

This is a repository copy of *Implementing autonomy in nuclear robotics; an experience-informed review of applying SACE*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/208165/>

Version: Accepted Version

Proceedings Paper:

Mulvana, Phillip, Marsland, Lacey-Jo, Boden, Tom et al. (3 more authors) (Accepted: 2024) *Implementing autonomy in nuclear robotics; an experience-informed review of applying SACE*. In: *Safety-Critical Systems Symposium (SSS'24)*. Safety Critical Systems Symposium, 13-15 Feb 2024 , GBR (In Press)

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Implementing autonomy in nuclear robotics; an experience-informed review of applying SACE.

Phillip Mulvana¹, Lacey-Jo Marsland², Tom Boden¹, Guy Burroughes¹, Richard Hawkins³, Matt Osborne³

¹RAICo, UK Atomic Energy Authority

²AtkinsRéalis ,

³Assuring Autonomy International Programme, University of York

Abstract Nuclear decommissioning is a complex, hazardous, and time-consuming process that requires highly skilled and trained operators. To address the workforce bottleneck and the growing inventory of nuclear materials, a Robotic Glovebox with AI capability that can assist in the preparation and processing of nuclear material is being developed. This innovative solution can enable safer, more efficient, and continuous decommissioning operations. To support the adoption of this new technology it is necessary to develop a safety case for the system. In this paper we describe how we have used an autonomous system safety case process (the SACE approach) to generate confidence in our initial AI glovebox design. This safety case example is being provided as input it into the Office for Nuclear Regulation (ONR) regulatory innovation sandbox and should help to establish a new paradigm in safety cases for autonomous systems in nuclear environments.

1 Introduction

Gloveboxes are sealed containers that allow safe manipulation of hazardous materials in a controlled atmosphere. They are widely used in the nuclear industry for various tasks involving radiological samples. However, gloveboxes differ in their design, material, and number of operators, depending on the specific tasks and requirements. This also affects the safety risks and challenges of glovebox operations.

The operators who work in gloveboxes are highly trained and skilled professionals who follow strict procedures and protocols. However, their training is costly and time-consuming, and their demand often exceeds their supply. This creates a backlog of work and increases the risk of human error and fatigue. Therefore, there is a need for automation of glovebox processes to reduce human-related inconsistencies, enhance safety performance, and increase productivity.

Automation of glovebox processes can provide several benefits for the nuclear industry. It can improve the safety of operators by minimizing their exposure to radiation and other hazards. It can reduce the environmental impact of glovebox operations by preventing the spread of contamination and ensuring proper waste management. It can increase the speed and efficiency of glovebox tasks by enabling

continuous and consistent operation without human intervention. It can also facilitate the decommissioning of legacy nuclear sites by accelerating the processing of hazardous materials.

Systems that use AI can exploit its capability to infer desired options in new, unseen settings. This is enabled by training the machine on a large corpus of media showing the target in known settings, so that it can ‘understand’ what the target is in new settings. The application of AI in this case seeks to train a system to identify a radiological sample in the glovebox environment. This allows the system to take the sample through a process of identifying, cutting, cleaning, and packaging, while accounting for dynamic environmental variables such as size, shape, and material properties.

This capability allows the system to continuously process samples without the constraints highlighted previously (e.g. operator availability, fatigue, training). This is an essential part of managing current and future demands in nuclear decommissioning.

However, autonomous decision-making introduces significant system complexity, which leads to potential hazards with causes that do not exist in conventional systems, such as:

- Object misclassification leading to execution of unsafe actions.
- Unintended interactions with the environment leading to unsafe outcomes.
- Failure to identify an unsafe system state.
- Exploitation of the desired process leading to unsafe outcomes.

AI elements of the system cannot presently lead directly to harm, but they can act as a secondary cause in ways that differ from the automated software that they replace.

These hazards and their root causes are predictably of concern to regulators, industry, and the public alike. The success of this project therefore depends on its ability to assuage the concerns of these parties.

Much like the hazards themselves, there are a defined number of secondary factors for the onset of these hazards, which include challenges such as insufficient/unrepresentative training, skewed/tampered training data or poor goal specification.

At time of writing there is no mature regulatory regime for autonomous systems in the nuclear environment. The Office for Nuclear Regulation (ONR) has clearly identified the same benefits and challenges as the RAICo technical leadership and has made space within their ‘regulatory sandbox’ environment to explore the challenges of developing safety assurance arguments for AI in the nuclear environment. This resulted in the formation of a panel of experts and stakeholders who stood to either shape or become a recipient of these autonomous systems.

This paper explores the process of creating a safety case for a robotic glovebox utilising AI to process nuclear material (referred to as the “Robotic Glovebox” hereafter). The development of the Robotic Glovebox aims to reduce risk to operators by eliminating the need for them to be present in high-risk areas and perform manual

operations. This paper accomplishes this by hypothesising the addition of goal-oriented AI to a real-world system that is currently trialling robotic decommissioning. This Safety Case has been developed in line with the Guidance on the Safety Assurance of Autonomous Systems in Complex Environments (SACE) (*Hawkins, et al., July 2022*), which relies on the non-autonomous aspects of the system design and safety assurance running concurrently to produce a complete safety case for an autonomous system.

The paper is structured as follows;

Section 2 considers the background and context of AI and robotics in nuclear applications.

Section 3 discusses the use cases, application and expected deliverables.

Section 4 outlines the selection of the safety case approach.

Section 5 details the implementation of SACE.

Section 6 draws conclusions on the nature of the approach in this context.

Section 7 details potential improvements to the approach.

Section 8 considers possible future opportunities.

2 Background

2.1 AI in the nuclear context

The nuclear domain moves with incredible intentionality with a relatively concrete position on risk acceptance, favouring technologies that have been tried and tested for a given application over minor increases in performance at the expense of adding novelty. Both the RAICo programme and the regulator, however, recognise the expansion of AI into all domains, whether through the automated generation of technical literature or through system design itself. In multiple workshops held with regulators, site license companies and supply chain it was recognised that assessing the impact of highly innovative technologies in AI would be proportionate to the benefits it may yield.

This panel of experts ultimately agreed that creating an exploratory safety case could test the limits of existing technical and governance regimes in the context of AI. This group pursued the RAICo proposal to create a quasi-real-world test case representative of real hardware and nuclear use cases with the notional application of goal-oriented AI.

3 Use Case

3.1 Use Case Description

This goal-oriented AI system would make use of two types of cameras to identify objects within the glovebox; wide field-of-view cameras attached to the glovebox and an object-focused camera situated on the end of the robot arm. The system will first use these to locate the target in its environment, then make decisions which allow the sample to move through the process.

These capabilities would be based on existing capabilities that RAICo possess in deep learning for machine vision which are currently used for research-based pick and place tasks. This capability has been created from a human-labelled training set which allows the system to learn the shape, size, and characteristics of several objects found in a nuclear-industrial environment and would thus be carried forward into the strawman.

3.2 The RAICo Strawman Glovebox

For the purposes of the strawman glovebox, this AI layer would be theoretically applied to an existing hardware project built to test remote sample processing using a modified glovebox fitted with two collaborative robot arms (Figure 2). This was done intentionally to reflect a generic glovebox in industry and ensure that this could serve as a base for the implementation of an autonomous system in the context of SACE.



Fig 1. CAD Renders of the RAICo Strawman Glovebox

This setup proved the ideal vehicle for a strawman safety analysis, allowing safety engineers to build upon its outputs without causing any real-world project or safety

risk that may result from the experimental nature of the project. The credibility of the safety case is strengthened owing to its ability to make use of a genuine bill of materials, conduct real analyses of the conventional elements and leverage the expertise of a real combined project team, as illustrated in Figure 3. This conventional glovebox will operate in an environment that has been designed for representative operations to handle the processing of low-activity nuclear samples which lay below a threshold that requires strong nuclear diligence; this uniquely positions the project to conduct safe research and development while retaining the authenticity of the strawman.

The degree to which this project is abstracted from a true deployment does, however, bring challenges; simplifications such as a high degree of control over the environment and object shape/size may not be representative of a final deployment of this system. Despite the potential limitations, it is believed this still offers a strong opportunity to both template safety case approaches for AS and demonstrate a credible regulatory position for the deployment of a complex system.

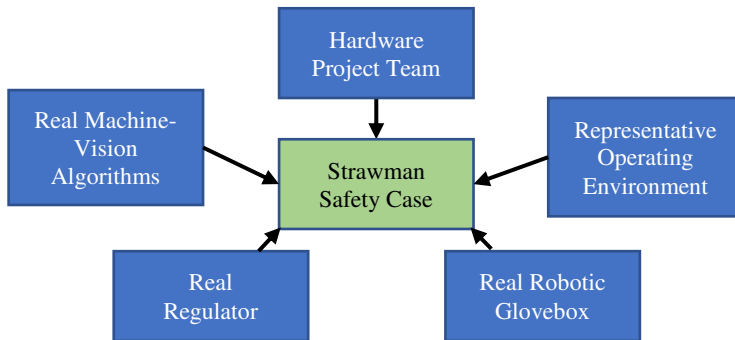


Fig 2. Synthesising of Elements to Create the Safety Case

3.3 *Expected deliverables*

This project offered the ability to assess a number of technical and governance approaches, here primary aims included:

- Evaluating a safety case approach for AI systems in nuclear and other safety-critical environments.
- The assessment of safety and regulatory considerations for autonomous systems in nuclear applications.
- Developing a safety case model for future autonomous systems within the RAICo programme.
- Investigating compatibility with current nuclear operational practices.

Secondary aims include:

- Identifying technical limitations and risk tolerance for autonomous systems in nuclear settings.
- Discovering opportunities for creating new safety assessment tools for autonomous systems in nuclear contexts.
- Building expertise in this emerging field of safety engineering.

4 Selecting an Appropriate Safety Case Approach

The project began by analysing regulatory and standards frameworks for ensuring the safe development of AI systems in high-integrity contexts. This analysis included a review of safety development lifecycles, comparing approaches such as UL4600 (Standards, 2023) for automotive development and *SCSC Safety Assurance Objectives SCSC-135B for general AI safety guidance*.

Evaluation of these options focussed on three key criteria:

- The extent to which they offered a comprehensive development lifecycle for the autonomous elements of a system, with an emphasis on compatibility with existing safety practices and the ability to provide robust safety decision-making processes.
- Independence from industry-specific biases, ensuring a fair assessment of risk reduction potential irrespective of the approach's origin.
- Compatibility with existing approaches, with the intent of managing challenges related to organisational inertia such as those which may make organisations inclined to remain within existing domain frameworks.

In summary, SACE emerged as a strong, independent framework suitable for building a safety case. It demonstrated adaptability to different systems and implementations, with a track record in various safety-critical domains such as medical, aerospace, and automotive. SACE's development was guided by respected institutions and backed by the Assuring Autonomy International Programme, making it the optimal choice for this project.

SACE (Hawkins, *et al.* 2022) is a methodology that provides detailed guidance on the creation of a compelling safety case for an autonomous system (AS). It comprises a set of safety case patterns and a process for systematically integrating safety assurance into the development of the AS and for generating the evidence base for explicitly justifying the acceptable safety of the AS. SACE builds on existing established system safety assurance processes and defines modifications, enhancements and additions to specifically deal with the safety assurance challenges of an autonomous system operating in a complex environment. As such, SACE is intended to complement activities undertaken as part of existing systems engineering and safety assurance processes.

5 Application of SACE

The safety case project has initiated the SACE process (Figure 4) and, at time of writing, has completed the Operating Context Assurance and Hazardous Scenarios Identification stages.

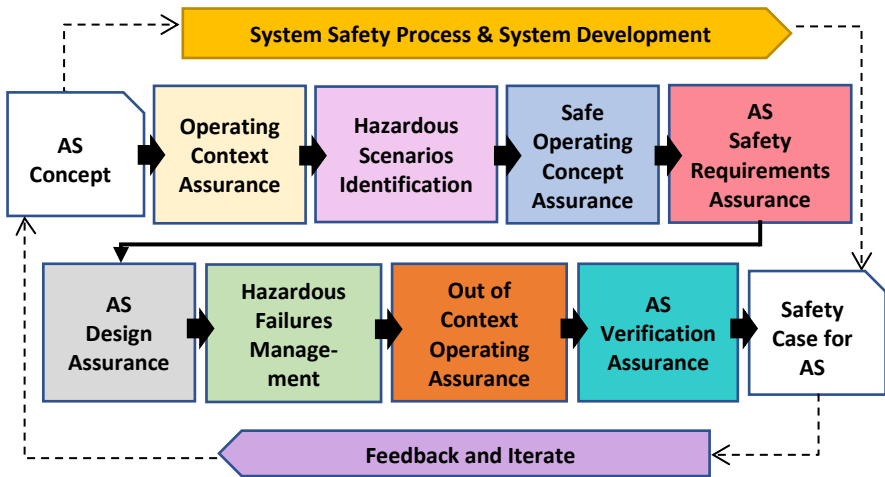


Fig 3. Outline of SACE Process

The guidance suggests that these early stages be undertaken at the functional level. Upon application within the nuclear context, however, a different reality is revealed. Our use case has suggested that, as autonomy is used to simplify operation, reduce complexity, or expand capability, very few (if any) autonomous systems would be designed from a truly blank slate, particularly in nuclear applications. In the context of this project, for example, there were constraints around hardware from the outset. Enabling the incorporation of this information in the early stages of safety case development may limit the number of Operating Scenarios and subsequent Hazardous Scenarios, thus enabling efficient utilisation of resource in authoring the Safety Case and a greater focus on the elements which could truly present hazards.

5.1 Phase One – Operating Context Assurance

Guidance from the ISA suggested there were strong and heavily iterative connections between these early phases of the SACE and our experience in implementing it confirmed this.

This became particularly apparent when developing the Context of Safe Operation (CSO) during Phase 1. The CSO consists of the Operational Domain Model

(ODM), the Autonomous Capabilities and the Operating Scenarios for the autonomous system. The development of any one of these areas directly influences the other two and the effect of this can be a process of developing one area to maturity only to find that a small change in another area forces redevelopment of the initial document. Whilst challenging, it is these connections which represent the true nature of autonomous decision making and the contextual importance of elements like the operating domain on the AS capability. No revised work was considered lost work in this process. With all changes logged and rationales captured, these were then used to inform the validation reports for each of the activities. In terms of providing a safety case that is complete, compelling, and accurate, reflecting on all these changes is essential.

As well as highlighting the relationships within the CSO, identifying Operating Scenarios also involved the consideration of existing approaches to nuclear safety cases and subsequent SACE activities. For example, equipment failures should not be considered as Operating Scenarios as these would, instead, be captured within a Failure Modes and Effects Analysis (FMEA). Additionally, outcomes of decisions made by the autonomous system (e.g. Robot collides with Object) were not considered Operating Scenarios as they are elicited as outcomes of the AS Decision Analysis.

5.2 Phase Two – Hazardous Scenarios Identification

This phase marked the first point at which safety issues began to be elicited and provided a concrete idea about the effectivity of our prior analyses and specification for use in future phases.

The Operating Scenarios identified within Stage 1 present the decision points of the autonomous system. Options are then elicited for each of these decision points to outline the potential actions of the autonomous system in response to the scenario. These decisions are analysed in the context of real-world state and system belief state; disparities between the two belief states present a space in which hazards can arise and enable the identified of hazardous scenarios associated with operation of the autonomous system.

It is noted that this decision analysis is similar to a Hazard and Operability (HAZOP) study in that it aims to identify the points at which a hazard may arise within a process/set of actions. It is, therefore, advisable that this analysis is undertaken with a multi-discipline team to capture knowledge gaps.

When applied as written in SACE, the decisions analysis identified a large number of potential outcomes. It was, therefore, beneficial to apply a severity rating to the outcomes (as recommended within SACE) to enable straightforward distinction between seriously hazardous scenarios and those which are simply inconvenient. This allowed more effective analysis to take place by identifying scenarios that led to harm to people.

Hazardous scenarios in SACE are elicited by applying environmental conditions that may influence the outcome of an autonomous decision. The Robotic Glovebox offers a self-contained micro-environment of pressure, humidity, lighting etc. meaning it is far more controllable than may otherwise be the case. This limited the number of applicable environmental conditions that could be introduced to the system that actively posed a hazard to a person. The constrained environment of the Robotic Glovebox is something that also became apparent when creating the severity scale for the hazardous scenarios; the extent to which hazards can impact a person was limited to potential wounding/laceration, indirect radiation exposure from glovebox operations, and to direct exposure to radiation by a complete system failure.

It is noted that SACE does not provide guidance around the means of validating that all appropriate environmental conditions have been considered in the identification of hazardous scenarios and, as such, there is uncertainty around the completion of the AS decision analysis. Completeness of the AS decision analysis for the robotic glovebox was demonstrated through peer review rather than taking a systematic approach.

The identification of these hazardous scenarios is essential because traditional safety analysis does not fully analyse the consequence or hazardous output of decisions that an AS can make. Decision analysis does this in a comprehensive way that accounts for misunderstandings within the AS.

By identifying the high severity hazardous scenarios, proper requirements can then be written that mitigate these scenarios, through either AS limitations or hardware choices.

5.3 Phase Three – Safe Operating Concept

The Safe Operating Concept (SOC) encompasses system-level safety requirements that dictate how an autonomous system should behave to mitigate hazardous events. Progressing through Phases One and Two, we observed the emergence of requirements when discussing the interactions between system elements. While there was a temptation to move directly to the requirements and hardware allocation, completing the preceding phases established a robust foundation for decision-making.

The link between hazardous scenarios, requirements and design choices signifies comprehensive hazard identification and assessment. This method instils greater confidence in the final safety case.

Implementing hardware decisions prior to the identification of hazardous scenarios and SOC may unnecessarily constrain the autonomous capability of the AS. These constraints limit the number of available options at the design stage and could, therefore, stifle innovation.

As previously discussed, SACE Phases One and Two favour functional-level analyses to influence safety-driven design. Making hardware decisions without a

clear grasp of hazardous scenarios and associated requirements can lead to sub-optimal design choices. Detailed knowledge of hazards and requirements is essential for effective optioneering studies.

Example 1:

During the AS decision analysis, the following hazardous scenario was identified:

<The Robot Arm is travelling with an Object> <with the Robot Arm Mounting Frame in the path of the Robot Arm> AND <The Robot Arm continues along its current path>.

This could challenge the integrity of the Mounting Frame and result in a breach of glovebox containment. From this, the following requirement was produced as part of the SOC:

“While travelling, the Robot Arm shall maintain safe separation from the Robot Arm Mounting Frame”.

Example 2:

During the AS decision analysis, the following hazardous scenario was identified:

<The Robot Arm is Grasping a radioactive Object> <with the incorrect End Effector attached> AND <The Robot Arm continues with Grasp>.

This could result in the spread of contamination within the glovebox through contamination of the End Effector and/or damage to the radio-active Object. From this, the following requirement was produced as part of the SOC:

“When Grasping an Object, the Robot Arm shall ensure that the correct end effector is attached”.

While many of the requirements derived from the hazardous scenarios can be decomposed and addressed through hardware decisions, some may pertain to machine learning function in SACE Phase Four. For instance, Example 1 (above) could lead to safety requirements related to the Machine Vision System. The safety assurance of such elements may be undertaken in accordance with the guidance for Assurance of Machine Learning for use in Autonomous Systems (AMLAS). An understanding of AMLAS is, therefore, recommended when undertaking implementation of SACE.

Requirements within the SOC prioritise prevention of the hazardous scenario over protection or mitigation. This aligns with the ONR Safety Assessment Principles with respect to the five levels of defence in depth (*Office for Nuclear Regulation, July 2019*); safety functions associated with normal operations typically relate to Levels 1 and 2 of the Defence in Depth hierarchy. The SOC also includes the identification of Reduced Operating Domains (RODs) in response to system or component failure modes, making it advisable to conduct a functional failure analysis at the beginning of SACE Phase Three.

6 Conclusions

6.1 *Nature of the Strawman Safety Case and the Need to Narrow the Project Scope*

Creating a comprehensive safety argument through the SACE framework is a resource-intensive endeavour that exceeds the scope of this project. Consequently, we've prioritised certain framework elements and documented our progress and lessons learned as we gain proficiency in each step.

Rather than duplicating well-established analyses for the entire operational scope of the AS, we focussed on four specific use cases of interest. This approach allowed us to gain a broader understanding of the SACE process while efficiently testing various components of the Safety Case framework.

The decision to narrow the scope occurred when identifying AS Operating Scenarios. The complexity of the Robotic Glovebox resulted in a vast number of activity diagrams and potential use cases. Identifying Operating Scenarios for the full scope was deemed too time-consuming and labour-intensive within the context of this project. The decision to reduce scope has meant that hazard analysis has not been performed for some tasks of the Robotic Glovebox. This however has had limited impact as some of the Use Cases carried forward (e.g. Travel to Destination and Grasp Object) were representative of many other activities within the Safety Case.

The "Prepare to Cut sample" operation was retained in the down-selection process as it was a clear candidate for presenting hazards across the AS and traditional safety elements. Both the project team and the regulators recognised the value of assessing its associated risks. Additionally, the "Object Classification" use case was retained due to the recognised hazard of object misclassification and its potential impact on other use cases.

Lastly, autonomously posting items out of the glovebox was known to involve hazardous operations, but it was not considered feasible given the current state of technology. It is considered a future candidate for more detailed assessment.

6.2 *Lessons Learned in Applying SACE*

A critical outcome of this project is its ability to inform the industry of potential strengths and weaknesses with both the SACE framework and the process of delivering an autonomous system at large. These strengths and weaknesses are expressed below in a list of lessons learned;

- The relationship between SACE and conventional system safety activities should be considered prior to any of the implementation activities. There are natural synergies between the two lifecycle approaches, and without explicit consideration of these projects may find themselves either delayed as they wait for information required from the conventional design, or overly constraining the system because of requirements over the AS.
- The SACE process cannot be implemented after hardware is allocated or a system has been built. The iterative process of defining an ODM, the operating scenarios and the autonomous capabilities directly drives design. There can be no mistaking the extensive influence that the safety assurance process has over system design. Dependent on the aspirations of the system or the complexity of the environment, the safety elements may either shape the design or preclude it entirely. It is in the design organisations best interests to be aware of this at project initiation.
- The application of AI to a given system may be at best detrimental to the process itself or at worst objectively add unsafety to an otherwise robust system. In summary, the desire to make a process autonomous should not override the drawbacks of doing so and, as previously stated, a robust justification should be provided for the implementation of an autonomous system. In a significant number of cases, it became apparent that autonomous decisions and functionality could easily be replaced with automated (i.e. pre-defined closed loop) behaviours.
- The securing of an effective governance framework in which to operate the autonomous system is as challenging and important as developing the technology and the safety case itself. Acceptance, to a large extent, is driven by trust factors such as transparency, intentionality, and education. Acceptance arriving from these factors should be considered with equal weighting to the safety case itself.
- An AI safety case is an additional stream of work, and whilst complimentary to the existing safety case, requires at least the same level of effort. Like all processes, generating talent and corporate memory for the AS design process will inevitably lead to time saving in the long term. Its short-term application can be expected to approximately double the time invested into safety assurance.
- Irrespective of intention to develop autonomous capabilities, there is benefit to all organisations undertaking safety evaluation in complex environments in considering the creation of a domain specific ODD. The ability to interrogate an ODD against specific conditions and design considerations is among the most powerful

elements of the SACE process and yielded consideration of failures and conditions outside the AS that were not present in previous dialogue.

- The use of specific language has been demonstrably vital within our application of SACE. Clearly defining the scope of autonomous operations and the associated decision points enables accurate identification of potential outcomes/hazardous scenarios. It is, therefore, recommended that terms are clearly defined and consistently applied throughout documentation. Failure to define semantically correct definitions may directly lead to incorrect safety requirements.
- The Goal-Structure Notation (GSN) that underpins the SACE approach calls for the validation of the output of each phase. There is, however, no specific guidance for performing this validation at any phase. As previously noted, this lack of clarity has the potential to introduce uncertainty around the accuracy and completeness of work undertaken within a given task/phase. The validation activities did offer an opportunity to review the contents of the phases' outputs in detail of which the importance cannot be understated. It also allowed for an opportunity to justify the decisions and progression of the Robotic Glovebox Safety Case as it was being developed.

7 The Future of the AI Glovebox Project

The project is approximately 40% complete and we have recognised a significant amount of effort required to complete the process is front-loaded into phases 1-3, with these steps forming the foundation of future activity. The project will continue to move forward with its reduced scope to assess overall fitness for purpose of the SACE framework within the nuclear sector. Should there be a desire to continue, this project scope can be returned and omitted use cases can be analysed for remaining hazardous scenarios. By this point, the process for analysing, mitigating, and implementing the remaining scope of the Robotic Glovebox Safety Case will be understood by the completion of the reduced scope. Updating the remainder of the Safety Case should be a comparatively low effort endeavour given no unexpected or unmanageable risks.

When complete it is hoped this project will have produced the following outputs:

- A template for implementing SACE in future robotic systems in a nuclear environment.
- A gap analysis of current regulatory requirements against those produced during the SACE and traditional system safety case process.

- A body of knowledge to be used by other safety engineers undertaking this work to provide detail over the intention and rationale of various steps within the SACE process.
- An accord with the ONR as to the viability of autonomous systems in the nuclear regulatory regime

8 Opportunities for Improving our Approach.

Progressing through the SACE process has highlighted some interesting technical and process-related work that could be pursued to either improve the process or generate a higher level of confidence in its outputs. A list of these elements can be found below.

8.1 Implementation Opportunities

- Considering responsible innovation as a key element of the complete safety argument for the system alongside conventional and AS safety processes.
- Detailing the process of interacting with the regulator. Creating a short document listing concerns from the regulator and potential answers from within SACE could be used to guide early-stage conversations in new or similar domains.
- Forming a panel of independent assessors comprised of customer and regulator representatives to conduct the assurance steps of each part of the SACE process.
- Automating the AS analyses that support the SACE process such that expected links between these elements are automatically formed ready for population with data.

8.2 Project/Hardware Opportunities

- Add additional complexity to the system by considering the fully autonomous interaction of two robot arms within the glovebox, ensuring their safe operation as independent parts of an autonomous whole.
- Instantiating this safety case in contexts where the operator doesn't have full control of its ODM. This would mean we have to accommodate a wider range of parameters within the safety analyses.

- Incorporation of new techniques/technologies such as chemical sampling replacing the need for swabbing.

8.3 Regulatory Opportunities

The project continues to work with the regulator to identify opportunities to better ease the coming of AS in the context of nuclear safety. We plan to better define these opportunities in future papers but highlight three key learnings below:

- Existing guidance for nuclear safety cases focusses on sites and facilities rather than smaller systems. If this approach is accepted and implemented more widely, guidance around the development of system safety cases would facilitate innovation across the industry.
- As a result of point one, this autonomous system and the approach used to define its safety characteristics has some incompatibilities with the current license conditions. We have identified there may be some opportunities to draw parallels between the two, but ultimately this would require a separate more focused study.
- If this approach is accepted and implemented more widely, it would be beneficial for Nuclear Safety Committees to have representatives with a detailed understanding of AI and the implementation of Nuclear Safety Cases for Autonomous Systems. Further guidance may enhance this.

9 References

Hawkins, et al., July 2022. Guidance on the Safety Assurance of Autonomous Systems in Complex Environments. *Assuring Autonomy International Programme (AAIP) The University of York, York, 2022* <https://www.york.ac.uk/assuring-autonomy/guidance/sace/sace-download/>

Office for Nuclear Regulation, July 2019. Nuclear Safety Technical Assessment Guide – NS-TAST-GD-094 Revision 2.2. *Office for Nuclear Regulation, 2019*

UL Standards and Engagement. (2023). Presenting the Standard for Safety for the Evaluation of Autonomous Vehicles and Other Products. <https://ulse.org/ul-standards-engagement/presenting-standard-safety-evaluation-autonomous-vehicles-and-other-1>

SCSC Safety Assurance Objectives SC135b for general AI safety guidance. Version 3.0. *Safety of Autonomous Systems Working Group.*

10 Acknowledgement

Funding and support from Robotics and AI Collaboration (RAICo)