger operational environment require extensive testing, evaluation, and strategic planning. Extant models are prone to diverse results with small changes cause through interpretation of the questionnaires. The call for papers did not solicit sufficient research papers and allow for a peer review which would have been advantageous. The outcome was satisfactory derived from the extensive discussions in the margins provoked by the presentations.

APPENDIX

Synopsis of the Papers and Associated Presentations

The host, Jose Borges, welcomed the delegates to Lisbon and wished for a successful meeting The chair of the IST Panel Michael Wunder opened the meeting with a brief overview of the STO and the position of IST within this body. The disciplines addressed in 'The Panel' and some of the work in hand was presented. Full detail of the structure and the associated activities is available from the STO web site. Misty Blowers and Federico Mancini as co- chairs of the specialist meeting started the proceedings. The following is a summary of the papers, inclined readers are directed to the STO website to clarify the detail where the papers are available to download.

Keynote 1: *Current Challenges for Autonomous Robot Systems . Pedro Lima* An overview of the situation and levels of development in the field of autonomous robotic systems was presented with videos demonstrating the progress. The presentation assessed the indoor and outdoor developments separately. A linear progress was considered following the activities sensing - control – guidance – planning - navigation, control-action. Guidance is at an advanced stage with the ability to infer preferred routes, and adapt to human intervention and resolve conflict. Speed of reaction could be improved. Object identification is stabilising but higher reliability is required. In outdoor applications interpretation of street furniture was raised as an issue with improvements in the classifier demanded. Other proposed enhancements included improve situational awareness when indifferent lighting conditions are prevalent. The ability to grasp objects from a mobile base is an order more difficult than that from a stationary platform introducing a further degree of freedom over the present capability.

Paper 1: *Autonomous and Dependable Multi-Agent Systems for the Mission Planning of Multi UAV Surveillance Missions: Domenico Pascarella* A tool which strives to find an optimal route to optimise the wider mission. It employs a dynamic and decentralised approach to meet the dependability criteria through connectivity. The optimisation uses a game theoretic approach expressed as a Markov chain with the aim to achieve a Nash equilibrium. Verification is achieved through formal analysis. The approach is laudable but it is questionable that the vehicle agent will have sufficient processing power in a highly dynamic environment.

Paper 2: *Assuring Autonomy: Ramesh Bharadwaj* The advances in deep learning should be harnessed by the military with the necessary enhancement for reliability. Trust is a necessary requirement in this context. Many examples are in the public domain which show false classification. Adversarial perturbation of the images needs to be considered, areas most susceptible to change can be identified by applying a manifold. In military environments the training data may not cover the complete state space. Selecting the topology of the models remain heuristic but the results are much better that anticipated. Further analysis is required to understand the effects of the different topologies. The analysis can be theoretically undertaken in a formal manner but would require enormous computing power

Keynote 2: *Blockchain as a New Framework for Unmanned System Swarms: Misty Blowers* A comprehensive treatise on the different methods to implement blockchains identifying attributes which are desirable in a swarming structure of autonomous vehicles. In the bitcoin architecture all transactions can be traced to the source and the economic incentive is to add elements rather than attack an existing block. Considerable work needs to be undertaken to map the principles of blockchains onto a swarming structure in an efficient fashion. Benefit can be obtained by employing the use of Distributed Ledger Technologies which incorporate smart contracts as distributed architectures. It is thought that the overheads required for implementation may rule the techniques impractical for use in unmanned vehicles.

Paper 3: *Distributed Integrated Modular Avionics: Miguel Barros* A distributed integrated modular avionics (DIMA) project was presented to satisfy the certification requirements. The model is based on a federated architecture and made robust with crypto signatures incorporated in the protocols. It has been proven and can be scalable but each new configuration is non-trivial.

Paper 4: *Towards a Trustworthy Foundation for Assured UAVs: Thomas Macklin* Unmanned vehicles will be subject to the whole spectrum of countermeasures but particularly from cyber. Software integrated into such vehicles are generally not supported by traditional scanning tools. The authors offer a raft of potential solutions mainly low level with a bias toward ad hoc systems. Protection of both training data and the topology is recommended, security control by injection was also raised.

Keynote 3: *Gaps in the Basic Research Needed for Distributed Autonomous Vehicles: Frederick Leve* The paper exposes the potential of failure when operating unmanned vehicles and alerts the user to the aspects which need further consideration. The impact of poor communications is a case in point which violates the all informed concept. A further aspect is when external navigation systems fail requiring distributed control amongst the nodes. Consideration of the inherent resilience when operating as a swarm was offered but contrasted with the safety in such formations. The author stated that decomposition to undertake formal analysis was unlikely. In the space domain, in particular, aspects of radiation hardening and visibility of other spacecraft was non trivial. Four areas were identified as requiring further research a) Control and dynamic system theory b) Physical intelligence c) Specification and verification in adversarial environments and d) Communications

Paper 5: *Unmanned Aircraft Systems Risk Assessment Review of Existing Tools and New Results: Diogo Duarte* The authors have completed a review of existing tools for unmanned air vehicles to establish figures for risk assessment. They have concluded that the Specific Operational Risk Assessment (SORA) is satisfactory, but has limitations arising from the empirical derivation of the equation employed. In the process a questionnaire is completed and wide discrepancies can arise from the interpretation of the operators.

Paper 6: *Managing Adversity Risks for Non-anthropogenic Systems: Ian Bryant* The group at Warwick University in the UK are expanding the theoretical work into practical implementations. An all encompassing construct which synthesises the elemental components, from the unlikely but catastrophic through to the common occurrence but low impact. The attempt is to enumerate the derived risks to provide the commanders with a readily understood figure. As part of the process the group have needed to establish some definitions. The inclined reader is directed to their definition of the various stages of autonomy.

Paper 7: *Risk Management Framework: Qualitative Risk Assessment: John Piper* The presentation was given by an experienced consultant with an holistic outlook. The associated paper describes the process with an emphasis on fact finding rather than fault finding. A strong recommendation is to introduce security features at an early stage. The outcome is a qualitative assessment developed through interviews and presented as a coloured matrix.

Paper 8: *A Reference Model for Unmanned Systems: Mario Marques* No universally accepted reference model is currently available to quantify mission assurance. The presentation focused on a mechanism to subdivide the problem to then quantify the assessment by vehicle weight. Several examples were cited where UAVs had been employed during successful missions..

Paper 9: *Mission Orientated Optimisation: Gustav Anderson* A superficial description of a tool was provided to map utilities onto the mission space in an optimal fashion. During the discussions it became evident that the tools could be applied more rigorously. A model is run in the pre-mission phase based on a Markov chain with realistic constraints imposed on the resources. During the mission phase the resources and constraints can be dynamically adjusted to reflect reality. Improvements to the model are then made post mission. The aim was to provide a metric that offers high level information about the information exchange. Unfortunately no detail of the model was provided.

Paper 10: *Base Protection with Autonomous Systems: Solveig Bruvoll* The presentation was a description of an exercise to incorporate UAVs and unattended sensors whilst defending a controlled environment. The specific case is Trident 2018, I anticipate a report of the findings in the near future.