Human Autonomy Teaming in Naval C2 – Insights from Dstl's Intelligent Ship Project

A J Tate, MSc, CMarEng FIMarEST[+]* & C J Cooke[+]
[+] Defence Science & Technology Laboratory (Dstl), UK
* Corresponding Author Email: atate@dstl.gov.uk

Synopsis

Future military forces and platforms will operate within increasingly complex operational environments, facing a more aggressive and diverse range of threats. This will continue to increase the volumes and speed of data, and hence information, that platforms and their human commanders will need to capture, process and respond to. This leads to a clear need to be able to effectively, and flexibly utilise the best of both human and Artificial Intelligence (AI) skill sets in future Command and Control (C2) and decision making more broadly.

The UK Defence Science and Technology Laboratory (Dstl - part of the UK Ministry of Defence (MOD)) has funded three phases of the Intelligent Ship project. The project explored Human-Machine, Machine-Machine Teaming (HM3T), considering relationships between team members, and the approaches needed to enable them, aiming to understand the challenges and opportunities this approach brings to future Naval C2 and defence more widely.

The project achieved this through the development of a range of intelligent machine agents and through the development, evaluation and demonstration of those agents, with human operators, within a systems level architectural 'sandpit' known as the Intelligent Ship AI Network (ISAIN). The project also started to consider underpinning design issues, such as: how to effectively design-in humans into a system; how to manage inevitable conflict between recommended Courses of Action (COA) and; how such a system should be configured and managed with varying Levels of Automation (LOA). This work was delivered by a multidisciplinary team of external suppliers and evaluated within Dstl's simulation facilities.

While there is current focus on removing humans from threat environments, automation and autonomy will continue to have limitations, which when coupled with UK policy, drives a continuing need, and a clear benefit from the use of human based analysis and decision making skills. The aim of this project was, therefore, to optimise military decision making rather than pursue solutions to full autonomy or crew reduction for their own sake.

This paper will build on a previous 2022 International Ship Control Systems Symposium (iSCSS) paper to provide an end of project perspective on the key insights and challenges identified. It will postulate what future research and development is required, and how the project lessons should be integrated into future C2 concepts and designs.

Keywords: AI; intelligent systems; automation; future command control; Human-Machine Teaming

1 Introduction

Future military forces and platforms will operate within, and against, an increasingly complex, diverse and technology focused set of threats. This will increase the volumes and rate of delivery of data, and hence information, that human commanders and their supporting systems need to capture, process and respond to. Decision making speed will also need to increase as threats themselves become faster, more numerous and less predictable, or identifiable.

These challenges are amplified when considering the application of future, and potentially more distributed, combat systems. Systems such as distributed high-powered sensors and Directed Energy Weapons (DEW) drive a need for better coordination across a force, or fleet, and also improved intra-platform connectivity between internal systems, for example, power system controllers, and navigation systems and the combat system itself.

Addressing these challenges inevitably leads Defence to consider the wider use of automation and autonomy, with an expectation of the wider implementation of Artificial Intelligence (AI) approaches within future Command and Control (C2) systems. While there is significant focus on autonomy in defence systems and platforms, this represents the focused use of AI within Decision making processes. The emergent paradigm of developing systems to handle decision making demands has been described as Decision Intelligence, defined by Gartner (Brethenoux, 2021) as a "practical discipline used to improve decision making by explicitly understanding and engineering how decisions are made, and outcomes evaluated, managed, and improved by feedback".

The benefit of combining machine and human based intelligence skills and strengths, coupled with the demands from policy, leads to a need to form effective and adaptable teaming relationships between intelligent machine agents, matching automation systems, and human operators – a Human-Autonomy-Team (HAT). The project has also used the term human-machine and machine-machine teaming (HM3T) to reflect the need to also enable multiple AI agents to work together within a team as collaborative AI.

Authors' Biographies

Andrew Tate is chartered senior principal engineer for Dstl (UK Ministry Of Defence (MOD)). He manages a range of naval focused Science & Technology research activities for UK MOD, focusing on power systems, sustainability and autonomy. He has an MSc in Marine engineering from Newcastle University and is a fellow of IMarEST.

Colin Cooke is senior principal scientist for Dstl (UK MOD). With 39 years of experience working on Science & Technology research activities for UK MOD, focusing on autonomous systems, and their underpinning technologies (e.g. autonomy, launch and recovery, system architectures, platform integration, C2, interoperability standards).

Defence needs to understand how to design, manage and integrate these systems in order to reduce the potential risks of complex HATs and to maximise the defence capability benefit they could bring. This requires a paradigm shift from designing systems and their interfaces with humans, to designing systems with human and machine teaming interaction as a fundamental element of the initial concept and design.

Dstl has previous briefed ISCSS of the outputs of the second phase of the Intelligent Ship project (Tate, 2022). This paper will take a wider retrospective view at the end of phase 3 activities to capture some of the key insights identified and to suggest key challenge areas and needs going forward, that could allow HAT and collaborative AI to be practically integrated into future C2 concepts. The research, while focused on Naval C2, has clear and transferable applicability and relevance to any future defence C2 system, network or platform where humans and AI agents need to work together in an effective HAT.

1.1. *Project aims & phases*

The Intelligent Ship project was initiated with an aim of starting to understand, develop and evaluate concepts that address the challenge of enabling, integrating and managing complex HM3Ts in generation after next platforms. It was funded under Chief Scientific Advisor's research funding under Dstl's Autonomy Programme.

The project was challenged to consider a clean sheet approach to avoid the constraints of current Naval C2 processes and system architectures. It attempted to understand and qualify the follow project vision:

Machine learning and AI will be more closely integrated and teamed with humans, leading to timely, more informed and trusted decision making and planning, within complex operating environments.

Each phase has had a specific focus which can be summarised as follows:

- Phase 1: **'develop enablers'** – This phase delivered a diverse range of AI agents across a range of naval functions via a Defence and Security Accelerator (DASA) industry call. It funded the underpinning development and experimentation 'sandpit', the Intelligent Ship AI Network (ISAIN)[1]. This phase also started the process of developing an underpinning tactical naval navigation agent (TACNAV), and delivered a study called Platform Design Risks & Opportunities (PeDRO)[2], which captured the context, impacts, risks & opportunities of wider use of machine intelligence & automation in naval platforms.
- Phase 2: **'Integrate & evaluate'** - this phase developed Phase 1 agents and a range of new agent concepts, before integrating them, with a number of agents from other programmes, into ISAIN. The whole system was then tested against a representative naval scenario using Dstl's Military Advisors as operators to qualitatively test issues such as Trust in Automation (TiA), usability and Situation Awareness (SA).

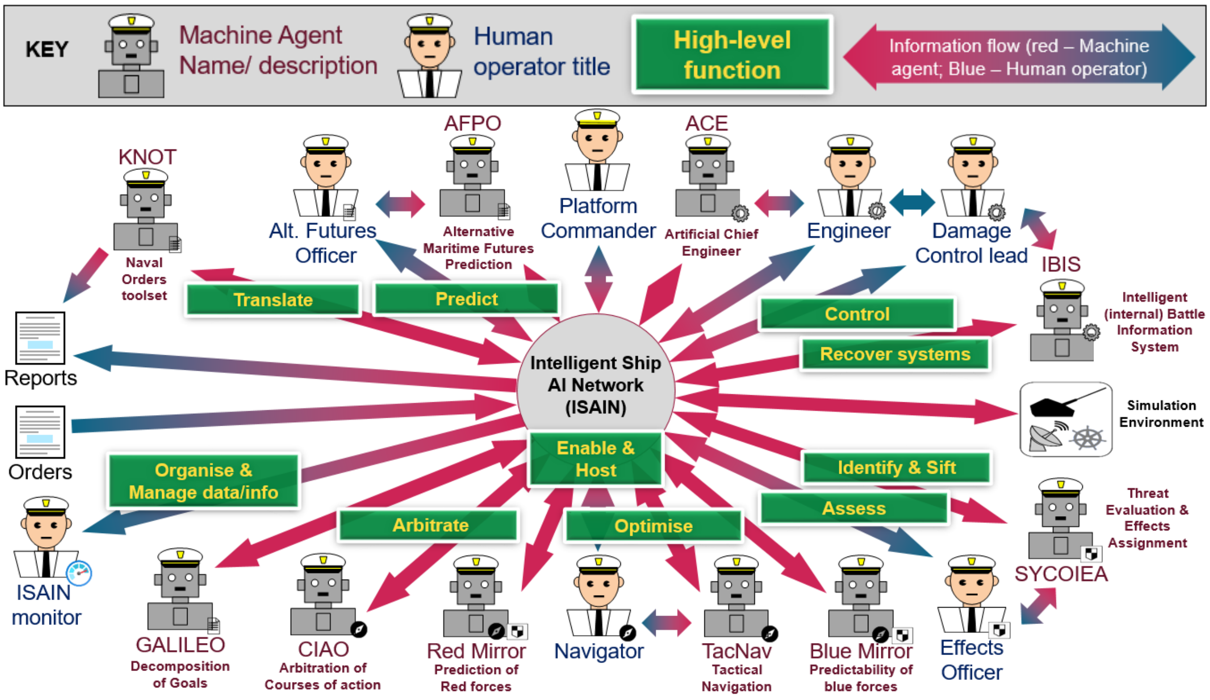

Figure 1: Phase 2 HAT system concept (humans and agents) & key functions

[1] ISAIN & TACNAV developed by CGI and owned by UK MOD/Dstl

[2] PeDRO informed further funding by BMT on their Highly Autonomous Warships Technologies (HAWT) project (BMT, 2021)

- Phase 3 '**Focus on the human in HAT'** – this phase aimed to use and adapt phase 2 agents and systems, but to design a system (both human and machine elements) from the ground up to meet future naval Anti-Air Warfare (AAW) capability needs. This resulted in greater focus on human-factors, roles and skill needs within the system; the requirements for human-machine interfaces and information management, and; the use of AI based arbitration approaches to manage potential conflict between proposed Courses of Action (CoA) from different agents.

It is important to note that the project purposely did not focus on, or aim to demonstrate, the potential to reduce crewing, rather to understand how a system could make best use of the strengths and skills of both machine and human intelligence working within a team. Inevitably, with greater technical maturity and trust in an Intelligent Ship type concept, crewing reductions should be possible.

2. End of project insights

The project previously presented a paper at INEC/ iSCSS 22 (Tate, 2022) which gave an insight into how the project was executed, details of some of the intelligent agents developed, and the evaluation approaches and key conclusions from the first two phases. The following discussion therefore, provides a retrospective discussion on some of the key insights at the project end. It is hoped that these provides general, and useful insights into the development, design, management and integration of complex HATs comprising of multiple agents and human operators (a HM3T) irrespective of the application or domain.

2.1. *Designing-in humans early*

The design of Naval C2 systems may start with assumptions around its human operators, their skills and needs, however, it is often the case that issues such as information management and HMI design, and corresponding training needs are considered later, perhaps too late, in any design process. The adoption of HAT, however, drives a demand to consider human operators early, and in parallel with the machine elements of a system, and has the potential to benefit from the best capabilities and skills of both the machines and humans.

The early adoption of intelligent systems has initially and is likely to continue to be focused on supporting existing operator roles and their core tasks. They have been typically designed to augment or directly replace a single operator and hence, have limited impacts on the overall roles, ranks and skills of the human team. An example of this would be a toolset/ agent that directly replaces an air picture compiler role.

The successful introduction of more sophisticated and broader capability intelligent agents offers both an opportunity and a risk that it drives changes within human organisations and hierarchical structures, changing the skills demand, the experience needed and ultimately roles. There is a clear potential impact on a range of Defence Lines of Development (DLODs) in areas such as training and people, as well as on more specific issues such as recruitment and retention.

An example of this is shown in Figure 2, showing a postulated conceptual AAW team structure and role descriptions developed under phase 3 of the project. These are based on a functional assessment of the AAW process, and the HAT system's design and policy needs, rather than existing Royal Navy (RN) team structures, roles or ranks

2.2. *Decision making with varying Levels of Automation*

Intelligent decision systems offer the opportunity to operate more dynamic Levels of Automation (LOA), where tasks could be allocated to humans and machines based on factors such as workload, threat level and the relative TiA in a given scenario or situation. An operator may need to take manual control of a system (potentially in, or on-the-loop) in peacetime operation to maintain SA and experience, and to build trust in (or understand the limitations of) the automation and autonomy. This creates a need to both understand how and when LOA should be changed, and how decision making is achieved at each LOA and is likely to change dependent on the situation.

A decision making framework was developed during phase 3 of the project as illustrated in Figure 3. It shows decision making across three LOA categories; fully automated, augmenting or decision support. It also shows how each category has a blend of human and machine interactions and how the best decision making approach depends on the characteristics of that decisions, e.g. Time available, level of risk or complexity. The project assessed an example complex AAW scenario to understand decision-making hot-spots. These are points of time when the complexity of a decision or the required speed of response places pressure on any C2 system, including its human operators. By using this framework an assessment of the balance between human and machine based decision making could be made. Examples of hot spots are shown in Figure 3, marked by the relative position of the green stars on each bar.

There are clear parallels to the current debates and experience with vehicle driver assistance systems, where operators need certainty over their responsibilities at any given point and need to maintain sufficient SA to take over control or intervene where and when necessary. For these reasons a clear policy is needed around how LOA are managed and changed in a military context, based on balancing the desire to optimise decision-making, while ensuring sufficient operator understanding and visibility of both the LOA and their decision making role within it.

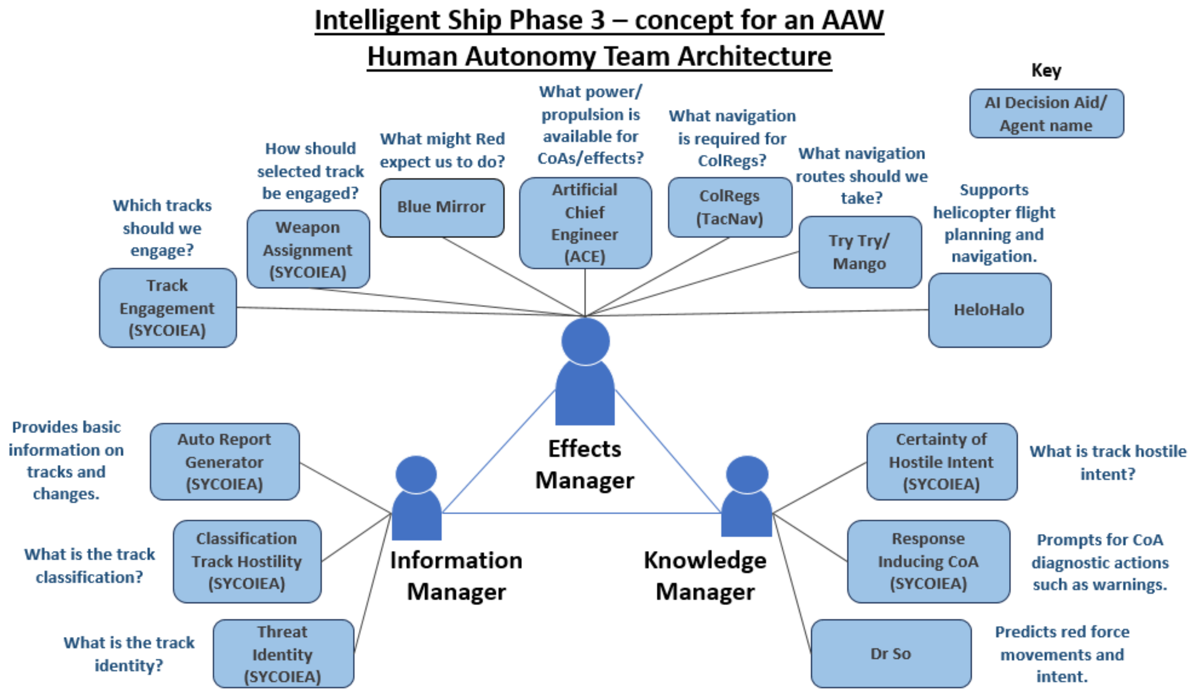

Figure 2: Proposed phase 3 AAW team structure & supporting Agents[3]

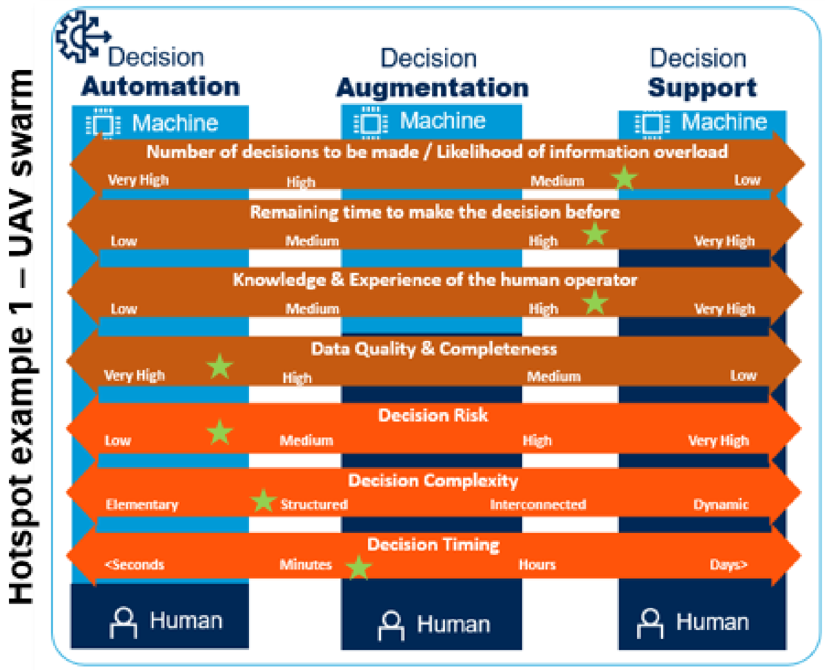

Figure 3: Phase 3 decision making framework with the characteristics of an example decision making hot-spot (Developed by Decision Lab/ Diem)

2.3. *HAT Training aspects*

There is likely to be clear benefit in developing a combined approach to training and development of both operators and AI agents in parallel. For the human this allows them to build understanding, TiA and experience of AI agents' (or collaborative team of agents') limitations in a given scenario or context. For each agent there is a training demand for AI approaches (e.g. Reinforcement Learning (RL)) that would benefit from training

[3] More details on the individual agents can be found in the previous INEC/iSCSS paper (Tate, 2022)

scenarios involving humans, either to learn from their responses, or from the various training scenarios the operators train within. This would also help to inform development of HMIs.

While it is potentially dangerous to anthropomorphize machine agents within a HAT, human operators need to understand the limitations of, and develop trust in the same way as they would for new human team members. They will also need to train and develop as a team to become effective. Experience of non-intelligent systems show that user frustration can quickly result in limited trust, and hence reduced reliance upon it and in some cases not use it all.

Inevitably new agents based on machine learning or other AI training techniques will be ‘inexperienced’ in certain scenarios. Creating or capturing a desirable level of real world training data is a challenge, while naval platforms operate across a broad range of environments and tasks, resulting in large data-sets, they infrequently repeat the same tasks and experience certain environments and threats through-life, especially, and fortunately, high threat scenarios. This can limit the applicability of many big data based AI techniques and would make it valuable to able to integrate experience from other ships in a class, across the wider fleet, international partners or even from retired platforms, into currently deployed systems. This challenge places a demand on modelling and simulation capabilities as discussed in section 2.6.

2.4. *Enhancing Machine-Machine interactions*

While the Threat Evaluation and Weapons/Effects Allocation (TEW/EA) process offered the project significant and complex opportunities for multiple agents and operators to interact, the project also wanted to explore wider interactions with a platform’s own systems. This drove the parallel development of TACNAV, platform control focused tools such as Rolls-Royce’s Automated Chief Engineer (ACE – providing power and marine systems automated control) and Fraser Nash’s Internal Battle Intelligent System (IBIS – optimising resilience and damage recoverability strategies).

By bringing these agents and their developers into a single integrated system and project, new links and opportunities were quickly identified, including new uses for tools, or for the data and information they produce. An example of this included using a real time analysis tool such as IBIS to generate predictions of future material state and residual capability after damage, and providing information or constraints to other agents proposed CoA. Another example was the ability to balance navigational options between the needs of safe passage (e.g. collision avoidance from TACNAV) and other directional demands, such as to provide line of sight for an effector (TEW/EA), all constrained by availability of power and propulsion (via ACE).

The concept of Tactical Energy Management (TEM) also becomes relevant, with better linkages between material state awareness, power management and navigation system potentially providing command or other agents with a better understanding of options and limitations with, for example, energy based systems such as Laser Directed Energy Weapons (LDEW).

2.5. *HAT System design & enablers*

The aims, development and key characteristics of ISAIN were discussed in a previous paper (Tate, 2022). ISAIN was configured to support HAT experimentation and was not optimised for deployment, but provided the project with the ability to configure, test, assess and monitor interactions between multiple AI agents and Human operators. In setting to work and then evaluating the HAT system developed in Phase 2, a number of lessons were learnt around the design, management and confirmation of a complex HM3T. These were investigated further in Phase 3 and reflected in the designs developed. This sections summarises some of the key enablers to facilitating a functional and effective HAT.

2.5.1. *HAT management*

Phase 2 of the project included a series of HAT evaluations events, allowing the system, and operator’s experience of the system to be captured and assessed. This identified challenges around the management of the system and the information within it and hence a potential need for a HAT manager. It was envisaged that this could operate as a ‘teammate’ within the HAT, sensing and understanding the system environment, the processes within it and the capabilities and needs of the wider team members. Functions this HAT manager could provide include:

- **Functional Distribution** of functions and/or tasks between humans and AI agents within a HAT
- **Oversight and Control** of the ways and extent to which humans within a HAT are able to maintain SA and control over AI agent activity and understanding.
- **Coordination** of the ways in which humans and AI agents within a HAT are able to coordinate their actions and SA in response to constraints such as timing, resource allocation, prioritisation and conflict resolution.
- **Cognitive Enablement** i.e. the ways in which humans and AI agents within a HAT mutually co-facilitate complex cognitive activities such as sense- and decision-making.
- **Adaptation** of a HAT to new behaviours based on new circumstances and demands.

 https://doi.org/10.24868/11145

- **Ethics** i.e. the ways in which humans and AI agents within a HAT apply/deploy ethical/ policy driven constraints.
- **Trust Calibration** in which a HAT supports interactions that enable evidence-based calibration of appropriate and context specific trust

2.5.2. *Arbitration*

Decision arbitration was a core focus area for the project. A HAT is built around a set of AI agents that make observations, predictions and recommend a CoA. The term arbitration is, therefore, used to describe the process of resolving conflicting outputs or proposed CoAs from two or more agents in an explainable way. The concept was initially explored through the development and evaluation of an agent called Compounded Intelligent Agents for Optimisation (CIAO), which was based on Reinforcement Learning (RL) techniques and developed by Decision Lab.

Enacting the phase 2 HAT system's design suggested value in the use of a capability like CIAO in multiple locations within the system, potentially operating at different levels of decision making. An example could include a Navigation arbitrator agent, arbitrating between the outputs of two alternative navigation agents, before a higher level arbitrator ranked CoAs based on constraints, such as the availability of power.

This idea was developed further in phase 3 by considering which enabling AI or non-AI techniques would be best matched to each decision area and/or decision level. Techniques were assessed against decision making factors such as the level of explainability required, the decision speed needed and the overall complexity of decision making. Options considered included:

- Rules-Based algorithms: a set of explicit instructions or conditions that dictate system behaviour in all situations. These are typically pre-defined by domain experts and engineers. It is a deterministic approach, and there is no behaviour that is "learnt" in contrast to machine learning methods;
- Optimisation algorithms: a broad class of algorithms that are used to find the best possible solution to a problem by maximising or minimising a given objective function.
- Reinforcement Learning: RL is a machine learning approach where an agent learns how to act in response to either positive or negative reward signals received by interacting with its environment. The RL agent tries to maximise the expected cumulative reward received across an episode.
- Bayesian methods: statistical techniques based on Bayesian probability theory. For decision-making systems, Bayesian methods provide a framework for making decisions by combining prior beliefs with observed data, allowing for more informed and adaptive decision processes. Bayesian belief networks are a popular method well suited to decision-making under uncertainty.

The final design concept investigated was the potential need for the highest level arbitrator to have a management function in order to manage the complexity of multiple arbitration agents within any system. When combined with the concept of a HAT manager, the overall high-level structure of the system can now be summarised as shown in Figure 4.

2.5.3. *Information management and Human-Machine Interfaces*

Creating a collaborative AI system requires access to reliable, traceable and high quality data. This enables operational use of the agents, but also underpins agent training and testing during the development of both the agent and the system. This is a key challenge for any system, despite modern platforms generating significant levels of data, the sensors and data loggers were often not designed or optimised with these applications in mind. Real systems also come with significant levels of 'noise' making it essential to test collaborative AI and HAT systems against this realistic data to understand the impacts of data quality on system outcomes. This creates a demand to consider future platform's sensing and data processing capabilities, earlier and in greater detail, potentially with a greater emphasis on sensor quality, redundancy and reliability.

Assuming access to all the required data, a collaborative AI system will generate additional data and information itself, especially if AI agents actively interact and collaborate. This will augment the system's 'data-lake' further, requiring careful management to minimise data, processing and network system's scale and hence costs. If intelligent decision making is more disaggregated, then there will need to be a focus on network bandwidth, the better use of edge processing and local automation, and careful control of the amount of data required to perform certain functions.

Data and information management is also required at the Human-Machine Interface (HMI). Humans need to see the right information at the right time to make timely decisions, but will also need to understand machine proposed CoA to gain trust and insight into the systems limits and capabilities. This drives taught design of HMIs with potentially more information available and less operators, but also a need to enable 'drilling-down' into the system to provide explainability when required. The project showed, for example, the impact on operator TiA and SA of not knowing whether an agent was unable to provide a CoA or was simply still calculating options.

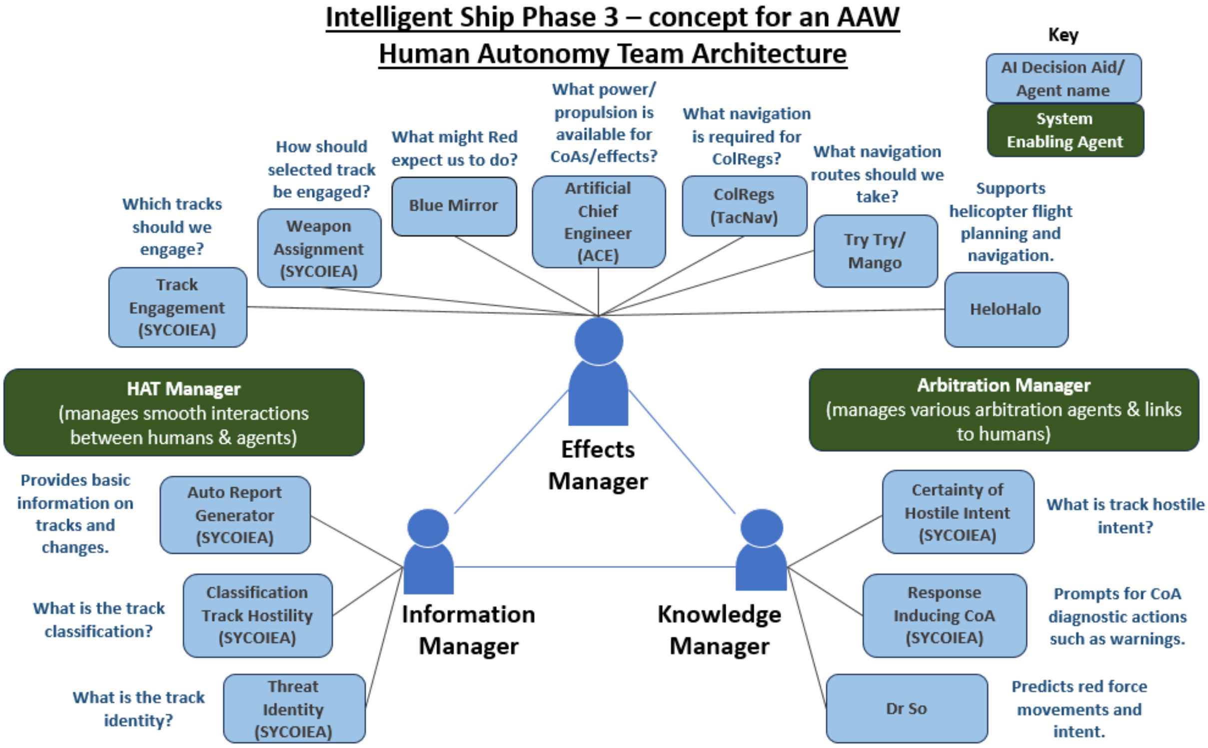

Figure 4: Intelligent Ship – proposed phase 3 AAW team structure with enabling agents added

2.6. *Modelling & simulation*

The project was underpinned by the ability to model and simulate the world in which the system operates, the data needed to operate it, and the impacts of its actions. Both cloud-based and site-based simulators were used, with cloud based solutions supporting initial remote development at each developer's site, and providing a level of high-level integration testing before installation within a physical test facility. For phase 2 this was Dstl's Command Lab, now in process of being upgraded to create the Maritime Experimentation Laboratory. This reduced demand on shared infrastructure from multiple projects, while benefiting from the higher powered computing environments available within a lab environment for final testing and evaluation. The use of open systems and standards throughout eased the complexity of integrating agents from different suppliers.

While the project found effective ways of managing its experimental needs, both through purposeful design of ISAIN and ownership of the labs and environments being used, it did not explore the wider, longer term enterprise challenges of developing, integrating and managing AI systems through-life. A range of modelling and simulation capabilities can be envisaged, from; high-level systems to undertake initial functional development and testing; real time semi-complex systems on which HAT systems can be tested and development and operational training can be undertaken, through to; high-speed, non-real-time, simulation systems to support rapid training of AI agents. This is further complicated by enabling access to data sources, linkages to other systems and potentially even to digital twins of certain systems.

Development and integration plans for an individual agent and the HAT system it is integrated into will also need to be developed. These need to define how and when it is tested, validated and assured, with what data and what level of interaction with humans, before road mapping how agents and systems are tested shore-side and when and how they are tested, and then deployed, on-board real platforms. Finally any links and opportunities to share capability with operational trainers and digital twins should be explored.

2.7. *Balancing investment, risk and maturity*

While the project focused on the design aspects on intelligent decision making systems, it also highlighted the clear need to undertake a balance of investment in which capabilities, tasks and functions would benefit from the integration of AI & HAT, considering their wider DLOD impacts.

It is unlikely that the highest level of automation and autonomy achievable in the future will be affordable, and there needs to be a pragmatic compromise between optimising designs and systems for leaner or zero crewing, making systems more robust or redundant, overall platform cost and the use of intelligent systems. Rules based control systems, or human based interventions may continue to be sufficient for many applications or tasks, be less costly and reduce the need to consider cyber-security and assurance.

Factors that will change the balance of investment with time will include:

 https://doi.org/10.24868/11145

- Relative maturity of AI and matching Automaton systems - automation, robotics and other systems will inevitably mature at different rates to AI control systems in different application areas, and in some cases the cost of matching automation could continue to be the main barrier to implementation.
- A need to balance capability with matching network and sensing needs – for agents to be 'aware' or able to learn from the environment and other constraints in their action space, they require connectivity to appropriate sensors and networks. The complexity and scale of these systems will impact platform design and cost, and also areas such as platform supportability, potentially adding workload onto human maintainers. Novel sensing approaches may help manage this challenge.
- Impacts on crewing – there is a need to prove the benefits of any automation or AI system against impacts to crew skills, roles and job engagement and satisfaction. While automation and autonomy could help to reduced crew sizes, care is needed to avoid adversely impacting recruitment and retention.
- System openness and flexibility – currently AI is good at specific tasks, it is less advanced in more general reasoning focused activities and in providing high levels of explainability critical to defence. This has led to early adoption in certain capability areas, but much longer development periods for the more sophisticated anticipated applications. This drives a need for open and flexible hosting, and modelling and simulation environments that allow future complex AI agents to be incrementally tested, trained and integrated into, and alongside, older systems.
- Managing reversionary operation – currently systems are generally designed to retain manual, human operated reversionary modes after failure or damage. This may not be ultimately possible in the future, due to the level of complexity needed to provide reversionary control, the cost and complexity of implementation, or due to the reduction or removal of crews from platforms. This will demand different approaches to safety, assurance, recovery and security of collaborative AI systems and ultimately change how the RN considers platform and systems survivability.

3. Summary & conclusions

As greater levels of intelligent decision making are integrated into platforms, it is inevitable that interactions between them will be required to maximise benefit. This requires C2 systems to consider this requirement early, and to reflect the expected needs, such as supporting simulation and modelling needs & the development of CoA Arbitration, that will enable such a transition well before the component AI capabilities mature and are incrementally deployed. This is both a technical, digitally based challenge (data, processing & networks) and a significant challenge to wider DLODs (e.g. to people, training and infrastructure).

Early consideration of human interactions, roles and skills in future HATs is also essential to ensure the best use of both human and machine skills-sets to maximise impact and capability. It is also necessary to realistically realise any practical crew reductions.

The Intelligent ship project has started to unpack some the opportunities and challenges of future naval C2 based on Collaborative AI and HAT, but can only be considered as the start of a longer journey. For example, experimentation to date has not used realistic data and hence there is limited understanding of data quality on system design, performance and human-machine interactions. Future work will be needed on how such a systems is managed and updated through-life, considering initial training, testing, validation and assurance, as well as developing systems' 'experience' through-life. There are significant HAT focused questions to explore around building trust, explainability and managing information effectively between human and machine elements of a systems, and around how and when a HAT system moves between different levels of automation.

At an enterprise level, defence ministries will need to consider and manage the wider DLOD risks of AI and autonomy, and understand modelling and simulation needs in the round, enabling initial testing of agents, systems and a HAT more generally, but also to allow through-life development and maintenance, and agent and operator training.

Finally, there is a clear need to have a high-level holistic assessment of where to focus development and investments in intelligent decision making, reflecting both capability demands, but also corresponding decision-making dynamics (speed, complexity & type), the maturity of AI techniques and realistic assessment of supporting sensing, data and automation systems required.

Acknowledgments

This project was funded by UK's Chief Scientific Advisor's (CSA) research budget and delivered by Dstl within the Autonomy research programme.

The authors would like to thank the contractor team that enthusiastically and collaboratively supported part of, or all of, the three phases of the project. This includes the following organisations: *Affect In*; *BMT; CGI; Daden; Decision Lab; Diem Analytics; Fraser Nash Consultancy; GE; HFE solutions; Modux; Montvieux; Nottingham Trent University; Polestar; Roke; Rolls-Royce; SCL; Strong Enterprises; Thales*, and; *Trimetis*. In particular the authors thank *CGI* for leading delivery and integration of phases 2 and 3.

References

Tate, A., Purvis, E., Fromm, C., Jaya-Ratnam, D. (2022). ‘The UK’s Intelligent Ship project – exploring future human-machine and machine-machine teaming’. *INEC/iSCSS proceedings*.

Unnamed. BMT. (2023). ‘Large Uncrewed Surface Vessel-LUSV – white paper’. *https://www.bmt.org/news/2023/charting-new-horizons-bmt-showcases-vision-for-lusv/*

Brethenoux, E. (2021). ‘Decision Intelligence Is the Near Future of Decision Making’. www.Gartner.com

 https://doi.org/10.24868/11145