

Human-machine teaming for Artificial Chief Engineer

M M Mansor* MEng PhD MIET

* *Rolls-Royce Solutions, UK*

* Corresponding Author. Email: maszatul.mansor@rolls-royce.com

Synopsis

There has been a marked increase in the interest for more capable optionally-crewed naval platforms over the last few years. These complex platforms will need to perform continuously in dynamically demanding scenarios, for extended periods. The capabilities of on-board digital technology and “edge processing” have also advanced considerably in recent years, providing more data and potentially previously unavailable analytically-derived information to be exploitable. These drivers have led to the development of a number of autonomous technology concepts. Artificial Chief Engineer® (ACE), an example of such technology, is a machinery control system concept that manages the power and propulsion systems and auxiliaries, targeted for uncrewed and lean-crewed vessels. The development of ACE has been previously focussed on the ability to operate at high levels of autonomy. This paper describes the development of the ACE technology in the context of variable autonomy and its teaming with human operators.

There are multiple options when it comes to allocating the autonomy level of a system operation. These options may be dependent on the specific decisions to be made, the subsequent authority of these decisions, and the role of the technology and human(s) involved, among other factors. A few key questions are raised which include: (1) How could the best traits of the human and machine be exploited to result in the most harmonious, successful operation? (2) How could trust be built in the operation models that combine both human and machines where machines are not only operating in subservient roles? (3) What are the pertinent areas in human-machine teaming that should be prioritised, or treated with additional caution (e.g., due to higher risks) to ensure advancement in the area will be beneficial instead of causing detrimental effects?

As part of DSTL’s Intelligent Ship Phase 2 programme, operating models with variable autonomy will be explored between ACE and its human operators. Starting with full autonomy and authority to make decisions, the interactions between ACE and its human operator(s) are investigated by means of discussions/feedback from a representative group of users made available through the programme. Operation of ACE at lower levels of autonomy is then explored, with human operators more involved in the decision-making process. It is expected that these human-machine teaming evaluations will provide guidance on the vital factors that need to be taken into consideration for the successful teaming between human operators and technologies like ACE. This paper concludes by identifying the high-risk or high-priority areas for future work and extracting insights gained from the direct feedback and lessons learnt from the programme. It is hoped that the work will also contribute to building trust in machines, leading to enhanced decision-making in our future missions.

Keywords: Uncrewed vessels; Lean-crewed vessels; Variable autonomy; Machinery management; Marine systems; Autonomous ships.

1. Introduction

Technology growth in digitisation has expanded considerably in recent years, providing more data and potentially previously unavailable analytically-derived information to be exploitable. Through data availability and computational advancement, more ‘intelligent’ systems are being developed. These systems are rapidly integrating into our daily lives and ship operations are no exception. In a naval context, these systems could enhance operational advantage, helping crew make better, more informed decision-making and/or shielding crew from dangerous, dull or dirty roles by performing the crew’s roles in their absence where it is digitally possible to do so. These instances necessitate the need to consider the interactions and the scope of authority between these intelligent systems (machines) and humans, not just during operation but also during design.

Author’s Biography

Maszatul Mansor is a Control Systems Engineer for Marine Electrical and Autonomy, Rolls-Royce Solutions UK. Her role is focussed on technology development, specialising in systems engineering, autonomous systems, intelligent control and optimisation. She holds a PhD in Control and Systems Engineering from the University of Sheffield.

It is generally agreed that there are multiple levels of autonomy (LoA), with no common definition within or outside the industry, although they do share similarities. For example, Parasuraman *et al.* (2000) defines ten levels of automation while the European Defence Agency's Safety and Regulations for European Unmanned Maritime Systems group defines six levels of control (UK Maritime Autonomous Systems Regulatory Working Group, 2019). Generally, as the level increases from zero or one to a higher number, human involvement decreases, and the machine is granted with higher levels of decision authority.

The level of autonomy of a system and the level of crewing for the system is interrelated albeit not the same. Uncrewed vessels are vessels that are in operation without the presence of humans on-board. Uncrewed vessels may be autonomous or controlled by humans remotely from a shore control centre, for example. Optionally crewed vessels are vessels that by design are capable of operations with and without humans on-board. A vessel may be crewed whilst being autonomous. Different parts of the vessel may also have differing LoA. For a fully-autonomous system, the system must be granted full decision authority of its operations. The LoA may differ depending on the application, scenario and task at hand. This paper will focus on integration of humans in operations where machines are provided with a degree of autonomy and authority.

Integrating humans and intelligent machines into a system is not straight-forward. The integration of intelligent technologies into systems often alters the overall system operation and may subsequently change the roles of all the actors within the system (Roth *et al.* 2019). Systems equipped with machines with high LoA may also be able to perform their own analysis based on the increasing amount of data being made available. To ensure the best overall system operation where machines and human work as a team, information-sharing and trust becomes more important. McLeod (2022) highlights the need for the human user(s) to maintain communication with the machine to ensure that the human has constant situational awareness of the operating scenario. Human users need to be able to understand what is occurring and the rationale behind actions. Similar to how humans form relationship and trust, the human users need to be able to trust the machines in their decision-making capabilities. This raises the following questions: How could trust be built in the operation models that combine both human and machines where machines are not only operating in subservient roles? How could the interface between the user and system be designed to encourage trust?

Today's methods of operation are likely to change in future, especially with the emergence of new technologies and threats. The degree and specifics of the change however is uncertain and the means to assess new ways to operate is not straightforward. Ideally, any newer operation models would adopt a mechanism which would best utilise the strength and weaknesses of both humans and machines. In other words, how could the best traits of the human and machine be exploited to result in the most harmonious, successful operation? Could the 'best' of the machines' capability be brought out through the user interface that will in turn enhance the human's capability?

This paper describes the development of the interface between a machine and a human in the context of variable autonomy, using Artificial Chief Engineer® as the machine of interest. The design approach taken to allow for human-machine teaming (HMT) operations is described, followed by lessons learnt from early engagement with potential users through DSTL¹'s Intelligent Ship programme. It is hoped that the work described in this paper, of which the scope only covers a small subset of the challenges within HMT, would enable the key sensitive areas to be identified for future work in related programmes.

2. Artificial Chief Engineer (ACE) within the Intelligent Ship Programme

2.1. Intelligent Ship

The recently completed 18-month long Intelligent Ship (IS) Phase 2 programme sought to evaluate and de-risk intelligent and automated technologies for its potential in a hybrid human and machine network (Defence Science and Technology Laboratory, 2022). These technologies have the potential to cause a step change in military decision-making, mission planning and automation. Rolls-Royce participation in this DSTL programme focusses on the development and integration of ACE into IS's Intelligent Ship Artificial Intelligence Network (ISAIN) and development of ACE in the context of HMT.

2.2. Artificial Chief Engineer

Artificial Chief Engineer® (ACE) is an autonomous machinery control system concept that manages the power, propulsion systems and auxiliaries. ACE aims to perform the role of the Marine Engineering Officer (MEO) in fully-autonomous operation, and act as a decision support system for lean- or fully-crewed operations.

¹ DSTL is a UK government agency delivering Science and technology research for the UK Ministry of Defence.

In this paper, ACE is the technology (machine) that is targeted to function within an integrated system which combines humans and machines working together.

ACE as a decision support system aims to offer enhanced (faster, more informed) human crew decision-making, for example, by ensuring availability of critical on-board systems and leveraging digital data. As an autonomous system with authority over the machinery, ACE aims to enable the path to lean-, or uncrewed vessel operations by optimising the machinery configuration for endurance or tactical advantage in the absence of humans.

Figure 1 illustrates ACE in the context of a vessel system. As part of an Autonomous Equipment Control System, alongside Equipment Health Monitoring (EHM) and Automation System (e.g. Integrated Platform Management System (IPMS)), ACE receives mission intent and other mission information from the Command Officer, Mission Manager or equivalent. This informs ACE of the mission power/propulsion requirements, priorities to optimise against, and the operating context that may influence preferences over machinery configurations. Within the Autonomous Equipment Control System, ACE receives equipment/machinery health status from the EHM and equipment/machinery status from the IPMS. This informs ACE of current or future restrictions of power and propulsion of the system, as well as their performance and capability against mission intent and requirements. ACE continuously assesses the system based on information provided and constructs the best management plan for the power, propulsion systems and auxiliaries. The ACE plan is enacted by the Automation System which directly manages to equipment/machinery.

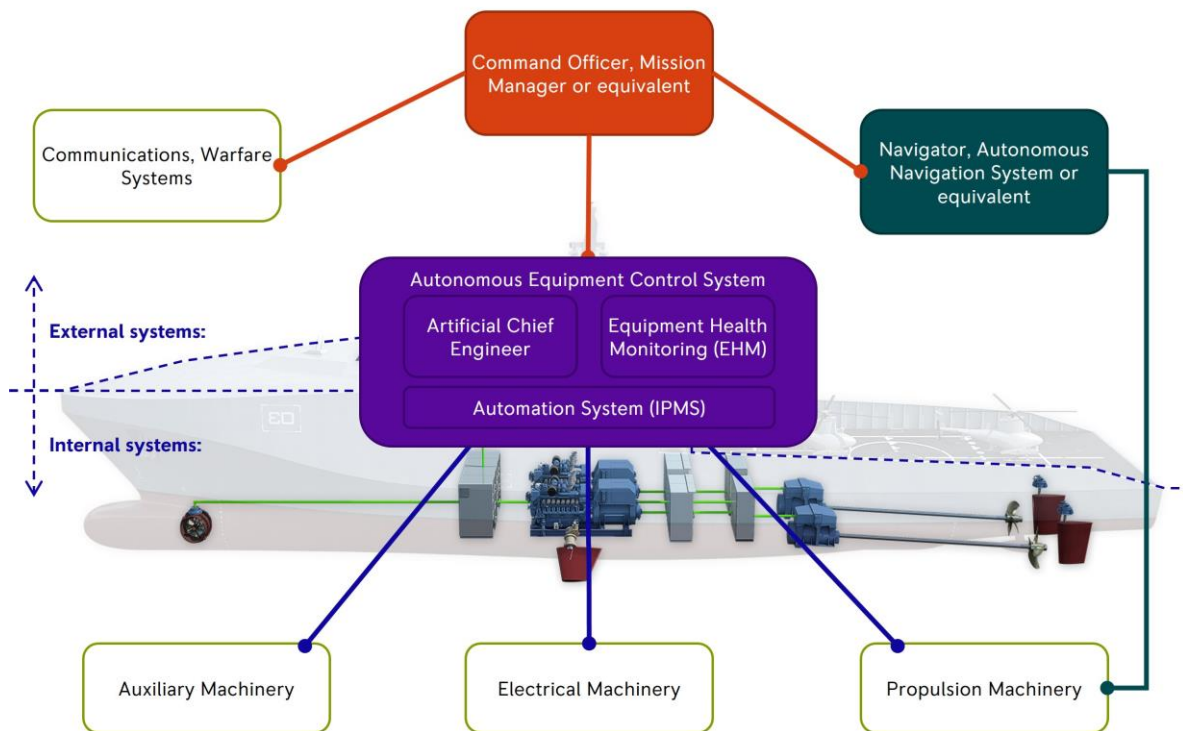


Figure 1: ACE on a vessel.

3. Designing for variable autonomy

ACE was adapted to be integrated into the ISAIN evaluation environment, which enabled ACE interactions with other machines and humans. The evaluation of ACE can take place in a stand-alone set-up independent of ISAIN or as part of ISAIN. The evaluations in this paper refer to the evaluations of ACE in the context of HMT; the wider evaluation against the IS concept is out of scope.

Based on the programme's objectives, published HMT guidelines, e.g. (Amershi, et al., 2019), and human factors considerations, the primary goal of the evaluations was to assess the effectiveness of information content for different levels of teaming. This goal was pursued by means of discussions with and feedback from Military Advisors (MAs) and current and previous serving Royal Navy employees, access to which is provided by DSTL and from within Rolls-Royce. The MAs are experienced Royal Navy staff/officers posted within DSTL.

There were two means by which feedback was obtained. One source of feedback was by means of survey responses based on user experience of ACE. A set of survey questions were used enable structured feedback for analysis. The second source of feedback was obtained by having ACE-specific discussions with MAs where

concepts, hypothesis and questions related to ACE were raised. These discussions also included a task walkthrough for a MEO.

To enable the evaluation of ACE-human teaming, different decision modes were considered. At the highest level of autonomy, ACE autonomously commands the best machinery configuration. At the lower level of autonomy, ACE could provide a veto option or seek human approval for its proposed course of action. The ACE implementation within the IS programme was designed to be able to function in three decision modes:

1. Information Only Mode
2. Veto Option Available Mode
3. Approval Required Mode

Figure 2 is a simplified illustration of Figure 1 with the different decision modes with varying LoA outlined. In Information Only mode, ACE enacts its decision autonomously directly onto the embedded machinery simulation. For Veto Option Available mode, ACE waits for a pre-defined period before enacting its decisions. During this period, a human operator may veto ACE decision and ACE will not enact its decision. In the Approval Required mode, ACE will only enact the decision once an explicit approval from the human operator has been received. Regardless of the decision modes, ACE require information in terms of mission intent and requirements and vessel/equipment statuses to be provided by a human and/or (other) machine(s).

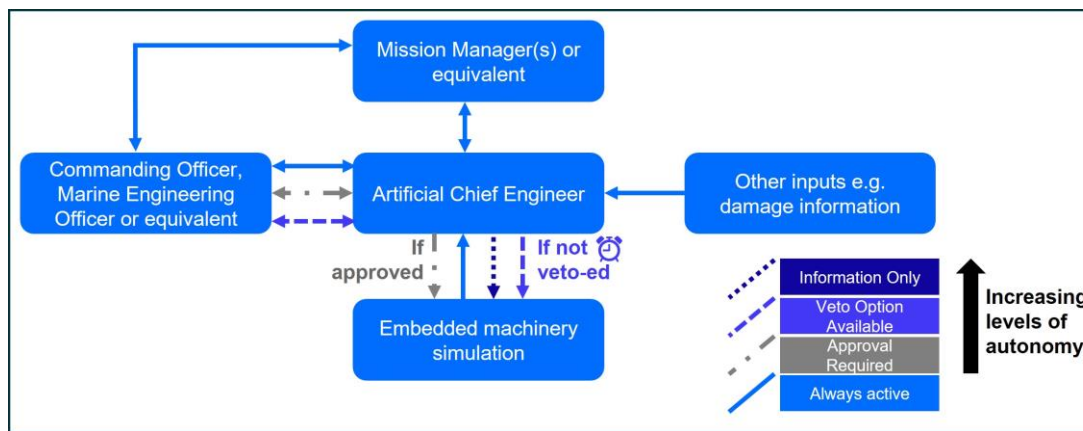


Figure 2: ACE interactions with a user at three different decision modes.

In order to enable the three decision modes to be evaluated, ACE graphical user interface (GUI) and internal algorithms were consequently designed and adapted to the IS vessel specifications and scenario and additional human-ACE teaming implications. For example, mechanisms for a human user to ‘Approve’ or ‘Veto’ a solution proposed by ACE was added into the GUI and the corresponding responses within the ACE decision-making algorithms were modified accordingly.

The purpose of the GUI is to provide a visual platform for a human to process and insert/alter relevant data for ACE. The ACE GUI comprises information on the operational context and the power and propulsion machinery status to help inform the user of the scenario, equipment status and ACE activities. The GUI design was based on the guidelines provided by DSTL (Ministry of Defence, 2021), which provided the users with a relatively familiar digital experience, and other published guidelines such as (Amershi, et al., 2019) with human factors in mind. Care was taken as to not overload the user with too much information. The GUI designs were iteratively improved following feedback from MAs during ISAIN preliminary evaluations and ACE-specific discussions, feedback from previously serving MEOs from within Rolls-Royce and based on similar/related systems within the vessels which Rolls-Royce was familiar with.

For example, Amershi, et al., (2019) advises the designer to make clear what the system of interest can do. ACE users were provided with an ACE User Guide ahead of the evaluations. In addition, tooltips were provided within the GUI to assist the user. Amershi, et al., (2019) also advises on showing contextually relevant information. Mission information was included within ACE GUI to provide information on the specific vignette which ACE is operating in.

4. Exploring variable autonomy

4.1. Scenario

A multi-event mission scenario was designed by other partners of the IS programme for use during the ISAIN evaluations. The IS was deployed as part of a task group to gather intelligence. The vignettes that were

considered were Entry Planning, Transit into contested area, Intelligence Gathering, Anti-Submarine Warfare, Air Defence and Damage Control. A similar set of vignettes were utilised during ACE stand-alone application evaluations.

The vessel of interest was simulated to be an 8000t Destroyer with a maximum speed of 26 knots and cruise speed of 18 knots. A generic design was utilised to minimise classification issues. Based on a peak hotel load of 2500kW (60Hz distribution), the vessel is modelled to be equipped with:

- One MT30-36 Gas Turbine (GT)
- Four 20V4000M53B Diesel Generators (DGs)
- Two 4.2MW electric propulsion motors.

The system operates in a COmbined Diesel-electric Or Gas (CODLOG) configuration, where the propulsion can be delivered by the GT or DGs via the propulsion motors. Both the GT and DGs can be used simultaneously to supply electrical power demands (hotel loads).

4.2. ACE evaluations

The ISAIN environment set up within the DSTL Command Lab provides a means to interconnect machines and humans to assess the IS concept where these actors work together to achieve the command aim. Within the ISAIN environment, ACE was connected to other machines (e.g. battle damage analysis system, tactical navigation system, and mission planner) and human actors (directly through ACE GUI or indirectly through other machines' user interfaces). The scenario outlined above was simulated to evaluate the interactions between machines and humans. Note that the evaluation for the general IS concept is not in the scope of this paper; this paper includes only the MAs' responses and feedback from their interactions with ACE during these evaluations. ACE was also designed to be operable without the ISAIN environment which enabled assessments independent of the variability that may be introduced by other machines or humans not directly interacting with ACE.

An example sequence of events given a change in mission requirements, are as follows (Information Only mode).

1. ACE receives new information resulting a change in mission priorities.
2. ACE updates its internal algorithms with this new information and constructs the most suitable machinery configuration(s).
3. ACE alters the machinery configuration and displays its actions (e.g. running equipment selection) and related information (e.g. platform capability advisories) onto the GUI for the human user and through ISAIN for the other intelligent machines.

The effectiveness of information content for each of the three decision modes were investigated based on the human operator roles and preferences. The users were requested to provide feedback on ACE decisions/actions, information provided within ACE and the way the information was provided by means of a set of survey questions. Examples of the survey questions are listed below.

1. Is ACE providing useful information in terms of its actions/decisions made?
2. Which information was particularly useful?
3. What are the missing key information, if any?
4. Is the GUI easy to navigate?
5. Are the prompts for human input appropriate?

4.3. User experience feedback

In Information Only Mode, ACE decision and decision rationale were reported to be mostly logical from the users' perspective. The information provided within the GUI was useful, clear and there was not too much information. The users found the ACE GUI easy to navigate around with any insertions and changes to ACE inputs easily accessible. However, there was a difference in opinion between the most useful information and what information was deemed missing. The users responded positively on the information designed to be always on display and highlighted the importance of mission context.

For evaluations which involved ACE decision modes at lower LoA, the key responses were on the insufficient information displayed to enable the human to 'Veto' or 'Approve'. Issues were also identified in the area of human factors.

Note that only Information Only Mode was used as ACE decision mode during the ISAIN evaluations. There have also been significant challenges in sourcing suitably experienced additional MAs that would be available to take part in the evaluations. As a result, the gathered data, whilst informative and provided insights that will steer future research, was insufficient for statistical significance.

4.4. Parallel discussions with potential user group

The HMT element of ACE and the ACE concept in general, were discussed in several sessions with MAs. Key takeaways included suggested areas of investigation on HMT, ACE decision-making algorithm, and ACE concept in general. The sessions confirmed some of the assumptions made for ACE operation and requirements. Additional considerations were also captured, some of which were incorporated into the designs for ACE implementation within the IS programme, while others noted for future technology development.

5. Discussion

ACE-human teaming poses an interesting challenge as there are many design options and uncertainties, as mentioned previously. Furthermore, the performance of the intelligent machines cannot be evaluated with confidence in the context of current practices. This is because the MAs are highly trained with today's systems and future roles may be filled by differently trained operators. This introduces conscious and unconscious bias towards the machines in terms of expected/assumed functionality and scope. The time provided for the MAs to familiarise themselves with the new machines are also limited although the supply of the ACE User Guide in advance of the evaluation seemed to be helpful.

Generally, ACE GUI did not overload the user with information however some content may be considered missing, especially when operating at lower LoA. In Information Only decision mode, ACE was found to be easy to interact with and most of ACE actions were logical to the users. The contextual information provided by ACE was also deemed useful albeit could be improved to enhance the situational awareness of users within the IS system. In the Veto Option Available and Approval Required decision modes however, the decision-making information provided to the user was insufficient and was difficult to interpret quickly. This points to issues relating to human factors.

The same interface design is used for all three decision modes, with exceptions of additional buttons to 'Approve'/'Veto'. The users in the Information Only decision mode may have relied on other parts of the ACE GUI to receive and understand ACE's decision-making process which is sufficient for that decision mode. The increased need for understanding, and the responsibility on the user side for the lower LoA decision modes however suggests that a different interface design may be required for these operating modes. Alternatively, similar designs could be adopted regardless of the autonomy levels of the system with the decision-making explanation to be made more prominent. The evaluation took place based on a single scenario although with multiple smaller events; the conditions of the scenario may not have necessitated the need for better information sharing at high LoA. The scenario simulated is also different to how the users would operate today, especially for stand-alone evaluations of ACE where there were no other MAs to communicate with and the user is only interacting with ACE. The visual aspects of the interaction could be improved and the ACE decision-making rationale could be made more clear. The user's need to rapidly understand the contextual scenario and make a decision may require specialist training and/or more targeted system design.

Trust is required for systems where humans and machines work together on equal terms. Good, effective information sharing may be a key element. However, care should be taken to ensure the human does not experience overload. Equally important is to not omit crucial information which may have been the case for ACE-human teaming at lower LoA. Although focussed on the HMT aspect of the overall IS concept, the wider IS study findings also highlighted the potential reduction of users' situational awareness due to delegation of authority to machines for lower threat vignettes (Jaya-Ratnam *et al.* 2022). These 'hidden' machine-machine interactions may also reduce the inclination to trust. Future designs of interfaces may require means to share otherwise hidden information to encourage trust. The ease of information interpretation by the humans should also be one of the priorities for development. This is because although the machine may provide the required information, a non-user-friendly representation may impair the potential of any rapid information absorption and hence the quality of human decisions. These extractions assume that the machines have been appropriately designed and developed to satisfy expected capabilities including the level of authority given.

It was noted during post-evaluation discussions that much like human-human interactions, a human-machine relationship may require time to allow for trust to develop (MA perspective). Next-generation future operators may also naturally have more trust in systems related to IS since AI-related technologies are more prevalent in recent years.

The assessment of the effectiveness of information content for different levels of teaming has highlighted areas for future focus and development. The usefulness of information content can be reliant on the human role facing ACE (e.g. Weapons Engineer versus MEO) however the underlying goal to satisfy command aim is prominent and any actions/behaviours should be to satisfy the command aim needs. The outcomes suggest there are potentially different information requirements for different LoA. Representation is crucial and includes visual cues and message style. These are early extractions based on a small subset of participants that may steer

future targeted set of evaluations to enable better human-machine interface designs for co-operative operating models.

There are other areas within HMT that was touched upon that was not the focus of this paper. The 'best' operation model between machines and humans is dependent on the system itself and the capabilities of its actors. There are still open questions surrounding ACE-human teaming and allocation of a decision mode, as the crewing of vessels that could be equipped with ACE is subject to change depending on the application. Based on the discussions with potential user groups and MAs, it is possible that there is no single 'best' HMT solution for ACE. It is likely that a list of HMT options can be developed based on the needs of different categorisations of applications, scenario and/or operation models as part of ACE-human teaming philosophy.

Another potential consideration for the future is to reflect on how the outcomes of the study will change should the simulation be extended. In this work, the participants are aware of the evaluation and a vignette on average runs for approximately fifteen minutes. This period is brief considering the operational time on vessels. Machines which share some of the responsibilities with its human counterpart may become advantageous during peak workloads. The wider IS concept study findings supports the potential benefits in reducing temporal demand during high-threat vignettes (Jaya-Ratnam *et al.* 2022). During quieter operational times however, the human's attention may be reduced should his/her role becomes dull or include long periods in a supervisory role. This may impact the overall team performance, especially since humans are known to perform less well when there is less to do and vigilance is required (Roth *et al.* 2019; McLeod, 2022). New responsibilities from the user role change may overcome this issue however it may be worthwhile to engage in a study which considers extended operation times and potential impacts on the humans, system, and mission aim.

6. Conclusions

The IS programme enabled new intelligent systems to be integrated and assessed into a virtual ship, enabling a low maturity prototype of ACE to be assessed and requirements adapted at the early stages of design. The programme has been a very useful platform for ACE concepts and developments to be tested and assessed, especially in the context of ACE-human teaming.

Operation of ACE at different LoA, their respective impact on human involvement and information sharing requirements as well as ACE authority to make decisions are key areas where engagement with DSTL and MAs have been very insightful. Although the number of users to enable user experience feedback was limited, key areas to focus ACE development have been suggested by the available users and separate but complementary discussions. It is noted that although the work described here is useful for future technology development, there is still much work to be done to truly incorporate HMT into the development process.

A key learning has been that lower LoA may potentially require more information to be shared with the user compared with in fully-autonomous mode. Despite the utilisation of the same underlying decision algorithms within ACE, the insufficiency of information shared with the human through the designed interface became more evident at lower LoA. This could be due to the need for the user to rapidly gain an understanding of the changing scenario and the information shared should be made clearer and more digestible. It may also be possible that additional scrutiny is imposed on ACE when the users hold the final authority of decisions, whereby altering the level of information and trust required by the humans to authorise ACE actions. This may be a useful consideration for the development of related systems for others in the industry.

Areas within ACE decision-making and ACE-human teaming for further considerations and focus have been identified and will be developed. Given the feedback from the MAs during the ISAIN and stand-alone evaluations, the focus of future work will include emphasis on the decision authority awarded for different conditions and improving the disclosure of ACE decision-making process through human factors considerations especially for operation at lower LoA. Continuous engagement with potential operators of ACE will be sought to validate this.

Acknowledgements

The author would like to acknowledge the support from DSTL technical partner and wider team including the military advisors. The author also thanks the Rolls-Royce team, specifically Pawel Merta, Gary Doherty, Roger Bovill, Rohit Nag, Eshan Rajabally and Alexander Finch, for their contributions in the programme. This work was supported by the DSTL under Grant DSTLX1000153693.

References

Amershi, S., Weld, D., Vorvoreanu, M., Fournery, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gill, R., Horvitz, E. (2019). Guidelines for Human-AI Interaction. In

- Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019 May 2 (1-13).
- Defence Science and Technology Laboratory 2022. *Competition document: Intelligent Ship Phase 2*. Defence and Security Accelerator. Available from: <https://www.gov.uk/government/publications/competition-intelligent-ship-phase-2/competition-document-intelligent-ship-phase-2> [Accessed March 2022].
- Jaya-Ratnam, D., Richards, D., Froom, C. (2021). Intelligent Ship Phase 2 – D22: Event 4 Report (Issue 2.0) 300000000041355/DEL/D22 UK OFFICIAL. CGI, DIEM, Thales, Decision Lab.
- McLeod, R. (2022). Human Factors in Highly Automated Systems. Chartered Institute of Ergonomics & Human Factors.
- Ministry of Defence 2021, *MOD.UK Design System* [online]. MoD UK. Available from: <https://github.com/defencedigital/mod-uk-design-system> [Retrieved December 8, 2021].
- Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A Systems and Humans*, 30(3), 286-297.
- Roth, E. M., Sushereba, C., Militello, L. G., DiIulio, J., & Ernst, K. (2019). Function Allocation Considerations in the Era of Human Autonomy Teaming. *Journal of Cognitive Engineering and Decision Making*, 13(4), 199-220.
- UK Maritime Autonomous Systems Regulatory Working Group. (2019). Maritime Autonomous Surface Ships (MASS) UK Industry Conduct Principles and Code of Practice (Version 3). Maritime UK.