



Summary Report

**Regulatory sandboxing: accelerating safe and secure deployment of artificial intelligence in nuclear installations**

## Summary Report

# Regulatory sandboxing: accelerating safe and secure deployment of artificial intelligence in nuclear installations

**In collaboration with:**



Doc. Ref. No.: ONR-FUT-001

Issue No.: 0

Publication Date: April 2026

# Contents

Contents.....	3
Acknowledgements .....	4
Disclaimer .....	4
Executive summary .....	5
1. Introduction.....	8
2. Project background.....	8
2.1 Context .....	8
2.2 Purpose of this project .....	11
2.3 Links with previous work on AI at national level .....	12
2.4 Links with cooperation on AI at international level .....	13
2.5 Project scope .....	14
2.6 Use cases .....	14
3. Project delivery structure .....	15
3.1 Overall project plan .....	15
3.2 Enabling activities .....	17
3.3 Regulatory sandboxing process.....	17
4. Technical findings and project learning .....	20
4.1 Technical findings .....	20
4.2 Project lessons learnt.....	20
5. Project outcomes.....	22
6. Conclusions and next steps.....	23
7. List of abbreviations.....	25
Annex 1 – Reference documents .....	26
Annex 2 – Technical Findings .....	28
Annex 3 – Evolution of innovation engagement .....	34
Annex 4 – Summary of project AI use cases.....	35
Annex 5 – Project impact and benefit realisation .....	38

## Acknowledgements

This work, led by the Office for Nuclear Regulation (ONR), was funded through the Regulatory Innovation Office (part of Department of Science, Innovation and Technology) through the AI Capability Fund. The delivery of the project benefits would not have been possible without the proactive collaboration, practical examples and technical expertise of organisations including:

- Defence Nuclear Safety Regulator
- Environment Agency
- Health and Safety Executive
- NNB Generation Company (HPC) Limited
- Nuclear Restoration Services Limited
- Rolls-Royce Submarines Ltd
- Sellafield Limited
- TÜV SÜD Ltd
- University of York / Centre for Assuring Autonomy

## Disclaimer

The views presented in this report are the Office for Nuclear Regulation (ONR)'s. The Environment Agency and Health and Safety Executive collaborated with the ONR authors on the content of this research report. Whilst project partners mentioned in the acknowledgement section were consulted on this report, because of the project timescales, they provided only limited input.

The findings presented in this report are based on the discussions in the regulatory sandboxing, focusing on two specific applications of AI and based on the information available when compiling the report from the project partners. As a result, the report does not represent a formal regulatory position. Whilst this work helps to inform regulatory principles for future AI applications, any future regulatory decision making (e.g. permissioning or compliance assessment) associated with systems containing AI used in the nuclear sector will be assessed by an independent regulatory team in line with ONR and Environment Agency regulatory policy, principals and processes.

# Executive summary

This report summarises a regulatory sandboxing project to accelerate the safe and secure deployment of systems containing artificial intelligence (AI) in nuclear installations.

## UK nuclear regulation and AI

The Office for Nuclear Regulation (ONR), the Environment Agency, and Defence Nuclear Safety Regulator recognise the potential for AI in nuclear installations to support the UK government's growth agenda on nuclear, including its net zero ambitions. The nuclear regulators also recognise the benefit AI can offer from a nuclear safety, security, safeguards, nuclear site health and safety, and environmental perspective, for example in terms of accelerating hazard and risk reduction for legacy facilities through innovative approaches.

To develop our approach to regulating AI, we have consulted with a wide range of stakeholders, undertaken research, engaged with government, dutyholders, other regulators, and relevant national and international bodies. We have responded to the government's white paper on the regulation of AI to outline our approach to regulating AI (see [A pro-innovation approach to AI regulation - GOV.UK](#)) as well as the relevant actions from the AI opportunity plan (see [AI Opportunities Action Plan - GOV.UK](#)). We have also collaborated with the Canadian Nuclear Safety Commission and the United States Nuclear Regulatory Commission to develop joint principles for the regulation of AI. (See [Artificial Intelligence | Office for Nuclear Regulation](#)).

ONR and the Environment Agency have proactively engaged with the nuclear sector to develop a regulatory sandboxing process to support innovation. This provides regulators and industry the opportunity to work together to explore how suitable assurance can be provided for particularly novel innovations where there is limited experience of deployment to draw on to determine what represents good practice. It also delivers broader lessons and understanding for regulating innovation and provides important input for the development of regulatory frameworks. The sandboxing process was piloted on two AI use cases in 2023. (See [Outcomes of nuclear AI regulatory sandbox pilot published | Office for Nuclear Regulation](#)).

## Project overview

This report summarises the findings from the ONR-led project selected by the Regulatory Innovation Office as part of the AI Capability Fund ([Regulatory Innovation Office: One Year On](#)), which provided funding to regulators for them to engage with a series of novel and experimental AI projects.

Building on previous work in the nuclear industry to support the responsible deployment of AI in the nuclear industry, the aims of this project were:

- to accelerate the safe and secure deployment of AI in nuclear installations; and
- to support a responsible and gradual shift of focus from applications with low autonomy and safety significance to beneficial applications of higher autonomy and safety / environmental significance of failure.

The project was delivered through regulatory sandboxing, looking at two specific applications of AI associated with applications of interest for the UK nuclear industry. These both utilised a specific class of AI (i.e. supervised machine learning) to analyse, interpret and categorise computer vision data. This ONR-led project was delivered by working collaboratively with nuclear and non-nuclear regulators (Environment Agency, Health and Safety Executive, Defence Nuclear Safety Regulator) as well as several dutyholders across the nuclear sector (Rolls-Royce Submarines Ltd, NNB Generation Company (Hinkley Point C) Ltd, Sellafield Limited, Nuclear Restoration Services Limited). The regulators recognise that working together effectively is key to ensuring clear and consistent regulatory expectations across industrial sectors, to support proactive cross-sector learning and reduce the risk of regulatory fragmentation.

### **Project outcomes**

This project helped identify several technical findings, which are detailed in the report. Key areas of learning from the use cases were around technical skills and competences, AI assurance, and safety case. These project findings will be shared and discussed in detail with the nuclear industry to support advancement of future practical applications of AI.

Other project learnings were also identified in this report and will support future regulatory sandboxing initiatives. Some of these project learnings cover reflections on the effectiveness of the overall regulatory sandboxing approach, the value of identifying concrete use cases and the importance of converging on clear use case problem and opportunity statements.

The main outcomes from this project include:

- reducing project uncertainties on the selected use cases, with no fundamental reasons identified as to why the approach taken by the proposed concepts would be incompatible with the existing regulatory framework;
- demonstrating the value of cross industry discussions, with learning from the health and safety, medical and automotive sectors;
- supporting a change in mindset on AI applications within the nuclear sector, by exploring more ambitious options compared to the original use cases proposed by the industry; and
- confirming the suitability of the overall UK goal-setting regulatory framework to support AI, and innovation in general.

To ensure the outcomes and learning from the project are enduring, ONR, in collaboration with the Environment Agency and, as appropriate, other regulators, will address the following actions:

- share the technical findings with relevant external stakeholders to inform applications beyond the scope of the use cases considered in this project;
- disseminate the learning from this project within regulators and explore opportunities for further upskilling on AI to support the regulatory readiness ahead of future regulatory assessment or compliance work;
- consolidate the project learning, by finalising a handbook to help our organisations in future sandboxing initiatives, both from an administrative and technical perspective;
- explore opportunities for wider frameworks and regulatory innovation 'products' which could be beneficial for future engagement on innovation, including but not limited to AI; and
- undertake a review of the project outcomes with project partners (in approximately one year), to discuss any further learnings and discuss advancement with the use cases.

The actions will be tracked to ensure they are completed in a timely manner.

# 1. Introduction

The purpose of this report is to summarise the insights gained from the completion of the project 'Accelerating the safe and secure deployment of artificial intelligence in nuclear installations, focusing on use of AI to analyse, interpret and categorise data'. The project was funded by the Regulatory Innovation Office (RIO) in 2025/26 through the Artificial Intelligence Capability Fund (AICF)<sup>1</sup>.

This report presents the project background and delivery structure, before focusing on the main project learnings and outcomes and then presenting conclusions and next steps. Further relevant detail is provided in the Annexes.

## 2. Project background

### 2.1 Context

#### **Regulatory context**

The UK has a goal-setting framework for regulation of the nuclear sector, which inherently provides an environment where dutyholders can adopt innovative solutions and technologies. This includes systems containing AI, provided that dutyholders maintain the high standards of environmental protection, nuclear safety, security and safeguards expected.

In developing the safety case for an application, the dutyholder is expected to show that risks arising from the activities have been reduced to 'As Low As Reasonably Practicable' (ALARP)<sup>2</sup>. In making that case, they may consider all the relevant risks arising from operations through the whole plant lifetime, the whole range of plant operations and indeed the related risks covered by the safety cases of related plants or operation. The Department for Energy Security and Net Zero (DESNZ) has issued a Ways of Working document<sup>3</sup>, developed collaboratively by government, ONR, the Environment Agency and nuclear industry representatives. It establishes nine principles to guide the application of ALARP and 'Best Available Techniques' (BAT) – the fundamental concepts underpinning UK nuclear safety and environmental regulation.

As regulators, ONR and the Environment Agency have a duty to enable the organisations we regulate both to comply with the law and to be able to grow and

---

<sup>1</sup> [Regulatory Innovation Office: One Year On | GOV.UK](#)

<sup>2</sup> [ONR Technical Assessment Guide. The purpose, scope and content of safety cases \(NS-TAST-GD-05\) | Office for Nuclear Regulation](#)

<sup>3</sup> [Nuclear industry: principles to guide the application of As Low As Reasonably Practicable \(ALARP\) and Best Available Techniques \(BAT\) | GOV.UK](#)

thrive in their business. This includes pro-innovation behaviours that enable the implementation of innovative approaches and emerging technology, including AI.

The regulators are working together to realise this potential by increasing understanding of the technologies, including how AI could be used to support regulatory activities.

In November 2025 the Nuclear Regulatory Taskforce published its final report<sup>4</sup>, which made 47 recommendations for industry, regulators and government to implement with the aim of improving nuclear regulation in the UK. Recommendation 41 proposed that government works to “establish a nuclear digital programme to increase take-up of digital technologies (including AI), which will modernise approaches including on safety, regulation and engineering”<sup>5</sup>. The report also references the ONR and Environment Agency regulatory sandbox as a successful example, and more generally, emphasises the potential benefits of adopting AI across nuclear programmes, such as through safety assurance and in predictive maintenance. In March 2026, the government published its response to the Nuclear Regulatory Taskforce report, accepting all 47 recommendations. ONR and the Environment Agency welcome this and will continue to work with government, industry, and other regulators to support the implementation of the recommendations.

### **Regulatory Innovation Office (RIO) and AI Capability Fund (AICF)**

The UK government, via the RIO, is keen for regulators to encourage the responsible use of AI by industry to boost productivity in sectors where the UK already has strengths and for regulators to develop the capability to regulate AI effectively by learning alongside industry.

In this context the AICF<sup>6</sup> provided up to £3.6million of funding across regulators for them to engage with a series of novel and experimental AI projects.

### **ONR’s approach to innovation**

Regulatory sandboxing is one of the opportunities for dutyholders to engage with nuclear regulators on innovation. Figure 1 provides an overview of how ONR’s innovation vision links our overall strategic aims (being a modern, agile and productive regulator) and organisational values (accountable, open minded, fair and supportive) with the innovation outcomes we are seeking. Figure 1 also provides insights on options for targeted engagements through innovation ‘products’.

One of the key outcomes sought through our innovation vision is about sustaining the change in mindset on innovation. That is, ensuring industry and regulators are

---

<sup>4</sup> [Nuclear Regulatory Review 2025 | GOV.UK](#)

<sup>5</sup> [Building our nuclear nation: government response to the Nuclear Regulatory Review 2025 | GOV.UK](#)

<sup>6</sup> [Regulatory Innovation Office: One Year On | GOV.UK](#)

always risk-informed, open-minded and outcome-focused when considering new approaches. In this respect, the internal innovation ‘products’ are aimed at supporting our openness to innovative ideas as well as supporting trials for new approaches (e.g. use cases on responsible use of generative AI / large language model to assist with regulatory activities). External innovation ‘products’ offer the opportunity to create a space where different internal and external stakeholders can engage in a positive environment and collectively progress understanding the opportunities offered by innovation and address challenges.

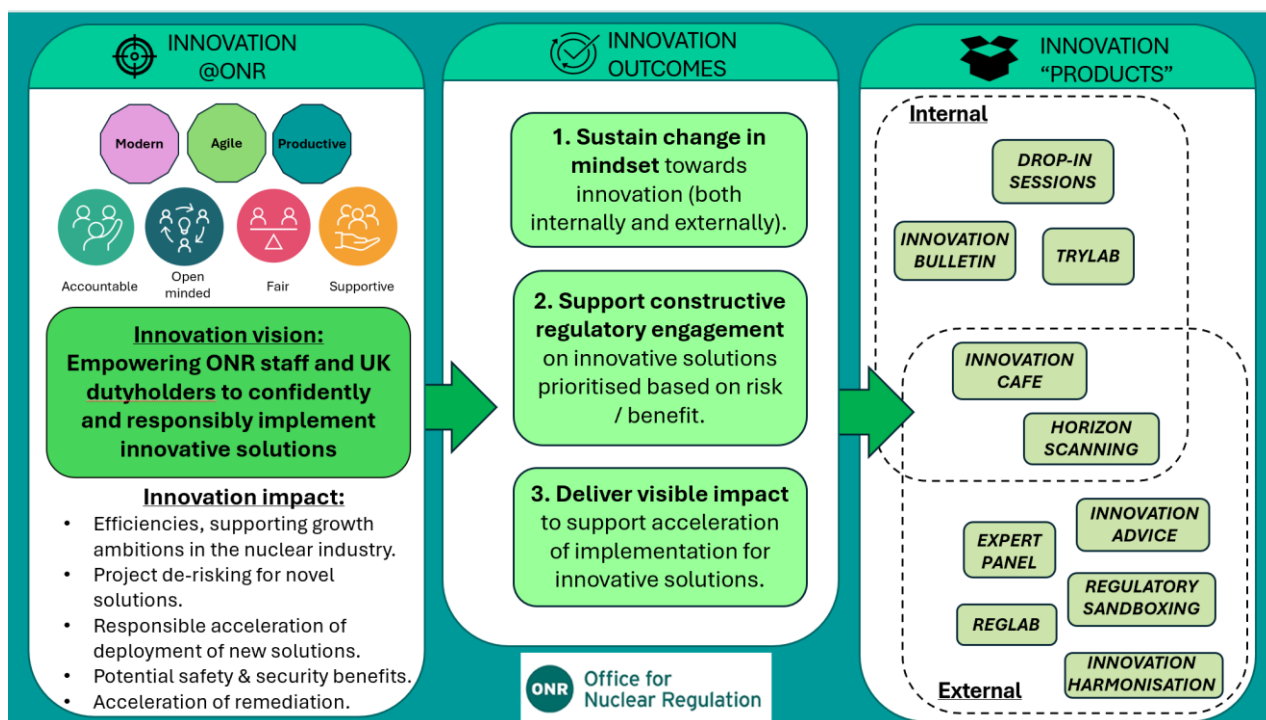


Figure 1: ONR’s innovation on a page <sup>7</sup>

Compared to other external facing innovation ‘products’, (e.g. innovation advice, expert panel, horizon scanning, innovation harmonisation), regulatory sandboxing focuses on the contextual information of a realistic specific application.

### Potential for deployment of AI in nuclear installations

Advancements in AI capabilities mean that new applications within international nuclear sectors that were previously not feasible are now being considered.

ONR and the Environment Agency recognise that the appropriate deployment of AI in nuclear installations may offer benefits, including:

<sup>7</sup> [Innovation | Office for Nuclear Regulation](#)

- the potential to improve nuclear safety, security, conventional health and safety, and environmental protection;
- the acceleration of the UK government's new build nuclear agenda; and
- the streamlining of hazard and risk reduction and decommissioning for legacy facilities through innovative approaches.

Whilst there is consensus both nationally and internationally on the potential application of AI in nuclear installations, currently there are only limited examples of deployment in operational applications<sup>8</sup>, particularly where AI is used to support safety-related functions.

It may be possible that the introduction of a new approach or technology may lead to short term increases in uncertainty. For example, in the case of AI, this may include different failure modes or a lack of recognised consensus on the appropriate methods of assurance. Any such uncertainty is expected to be managed within a demonstrably tolerable and ALARP risk envelope.

ONR and the Environment Agency expect the dutyholder to have adequate control of risks throughout the whole plant lifecycle. It may be appropriate to implement additional risk preventative or protective measures during the commissioning and initial operational phase until the understanding of the risks associated with the deployment and operation of the AI-based system is developed. New approaches and techniques may, in time, become good practices that regulators and the industry recognise as being relevant for achieving ALARP and BAT in certain circumstances.

## 2.2 Purpose of this project

The two main aims of this project were:

- to accelerate the safe and secure deployment of machine learning (ML) in nuclear installations - achieved through engagement between nuclear regulators, licensees and AI experts on examples of high-impact applications to support the government's nuclear industry ambitions; and
- to support a responsible and gradual shift of focus from applications with low autonomy and safety / environmental significance to beneficial applications of higher autonomy and safety / environmental significance of failure (see Figure 2 on the next page, shifting from bottom left corner towards the centre).

---

<sup>8</sup> [Considerations for Deploying AI Applications in the Nuclear Power Industry \(NR-T-1.26\) | IAEA](#)

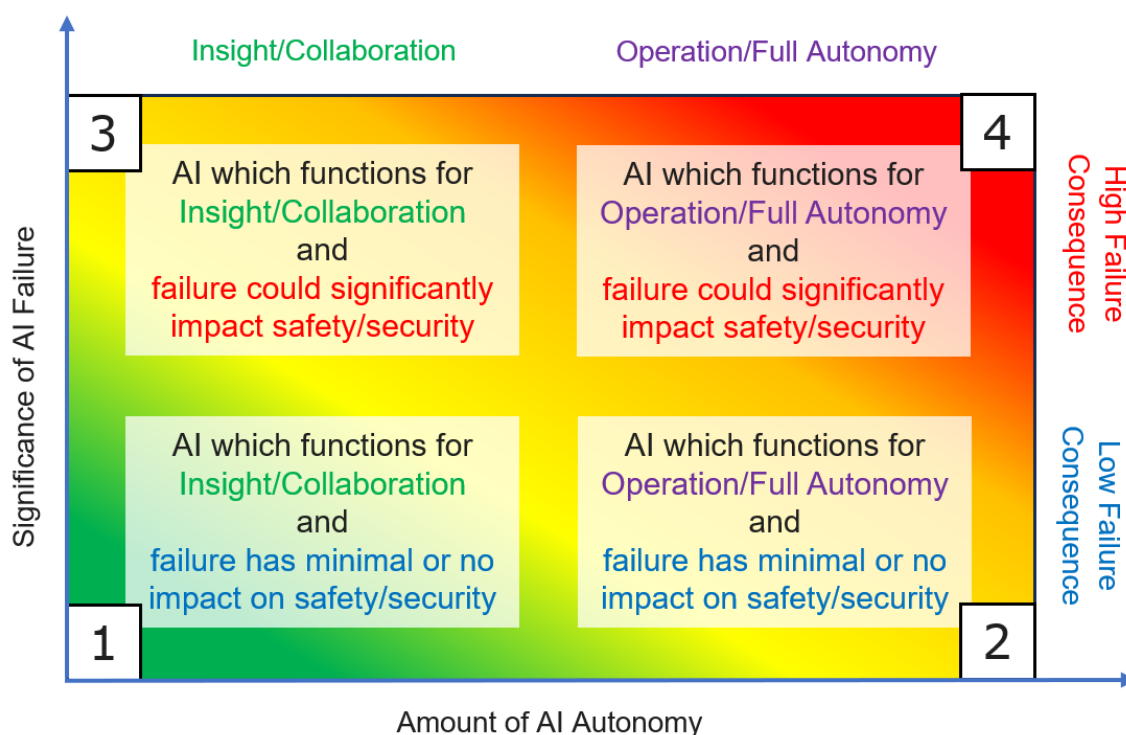


Figure 2: Four box model from ONR/US NRC/CNSC trilateral paper<sup>9</sup>

## 2.3 Links with previous work on AI at national level

This project recognises the learning from the previous regulatory sandboxing pilot on AI in the nuclear sector<sup>10, 11</sup>. Examples of the learning relevant to this project include:

- ML training featured strongly in the previous report stating: ‘AI systems can be very sensitive and the introduction of new data can have a big impact on the performance; for example, introducing data with different levels of noise can impact the output negatively. At the same time, natural changes in training data may make the system more robust. The use of multi-stream data could be beneficial as its output will be less susceptible to small changes in any one data stream.’
- Skills and experience were covered in the previous report stating: ‘Operational experience: end users of systems containing AI that could impact safety, security and environmental protection should have clear responsibilities. They should have the operational and application knowledge, including

<sup>9</sup> [New paper shares international principles for regulating AI in the nuclear sector | Office for Nuclear Regulation](#)

<sup>10</sup> [Regulators’ Pioneer Fund \(Department for Science, Innovation and Technology\): Pilot of a regulatory sandbox on artificial intelligence in the nuclear sector Outcomes of nuclear AI regulatory sandbox pilot published | Office for Nuclear Regulation](#)

<sup>11</sup> [Outcomes of nuclear AI regulatory sandbox pilot published | Office for Nuclear Regulation](#)

understanding the limits and conditions of operation, to ensure any inherent uncertainty in the AI systems leads to decisions that maintain conservatism’.

- Modification and future behaviour were covered in the previous report stating: 'change of state of system components, such as ageing of motors. Monitoring systems should be used to monitor any drift in behaviour of the AI system, to ensure it is still delivering its safety, security and environmental requirements and not inadvertently defeating safety and environmental protection measures. This needs to be captured in suitable and sufficient procedures.'

The main differences introduced during the new regulatory sandboxing were the level of detail in the applications provided by the industrial partners as part of the technical discussions. This enabled relevant aspects of the use cases to be explored to the detail required to draw out key findings in relation to AI assurance.

This work also considered other discussions and targeted work on AI. ONR and the Environment Agency are involved in various activities including:

- ONR targeted research on regulating AI and ML in the nuclear industry <sup>12</sup>;
- collaboration with other UK regulators through the AI and Innovation Subgroup of the UK Health and Safety Regulators Network; and
- learning about developing AI applications from industry, e.g. via fora such as the Nuclear Institute's AI4Nuclear Group and the Nuclear Decommissioning Authority (NDA) -sponsored Nuclear Waste and Decommissioning Research Forum.

## 2.4 Links with cooperation on AI at international level

This project also builds on the work that ONR and the Environment Agency have contributed to at an international level, aimed at building consensus and harmonising the approaches in relation to regulation and deployment of AI, including:

- pilot international regulatory sandboxing coordinated by the Electric Power Research Institute<sup>13</sup>; 9
- first RegLab<sup>14</sup> on AI organised by the Nuclear Energy Agency (NEA);

---

<sup>12</sup> [The Impact of AI/ML on Nuclear Regulation. Adelard | Office for Nuclear Regulation](#)

<sup>13</sup> [Program on Technology Innovation: International Nuclear Innovation Sandbox Pilot with a Regulatory Focus: "RegLab" | EPRI](#)

<sup>14</sup> [International RegLab Project | NEA](#)

- international engagement with other peer regulators, for example, ONR's work with the United States Nuclear Regulatory Commission (USNRC) and Canadian Nuclear Safety Commission (CNSC)<sup>15</sup>; and
- other work with international organisations (e.g. subgroup on AI within the International Atomic Energy Agency (IAEA)<sup>16</sup>, ONR's work with IAEA on a technical document on the safety and security implications of AI in nuclear).

## 2.5 Project scope

The term AI refers to a large category of applications, including machine learning, deep learning, generative AI and large language models. There are also different learning patterns, including supervised learning, unsupervised learning and reinforced learning.

The focus of this project was on a specific type of AI, i.e. ML with supervised training with no continuous learning. For this project, the use of ML can be conceptually described as two distinct stages:

- In the first, algorithms use labelled (i.e. supervised) data to learn the relationships that informed the assigned labels; and
- In the second, the trained model can then be applied to unlabelled data to generate classifications or predictions, consistent with the provided labels.

In addition, the use of AI in both test cases focused on delivering a single task (i.e. narrow ML). This facilitated the exploration of the factors that influence AI application and assurance whilst keeping a reasonable limit on system complexity.

## 2.6 Use cases

The ambition of this project was to start from applications of interest for the nuclear site licensees in the medium term. In selecting the applications, we wanted to make sure they provided learning across various sectors of the UK nuclear industry, from operating facilities (civil and defence) to new build and waste management. In this context, two use cases were selected:

- Use case 1. Use of AI in non-destructive testing (NDT), e.g. interpretation of weld radiographs; this is of interest to new build and operating facilities, including defence; and
- Use case 2. Use of AI for nuclear waste characterisation, including sorting and segregation of intermediate-level waste (ILW) items from low-level waste (LLW) items, as well as determining waste properties, e.g. sludge rheology and radiological property; this will support optimised waste management and

---

<sup>15</sup> [Considerations for Developing Artificial Intelligence Systems in Nuclear Applications. CNSC/ONR/USNRC | Office for Nuclear Regulation](#)

<sup>16</sup> [The International Network on Innovation to Support Operating Nuclear Power Plants | IAEA](#)

is of interest to the nuclear decommissioning and waste management facilities.

Industry discussions suggested these cases could improve safety and environmental performance. They could also speed up new build and deliver potential significant savings. Further detail of the use cases is presented in Annex 4.

The two use cases derive from applications being considered by four dutyholders operating in different parts of the nuclear sector.

## 3. Project delivery structure

### 3.1 Overall project plan

The project plan stretched over seven months from the initial government announcement of the successful projects for the AICF on 4 September 2025, to project completion on 31 March 2026.

Figure 3 provides an outline of the key activities involved. Section 3.2 provides some further details on the enabling activities which made the discussions in the regulatory sandboxing possible. Section 3.3 presents the core of the activities in the regulatory sandboxing, with the two workstreams. Activities planned after the closure of this project are presented in the conclusions in Section 6.

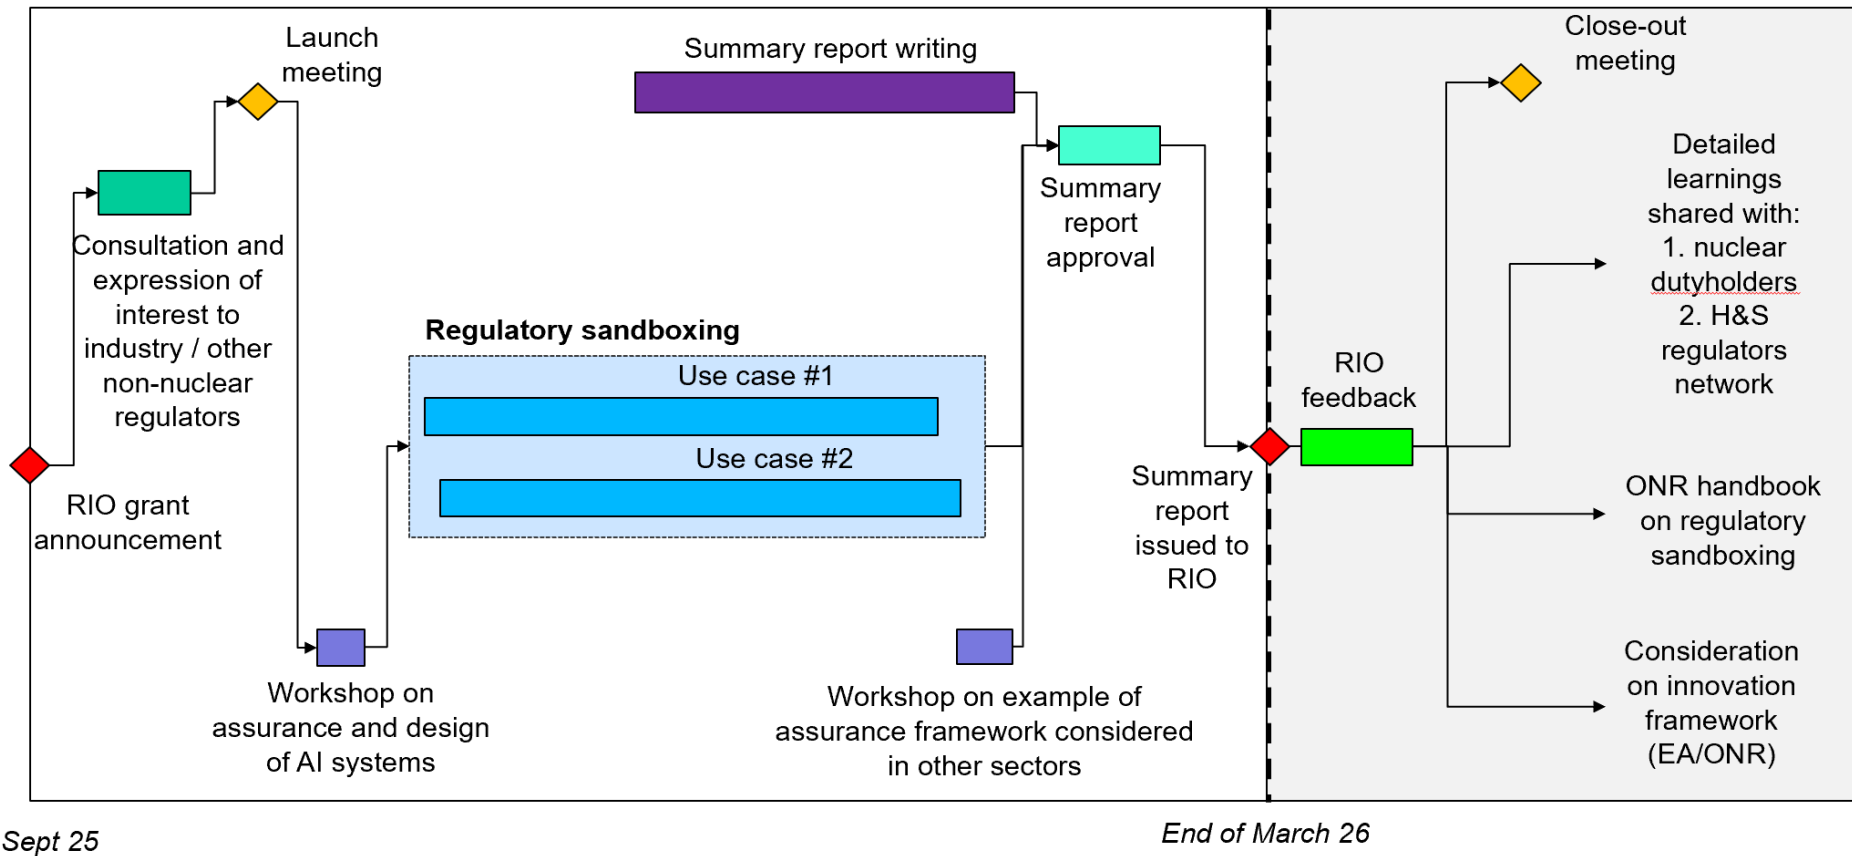


Figure 3: Key activities and project milestones

## 3.2 Enabling activities

Whilst the core of this project was the regulatory sandboxing activities, there were a set of other enabling activities essential for success, including:

- Consultation and expression of interest. Consultation activities with various external stakeholders identified key industrial partners and regulators who could contribute to the project in the given timescales. This phase stretched over three weeks before the launch event.
- Launch meeting. Following initial consultation, ONR and the Environment Agency hosted a formal launch meeting, publicised on social media, where the programme of work and the engagement expectations were set.
- AI design and assurance workshop. This enabling workshop helped to define the basis terminology and provide fundamental context to the follow up discussions, including identifying key references to be considered (see Annex 1).
- Example of AI assurance framework used in other industries. This event provided an opportunity to learn from the experience in ML assurance in other sectors, learning from an established framework developed for the automotive sector (Assurance of Machine Learning for use in Autonomous Systems - AMLAS<sup>17</sup>). The event was open to all project partners and the UK Health and Safety Regulator Network.

## 3.3 Regulatory sandboxing process

The core of the project was delivered by means of regulatory sandboxing. This provides regulators and dutyholders the opportunity to work together to explore how suitable assurance can be provided for particularly novel innovations where there are limited relevant good practices and experience of deployment to draw on.

The regulatory sandboxing process used in this project was adapted from the pilot approach<sup>18</sup> and informed by international practice<sup>19, 20</sup> alongside some of the specific features of the project use cases. Figure 4 presents the process for sandboxing of each use case.

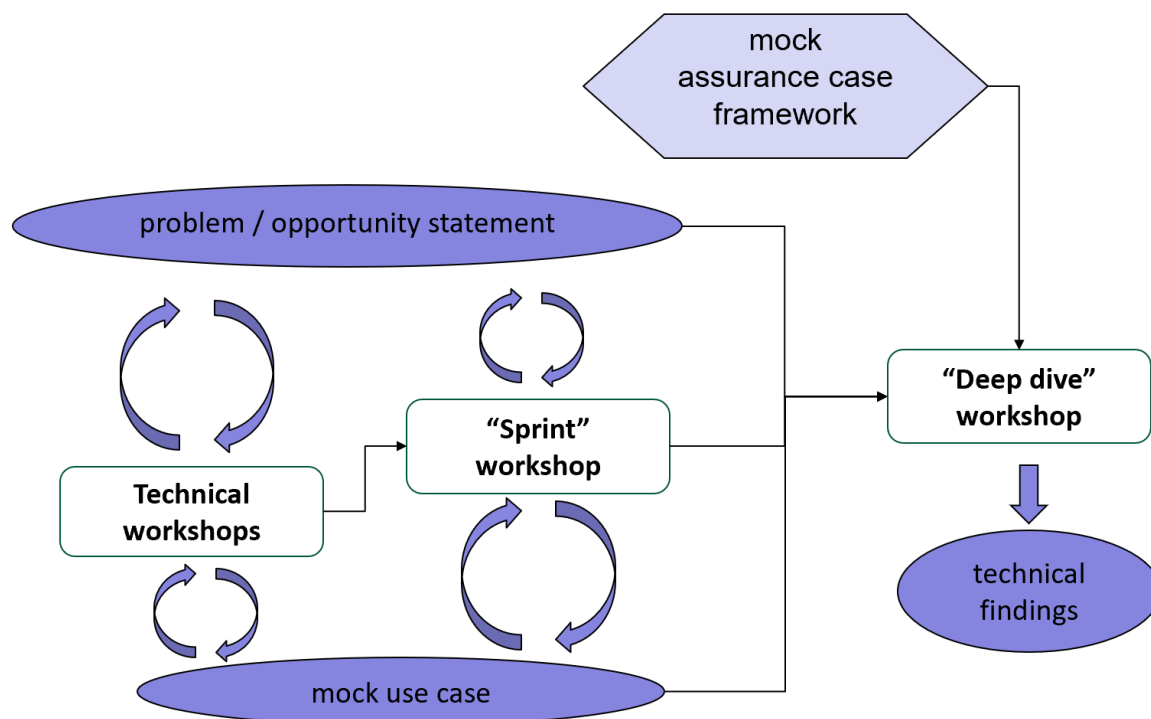
---

<sup>17</sup> [AMLAS - Centre for Assuring Autonomy | University of York](#)

<sup>18</sup> [Outcomes of nuclear AI regulatory sandbox pilot published | Office for Nuclear Regulation](#)

<sup>19</sup> [International RegLab Project | NEA](#)

<sup>20</sup> [Reglab Handbook | Ennuvo](#)



**Figure 4: Regulatory sandboxing process used during the project**

The key activities that formed the regulatory sandboxing were:

- Technical workshops. Two-hour online sessions for each use case to confirm partner ambition, outline key requirements and evidence, and identify areas for dutyholder-regulator discussion (including safety, security and environmental considerations).
- Sprint workshops. Half-day online sessions to refine the AI problem and opportunity statement and the draft assurance case, and to identify key opportunities and challenges. This informed the choice of deep-dive topics.
- Deep dive workshops. One-day face-to-face sessions for regulators and dutyholders to explore each AI application in more detail, develop key assurance-case fragments, and agree challenges and mitigations. Discussions also compared the AI-based approach with conventional solutions in the context of ALARP and BAT.

The outputs of the regulatory sandboxing include:

- problem and opportunity statement, which is the summary of the opportunity offered by the application selected as well as the key challenges to be addressed as part of the discussions;
- use case descriptions, which provide the contextual information about applications, helping participants having targeted discussions on critical issues;

- mock assurance cases, which provide an assurance case framework to act as the basis for the discussions and to explore the feasibility of developing assurance cases for the use cases under consideration; and
- the learnings from the discussions from the various workshops, including the identification of the evidence needed to support the key assurance case claims and whether it is likely to be available.

The process to converge on the problem and opportunity statement and on the use case (including the related mock assurance case) was iterative, reflecting the gradual learning process and incorporation of various perspectives.

The deep dive discussions were framed using a mock assurance case framework. This was derived from a classic claims, argument and evidence (CAE) structure used in nuclear safety cases, to set out a key fragment of the relevant assurance case. The key claims used in this framework were:

- the requirements are sufficiently defined for the system;
- the system meets the requirements initially; and
- the system meets the requirements in the future.

## 4. Technical findings and project learning

This section highlights the main technical findings and project learning from the regulatory sandboxing, consolidated across both use cases through discussions with the project partners.

### 4.1 Technical findings

In line with the project aim to accelerate the safe use of AI in the nuclear sector, the findings in this report focus on those technical findings that have wider relevance for dutyholders considering the use of AI.

Based on the limited information available in the regulatory sandboxing and within the bounds set in the report disclaimer, no fundamental reasons were identified as to why the approach taken by the proposed concepts would be incompatible with the existing regulatory framework. However, whether the use cases discussed can demonstrably meet safety and environmental requirements will depend on how the detailed design and execution of each use case is developed, including the context within which the use case is applied and the degree of assurance that is incorporated into each system.

Some technical points emerged from the discussions, the consideration of which will help the development of these technologies and potential areas of regulatory focus. These technical findings are presented in Annex 2 and are grouped into the following high-level themes:

- technical competency and organisational capability;
- supply chain management;
- assurance of machine learning;
- safety, security and environment case;
- interface between human and AI; and
- AI application context.

### 4.2 Project lessons learnt

The main lessons learnt fall into the following areas:

- Effectiveness of regulatory sandboxing. There was general consensus that sandboxing provides a constructive space for regulators and dutyholders to collaborate on innovative solutions. In this project, it helped stakeholders converge on those high-level technical findings from the two use cases that are relevant to other applications.
- Value in specific details of the use cases. Our experience from this project is that there is value in having a specific context of applications (including the

concept of operation, operational philosophy, workflow, and the role of the operator) to guide the discussion as part of the workshops and achieve concrete outcomes. Having real examples from the industrial partners was a key enabler to deliver concrete outcomes.

- Value in breadth of use cases. The sandboxing showed the value of covering different use cases (two use cases from four applications across multiple dutyholders and in different parts of the nuclear sector), with varied safety claims, roles and operating philosophies. Exploring more safety related and environmental related AI roles helped clarify what is, and is not, likely to attract regulatory attention.
- Multi-disciplinary effort. Regulatory sandboxing relies on input from multiple stakeholders and disciplines (e.g. safety, security, environmental protection, engineering, safety case, human factors, chemistry, nuclear liability, pre-disposal management and BAT). It also requires coordination across non-regulatory support functions, including communications, finance, commercial/contracting, project management and administration.
- Challenges associated with the timescales. The short window of opportunity required effective management of risks to ensure timely delivery. The short timescale had some impact on the depth and breadth of the project, e.g. in relation to:
  - considering in detail some other relevant examples from the UK nuclear industry;
  - widening the opportunity for the supply chain and other experts to join the discussions; and
  - addressing technical topics beyond those selected for the discussions in the deep dive workshops.

However, none of the above had a fundamental impact on the positive outcomes of the project.

- Effort associated with regulatory sandboxing. The coordination of multiple stakeholders across the industry and other regulators, as well as the technical input requires a dedicated funding route. ONR, in collaboration with the Environment Agency, will consider options going forward to enable the delivery of similar activities (including alternative options and innovation 'products' which can achieve similar outcomes, see Annex 3).

## 5. Project outcomes

Following on from the main technical findings and project learning we identified the following impacts derived from this project:

- UK regulatory framework and innovation. This project confirmed the advantage of the UK's goal-setting, outcome-focused framework. It can better accommodate emerging technologies such as AI than rules-based regimes, where legislation, standards and guidance often need time to mature.
- Cross-sector cooperation. Input from non-nuclear regulators and partners from other sectors (e.g. automotive and medical) helped build on good practice and reduce the risk of regulatory fragmentation. It also supported a more holistic view of risk, combining insights from nuclear operators with relevant lessons from industries with different hazard profiles.
- Reduction of use case uncertainty. Based on the feedback from project partners, the discussions on both use cases were able to reduce project uncertainty associated with the use cases considered. Whilst the timescale for the implementation depends on factors beyond the remit of regulators, there was consensus on the value of this initiative in supporting acceleration.
- Increased regulatory AI knowledge and skills. The regulators involved provided positive feedback on the opportunity this project offered to increase their understanding of key fundamental principles on AI, as well as gaining confidence in the dutyholders approaches to these technologies. This learning will be shared across their organisations to provide additional awareness and understanding to support future regulatory work.
- New opportunities offered by AI. One use case showed how AI could provide new insights and enable alternative operating concepts, for example by combining multiple sensors to inform operation. Further exploration of similarly transformational applications could unlock wider safety, security and environmental benefits.
- Opportunity for a new regulatory engagement 'product'. This project outlined the effectiveness of the regulatory sandboxing approach to support targeted discussions on specific innovative applications. These activities require significant involvement of both regulators and dutyholders, which often requires external funding support. Annex 3 presents some initial considerations on alternative options which could help in certain scenarios, and will be considered after the closure of this project.
- Innovation mindset. Project partners reported that the sandboxing helped sustain a more innovation-positive mindset, with industry and regulators remaining risk-informed, open-minded and outcome-focused. This was enabled through the sandboxing activities, which provided the space and time to support open, proactive engagement and the exchange of points of view between regulators, dutyholders and AI experts. This mindset also enabled

the exploration of more ambitious options compared to the original use cases proposed by the industry.

Annex 5 presents additional details of how the project delivered impact and benefit.

## 6. Conclusions and next steps

The project was successful in supporting the acceleration of responsible deployment of artificial intelligence in the nuclear industry. By enabling an open discussion between regulators, dutyholders and AI experts, ONR and its project partners managed to efficiently increase their understanding of the opportunities offered by a class of AI (i.e. machine learning with offline supervised training) that is of interest to the UK nuclear industry (i.e. in applications using computer vision for data analysis and classification).

The interaction with AI experts, the nuclear industry and regulatory partners in a regulatory sandbox was helpful to increase the overall understanding of how AI can safely and securely support efficiencies in nuclear, and to identify areas to address ahead of deployment. Based on the limited information available in the regulatory sandboxing and within the bounds set in the disclaimer, no fundamental reasons were identified as to why the approach taken by the proposed concepts would be incompatible with the existing regulatory framework. However, whether these use cases can demonstrably meet safety and environmental requirements will depend on how the detailed design and execution of each use case is developed, including the context within which the use case is applied and the degree of assurance that is incorporated into each system.

This regulator-led project, delivered over seven months, shows how we can work with multiple external stakeholders, including ten project partners, in a short time-period to help support innovation. We believe this demonstrates our intent to be modern, agile and productive regulators.

We expect our impact to go beyond the two test cases considered in this project, for example by:

- unlocking further discussions on innovations with UK nuclear site licensees, through the learning from this project and the openness shown by the nuclear regulators involved;
- strengthening cooperation with other regulators (including beyond the nuclear sector), reducing the risk of regulatory fragmentation and fostering learning across sectors; and
- taking advantage of the value added by drawing on AI expertise and experience from other high-hazard industries early in the development and implementation of AI-based systems.

To ensure the outcomes and learning from the project are enduring, ONR, in collaboration with the Environment Agency and as appropriate other regulators, will address the following actions:

- share the technical findings with relevant external stakeholders in the UK nuclear industry, to inform applications beyond the scope of the use cases considered in this project [**Owner: ONR and Environment Agency. Activies include: Close out event open to public in Quarter 2 2026; briefing at the workshop with Heads of Innovation across the nuclear industry in Quarter 3 2026; sharing of the main findings with the UK Health&Safety Regulators Network by Quarter 3 2026. Target completion date: Quarter 3 2026**];
- disseminate the learning from this project within regulators and explore opportunities for further upskilling on AI to support the regulatory readiness ahead of future regulatory assessment or compliance work [**Owner: ONR and Environment Agency. Activies include: Innovation Cafés and Innovation Bulletin. Target completion date: Quarter 2 2026**];
- consolidate project learning, by finalising a handbook to help our organisations in future sandboxing initiatives, both from an administrative, project and technical perspective [**Owner: ONR. Activies include: Internal review, learn and improve from this project and development of regulatory sandboxing handbook. Target completion date: Quarter 4 2026**];
- explore opportunities for wider frameworks and regulatory innovation ‘products’ which could be beneficial for future engagement on innovation, including but not limited to AI (see Annex 3) [**Owner: ONR and Environment Agency. Activies include: Reviewing the options presented in Annex 3 and progressing resolution. Target completion date: Quarter 4 2026**]; and
- undertake a review of the project outcomes with project partners (in approximately one year), to discuss any further learnings and discuss advancement with the use cases [**Owner: ONR and RIO. Activies include: One year on project review. Target completion date: Quarter 1 2027**].

## 7. List of abbreviations

AI	Artificial intelligence
ALARP	As low as reasonably practicable
AICF	Artificial Intelligence Capability Fund
BAT	Best available techniques <sup>21</sup>
CAE	Claims, arguments and evidence
CNSC	Canadian Nuclear Safety Commission
IAEA	International Atomic Energy Agency
IEC	International Electrotechnical Commission
ILW	Intermediate Level Waste
LLW	Low Level Waste
ML	Machine learning
NEA	Nuclear Energy Agency
NDA	Nuclear Decommissioning Authority
NDT	Non-destructive testing
ONR	Office for Nuclear Regulation
RIO	Regulatory Innovation Office
RSR	Radioactive Substances Regulation
USNRC	United States Nuclear Regulatory Commission
UK	United Kingdom of Great Britain and Northern Ireland

---

<sup>21</sup> The “best available techniques” (BAT) is defined as the latest stage of development (state of the art), with special consideration given to: a) comparable processes, facilities or methods of operation which have recently been successfully tried out; b) technological advances and changes in scientific knowledge and understanding; c) the economic feasibility of such techniques; d) time limits for installation in both new and existing plants; e) the nature and volume of the discharges and emissions concerned. It therefore follows that what is BAT will change with time in the light of technological advances, economic and social factors, as well as changes in scientific knowledge and understanding.

## Annex 1 – Reference documents

A key purpose of a regulatory body such as ONR is to ensure that dutyholders manage and control risks effectively, to the standards required by law.

Some law is prescriptive – spelling out in detail what must be done. However, UK legislation on safety, security and environmental protection is largely goal-setting – setting out what the dutyholder must achieve, but not how it must be done. When judging goal-setting law, decisions are often guided by those good practices that would normally be expected to meet the standards required.

Good practice in the deployment of AI in the nuclear sector is emerging and is yet to be validated by extensive use and regulatory scrutiny. However, several key documents and frameworks are already shaping the form of good practice in the development, design, operation, maintenance and decommissioning of such systems. A number of these were drawn on during the sandboxing activities and are summarised here as the project stakeholders found their content informative:

- IAEA document ‘Considerations for Deploying Artificial Intelligence Applications in the Nuclear Power Industry’ (NR T 1.26)<sup>22</sup>. This document provides a broad introduction to the current use of artificial intelligence in nuclear applications, covering some key aspects of the AI lifecycle management, providing considerations on use of data and identifying further high level considerations (e.g. explainability and the human interface). A follow up IAEA TECDOC covering safety and security implications of use of AI in nuclear is under preparation at the time of writing this report.
- Trilateral principles paper co-authored by the Canadian Nuclear Safety Commission (CNSC), the Office for Nuclear Regulation (ONR) and the US Nuclear Regulatory Commission (NRC). ‘Considerations for Developing Artificial Intelligence Systems in Nuclear Applications’<sup>23</sup>. This document outlines high-level principles for the deployment of AI as noted by the three regulators. It focuses on the use of existing relevant engineering principles, human and organisational factors, system architecture, lifecycle management and documenting the AI assurance case.
- The Institution of Engineering and Technology (IET), ‘The Application of Artificial Intelligence in Functional Safety’<sup>24</sup>. This document focuses on ten key ‘pillars’ that highlight the risks of using AI in functional safety systems. It outlines the additional considerations required in the engineering processes and provides guidance on building assurance cases for AI in safety related

---

<sup>22</sup> [Considerations for Deploying Artificial Intelligence Applications in the Nuclear Power Industry \(NR T 1.26\) | IAEA](#)

<sup>23</sup> [Considerations for Developing Artificial Intelligence Systems in Nuclear Applications. Sept 2024. CNSC/ONR/USNRC | Office for Nuclear Regulation](#)

<sup>24</sup> [The Application of Artificial Intelligence in Functional Safety | IET](#)

applications and the differences between conventional and AI algorithms. The document highlights the importance of hazard analysis and risk assessment, specification and architecture, selection of appropriate data, and system verification, validation and assurance. It draws heavily on existing good practice guidance contained in IEC 61508 Parts 1 - 7. 'Functional safety of electrical/electronic/programmable electronic safety-related systems'.

- Environment Agency research paper: Radioactive Substances Regulation (RSR) Digital and Robotic Innovation Review, May 2024. The Environment Agency commissioned a research report to gather examples of digital and robotic innovation relevant to RSR and to consider the applicability of RSR environmental permit conditions and principles. It found evidence of an increasing number of digital and robotic innovations (including AI) relevant to nuclear regulation; no significant issues with permit conditions were identified, RSR regulatory principles were found to be largely non-prescriptive and can generally accommodate digital and robotic innovation (including AI), but may require some minor amendments for clarity.
- ONR research paper 'Impact of AI/ML on nuclear regulation', June 2021<sup>25</sup>. This paper reports on a project to advise ONR on the suitability of existing UK nuclear regulation to the application of AI and ML, and to provide a route map towards gaining the benefits of these technologies in UK nuclear facilities. It presents a good overview of AI systems and assurance requirements. It finds that ONR's Safety Assessment Principles are themselves strong, however the associated guidance will require interpretation and the development of additional rationale in order to be meaningfully applied to AI/ML systems.
- ONR research paper on AI/ML, April 2026<sup>26</sup>. ONR commissioned a second research report into the application of AI and ML to the UK nuclear industry. It found that AI has great potential to improve the safety and efficiency of current and future nuclear plants and facilities. However, where there are significant consequences arising from failure of AI, ongoing research has to date been unable to identify suitable ways to assure safety and security. There are also significant challenges where workers interact with systems containing AI. Despite these challenges, the report found that if used in conjunction with non AI technologies, it is possible to successfully deploy systems containing AI in a wide range of applications in the nuclear industry. In addition, ONR is now seeing indications that the use of suitable safety and architectural design approaches may enable the demonstration of an adequate level of safety for systems containing AI in some applications.

---

<sup>25</sup> [The Impact of AI /ML on Regulation. Adelard. 2021 | Office for Nuclear Regulation](#)

<sup>26</sup> [ONR Research reports: Route map for regulating and deploying AI/ML systems within the UK civil nuclear industry. Characterisation of safety of AI technologies and AI framework | Office for Nuclear Regulation](#)

## Annex 2 – Technical Findings

The following technical points emerged from the regulatory sandboxing discussions and have been grouped into the following high-level themes:

### **Technical competency and organisational capability**

- Competency and organisational capability. Competency and organisational capability form the human foundation of trustworthy AI deployment, ensuring that every actor - from designer to operator to regulator - remains capable of understanding, challenging, and assuring AI performance throughout the system's lifetime. The discussions in the project emphasised the need to combine specific specialist AI knowledge with in-depth understanding of the use case. This applies equally to dutyholders and regulators, as a combination of AI and application specific knowledge is needed to understand how the performance of the AI component could lead to undesirable safety or environmental consequences and how to manage the AI component's performance within the whole system to avoid this. A combination of knowledge in various areas is essential for the correct formulation of a project specification for the use of AI. Building organisational capability means embedding AI expertise within management structures, decision-making processes, quality systems and training programmes.
- Operator awareness. For operators, competence extends beyond domain-specific skills. It would be beneficial for them to understand the fundamentals of how the AI system functions, its intended scope, its failure modes and limitations. This understanding supports correct use, interpretation of outputs (including their validation where needed) and recognition of anomalous behaviours to support situational awareness.
- Long term skill retention of stakeholders in sector. Regulators and industry acknowledged that the safety, security and environmental cases for AI applications will need to adequately consider how the increased reliance on AI could impact the skills, situational awareness and safety culture of the staff involved, both in the near-term and long-term. This is particularly important in applications where: i) humans are expected to retain responsibility for making the final decision regarding use of the output recommended by the AI system, and ii) the unavailability of an AI system may impact safety (i.e. human intervention is necessary to avoid an undesirable safety or environmental consequence).

## Supply chain management

- Intelligent customer.<sup>27</sup> Dutyholders may seek AI solutions provided by contractors. Therefore, the dutyholders will need to make informed decisions when using the products and services provided by contractors, acting as an intelligent customer, requiring a skillset or access to specialist knowledge which is currently limited within the nuclear sector. The licensees and other dutyholders in the UK nuclear industry are ultimately responsible for the safety, security and environmental protection of their installations and should develop internal capability to assess systems containing AI / ML solutions. They should also understand their failure modes and evaluate whether these are adequate solutions that meet their own operational and regulatory requirements. This will require a potential new set of skills, knowledge and experience commensurate to the reliance on safety and environmental protection placed on the AI component within the overall system.
- Liability and accountability. One of the challenges in acceleration of adoption of AI in nuclear may be the dutyholder's contractual relationship with the specialist supply chain. Safe and secure operation of the system depends on the combination of the application-specific data, the system development and design, and verification and other assurance techniques implemented. Intellectual property issues may arise with third-party vendors supplying AI components or training datasets when they are encouraged to provide adequate evidence of compliance with requirements and transparent performance metrics to support integration and assurance claims of the wider system. In addition, long-term assurance depends on how modifications and changes, either intentional or unintentional, are managed. A structured change-management approach is essential, which depends on robust engineering governance and supply chain management. Therefore, supply chain arrangements need to support secure and traceable update procedures that may be required if changes alter the system behaviour in ways that could invalidate earlier assurance claims. As contained in ONR's pro-innovation approach to AI regulation<sup>28</sup>, we expect that effective oversight is facilitated through structured, integrated, and diverse means such as self-assessments at facility and department level, internal independent oversight, sophisticated governance structures, external assessments, and peer reviews. Dutyholders may wish to discharge some governance responsibility to supply chain organisations; however, legal accountability remains with the dutyholder. Consideration needs to be given to these aspects at early stages of the supply chain engagement.

---

<sup>27</sup> [ONR Technical Assessment Guide. Licensee Core Safety and Intelligent Customer Capabilities \(TAST-GD-049\) | Office for Nuclear Regulation](#)

<sup>28</sup> [ONR's pro-innovation approach to AI regulation | Office for Nuclear Regulation](#)

- **Security.** Security of the AI supply chain should be treated as a continuous 'security-by-design' obligation rather than a one-time check. AI systems are uniquely susceptible to adversarial attacks, and it is therefore critical to maintain rigorous provenance for all third-party data, models and software dependancies. The feasibility of achieving and demonstrating security is contingent upon implementing strict supplier due diligence.

### **Assurance of machine learning**

- **Assurance.** Assurance of AI components requires a structured approach that considers the suitability of the AI system, the adequacy of the data, the robustness of testing and verification evidence, and the provision of an AI platform and good quality supporting tools. As part of the project, we looked at assurance approaches developed in the medical industries<sup>29</sup>; in the near-term, this is likely to require each application to define its own evidence framework proportionate to the system role and potential risk contribution.
- **Narrow ML.** The focus on a particular type of AI, specifically ML that performs a single task based on supervised training data (referred to as narrow ML), helped to clarify the difference in the challenge compared to wider categories of AI (e.g. continuous learning, reinforced learning, generative AI). Addressing some of the challenges more generally associated with AI (e.g. validation) can be less onerous for narrow ML. For example, when a narrow ML trained in a supervised way is used to identify a defect in a radiography, it is easier for a human to validate their results compared to using a generative AI for multiple purposes.
- **Training data.** In some applications (e.g. using AI to detect defects in NDT) there may be scarcity of data of a certain category for training the ML on relevant examples. In NDT, this is due to the high-quality manufacturing processes and the low number of defects. This should be considered as part of the initial definition of the system, to identify whether use of data from other applications or synthetic data is appropriate.
- **Monitoring of AI application during operation.** Whilst this is an important aspect for any component, for AI this becomes even more crucial because of the strong dependence on data. Both intentional changes (e.g. planned upgrades or maintenance) and unintentional changes (e.g. shifts in data distribution) can alter system behaviour in ways that may invalidate earlier assurance claims. Defining the scope of revalidation is critical since not every change requires complete revalidation. A proportionate impact assessment framework should be defined to determine whether modifications have impact (e.g. on system inputs, data quality, model behaviour, or downstream safety functions).

---

<sup>29</sup> [AMLAS assurance methodology - Centre for Assuring Autonomy | University of York](#)

## Safety, security and environment case

- Role of AI application in safety cases and environmental cases. The discussions around the two use cases emphasised the importance of holistically balancing the benefits of the AI system with the potential risk and consequences associated with the case specific application. In determining the role of the AI component, dutyholders should consider the overarching concept of operation (e.g. the role of humans in the supervision of the AI task), requirements on the AI systems (both functional and performance requirements), the claims in the safety / environment case (e.g. existence or not of further layer of defence in depth). In some of the use cases in this regulatory sandboxing, ML did not play a significant role in the overall concept of operation, and therefore the consequences of their potential failures were partially mitigated by other steps in the workflow, reducing the safety reliance on the AI component.
- ALARP and BAT. Established safety and risk assessment methods can still be used for systems with AI elements to control relevant risks, so they are reduced as ALARP<sup>30</sup>, using BAT<sup>31</sup>. These risk assessments should consider both the benefits and risks of using AI compared with other options, not just at an isolated point in time of operation of the system, but the overall risk landscape arising from the use of the system over its lifetime. Safety and environmental submissions for innovative solutions should seek to capture strategic benefits from deployment of the innovative solution to support the lifecycle approach to risk management. Often there is a need for the wider system to include protection if the AI fails. However, AI can fail in less predictable ways than traditional software, and its behaviour cannot always be fully defined in advance. When system requirements are incomplete, risk assessment lacks a clear reference point. This makes it difficult to identify failures and to balance risks appropriately, particularly for AI systems, where defining requirements and assessing risk are already more complex. Risk assessments should also compare AI systems with alternatives to show that risks are tolerable and proportionate. Because AI is relatively new, there is often limited real world data to support these comparisons. Dutyholders should therefore recognise that AI may introduce new risks, such as incorrect decisions leading to waste being placed in the wrong category.

## Interface between human and AI

- Human and AI co-working. The use of AI changes how people and systems work together. To use AI safely and effectively, clear arrangements are needed for how humans and AI share tasks and decisions. Human involvement should be built into the system's design. People and AI should have clearly defined roles and decision limits. Systems should give users

---

<sup>30</sup> [ONR Technical Assessment Guide. Regulating duties to reduce risks to ALARP \(NS-TAST-GD-005\) | Office for Nuclear Regulation](#)

<sup>31</sup> [Radioactive substances regulation \(RSR\): objective and principles | GOV.UK](#)

clear information about what the AI is doing, its limitations, and when human intervention is required. The right level of trust is essential. Over-trust can lead to errors, while under-trust can reduce the benefits of automation. Trust is supported by transparent behaviour, clear explanations, and visible levels of uncertainty. Human oversight also brings challenges, such as maintaining attention, avoiding fatigue, and managing handovers in urgent situations. Users therefore need appropriate training to understand AI outputs, recognise problems, and judge when AI may be unreliable. Ongoing review and training is needed to keep the system safe and effective over time.

- Different role of human subject matter experts. Rather than replacing human specialist knowledge, AI-based systems could support the human in performing the required operation or function. The role of the subject matter expert will continue to be crucial in ensuring the adequate behaviour of the system, including, for example, when defining the AI training set, validating the AI-supported system, monitoring the system during operation, and manually reviewing and overriding outputs of the AI-based system. Finally, the AI application in real operation has to be supported by an appropriate and proportionate demonstration that safety and environmental requirements have been met. This will continue needing input from suitably qualified and experienced persons. Depending on the specifics of the application, input from subject matter experts with deep expertise may be required.

## AI application context

- False positive vs false negative. There is a balance to be made between false positives and false negatives (or ‘operating point’) which needs to be considered when defining the requirement specification for the system containing AI. For example, a false positive for the NDT use cases would be the system identifying a defect that does not exist, whereas a false negative is the system not identifying a defect. This can create a trade-off between efficiency and safety. False positives can generate extra re-work, and false negatives introduce safety risks (missed defects) which need to be considered and mitigated, e.g. by alternative means. Both will require additional effort (either re-work or increased supervision). Depending on the context, consequence and risk, this balance needs to be carefully considered.
- Transferring AI application to a different use. As part of the project, we discussed using an AI solution to deliver a similar function across different applications (different context with a similar safety claim or performance objective). The transfer to a new application will require a re-assessment of the system’s assurance case and suitability to support the new safety and operational claim. This is because there can be a number of project-specific factors that may play a key role in the AI system performance (including the relevance of the training set for the application). From the discussions in the use cases, whilst the AI system would need to be re-validated, general learning from an application could be transferred to other applications expediting their implementation. The transfer of learning may be limited if the

level of dependence on AI in decision making is greater than in previous applications. For example, if a new waste characterisation application is developed, the extent of reliance on AI in decision making will need to be defined and a commensurate level of controls applied to manage the risks associated with that new application.

- Learning from early operations and considering design flexibility. There is an opportunity to use active demonstrations, trials, or early operational phases of systems that use AI to gather insight and guide future development, provided robust controls are in place to ensure they can 'fail safely'. For example, during autonomous sorting or segregation activities, operators may be able to assess the rate of false positives and false negatives generated by the AI component. This evidence can then inform what proportionate engineered and procedural controls should look like in the operational phase to manage risks such as misconsignment. Where there is significant initial uncertainty about the reliability of the AI component, it will generally be prudent for operators to adopt more stringent controls. As operational learning increases, some controls may be safely reduced or removed where they can be shown to be unnecessarily onerous. To make full use of this potential benefit, dutyholders may need to consider from the outset whether their systems containing AI or facilities can be designed with sufficient flexibility to accommodate future adjustments as understanding improves.

## Annex 3 – Evolution of innovation engagement

The project highlighted the value of regulatory sandboxing to support innovation, and ONR will continue working on opportunities for similar initiatives in the future.

To date, funding for regulatory sandboxing has been granted from government departments through competitive processes (Regulators' Pioneer Fund and the AICF). We have recognised that this method of funding can restrict the opportunity to support future sandboxing in the timescale needed by the UK nuclear sector and therefore are considering options to resolve this, including:

- Option 1 - New funding mechanisms for future regulatory sandboxing. Rather than from government, the funding stream could come from industry. For example, this could be facilitated through the head of innovation forum, based on agreed priorities across the industry;
- Option 2 - New innovation 'product'. There is an opportunity to consider a new innovation 'product' that is more detailed than an expert panel and innovation advice, but less onerous than regulatory sandboxing. The focus of this product could be on reducing uncertainty on an innovative solution by stress testing it at early stages of its development, to identify possible fundamental shortfalls. This could involve a more restricted amount of external stakeholders; and
- Option 3 - Industry-led sandboxing. The industry could set up sandboxing initiatives to address problem statements which helps their challenges and timescales, as opposed to participating in regulatory sandboxing driven by ONR and the Environment Agency. The regulators could be invited to contribute to specific points of interest. The output and the impact of the learnings may be different compared to regulatory sandboxing, because of the difference in the setup and objectives.

## Annex 4 – Summary of project AI use cases

**Table 1: Project AI use cases**

Item	Use Case 1	Use Case 2
Application	Non-destructive testing	Waste Characterisation and management
Dutyholders	Rolls Royce Submarines Limited and NNB Generation Company (HPC) Limited	Sellafield Limited and Nuclear Restoration Services Limited
Regulators	ONR, Health and Safety Executive and Defence Nuclear Safety Regulator	ONR and Environment Agency
Type of AI	Supervised machine learning	Supervised machine learning
Problem statement	<p>Radiographic assessment of welds is important to ensure no leakage, making it necessary to detect defects such as pores, wormholes and lack of fusion. This is currently done by a human operator using radiography.</p> <p>The main hazard is a leakage of radioactive material, which can lead (directly / indirectly) to a nuclear safety hazard.</p> <p>The focus of this work will be on the case of using AI to carry out a first radiographic assessment to identify and classify into ‘cleared’ and ‘with unacceptable defect’ radiographs. A suitably qualified and experienced radiographer would then proceed to review all films that have actionable defects and only periodically review the cleared</p>	<p>Current practices for nuclear waste characterisation use a variety of approaches to obtain information about the properties of radioactive waste that underpin decisions made on its subsequent management by people.</p> <p>In some instances, characterisation involves challenging environments and can be subject to human error. Some characterisation tasks are repetitive; others may incur operator dose. A variety of other factors can also feed into uncertainty around waste characteristics.</p> <p>The focus of this work is on the case of using AI for nuclear waste characterisation that improves some or all the topics described above. It includes sorting and segregation of intermediate-level waste (ILW) items from low-level waste (LLW) items or non-radioactive waste, as well as determining waste properties, e.g. sludge rheology and radiological property.</p>

Item	Use Case 1	Use Case 2
	<p>radiographs. This could potentially be done in a way where robotics automate concurrent machining and welding</p>	
<p>Application summary</p>	<p>Use an AI component to perform the first inspection to detect defects in radiographs and classify them into 'cleared' and 'with defects'. Use another AI component to classify these defects into 'actionable' or 'acceptable'. A human radiographer would only sample review the AI decisions.</p> <p>Target of 90% probability of detection of actionable defects at 95% confidence level, which is the current industry standard</p> <p>The system is being developed in-house. The AI components are fully supervised machine learning based and static (once the model has been trained it is fixed and does not learn anymore from the data). The models are trained and tested on images of each defect type without any 'normal' images being in the train or test set</p>	<p>Use an AI component to characterise waste sludge and advise the operator how much to fill the waste storage skip. The system uses computer vision to analyse video footage of sludge drops, measuring deformation to infer shear strength. The AI system segments images to track perimeter changes and uses established algorithms for analysis. The AI provides advisory outputs (e.g. a numerical measure of runniness), with qualified personnel making final decisions on skip filling</p> <p>Target accuracy: The preliminary performance objective is for the AI component to correctly characterise at least two out of every three sludge drops when assessed in sets of three. The system will be developed in house using convolutional neural networks using supervised learning and a static model. The training dataset has not been developed yet, but it is likely that it will be created using a test rig where different materials of known rheological properties are dropped. The AI component performs visual identification of specific items, replacing the manual step currently performed by operators. A secondary non-AI gamma detection system verifies the radioactivity of the AI-identified items</p> <p>The existing manual process achieves zero misconsignment errors, as it combines careful human inspection with gamma checks, and the objective is to keep the same. Target accuracy: Current ML detection performance is approximately 80%. The key objective is to increase the throughput of identifying those items without compromising accuracy</p>

Item	Use Case 1	Use Case 2
System Architecture	Forms part of overall defence-in-depth approach to environmental, safety, security and safeguards-related loss or harm reduction. Defence-in-depth is realised primarily through combining independent levels of protection that would have to fail consecutively before harm arises	Forms part of overall defence-in-depth approach to environmental, safety, security and safeguards-related loss or harm reduction. Defence-in-depth is realised primarily through combining independent levels of protection that would have to fail consecutively before harm arises
Deep dive topics	<ol style="list-style-type: none"> <li>1. Requirements for AI component</li> <li>2. AI component meets its requirements</li> <li>3. Post-monitoring and re-validation</li> <li>4. AI-human teaming, user trust and decision support</li> </ol>	<ol style="list-style-type: none"> <li>1. Requirements for AI component</li> <li>2. AI component meets its requirements</li> <li>3. Post-monitoring and re-validation</li> <li>4. AI-human teaming, user trust and decision support</li> </ol>
Notes	High level estimation of benefit in the region of £20 million for a typical new build project	<p>Sludge rheology: A potential saving of £100 to 500 million has been publicly published by Sellafield Ltd in their Annual R&amp;D Report, page 6, here: <a href="#">Sellafield Ltd RD&amp;I Review 2024/25 - 1</a></p> <p>Sort and segregation: Additionally, at the 2026 Nuclear Waste and Decommissioning Research Forum Seminar (11-12 March 2026), NDA shared an estimate of more than £400 million savings for the technology if developed and applied to the wide range of mixed wastes within the estate: "Autonomous sorting and segregation could improve the accuracy of waste sentencing, remove people from harm and save £400m+ in waste storage and disposal costs"</p>

## Annex 5 – Project impact and benefit realisation

This annex provides additional details of how the project delivered impact and benefit. This is presented in different formats, with an increasing level of detail:

- Table 2 provides a link between the project outputs and outcomes and the AICF objectives, by exploring how the project aims have been realised;
- Figure 5 summarises in a graphical format how the project activities, outputs and outcomes meet the project aims; and
- Further details on the link between the activities of this regulatory sandboxing project and project objectives are available in a claim argument and evidence (CAE<sup>32, 33</sup>) format at [Innovation | Office for Nuclear Regulation](#).

The options above provide complementary representation of the project benefit realisation, which may suit the preferences of different readers.

---

<sup>32</sup>[Claims Arguments Evidence Framework \) | Adelard](#)

<sup>33</sup> P Bishop, R Bloomfield and S Guerra. The future of goal-based assurance cases. In Proceedings of Workshop on Assurance Cases. Supplemental Volume of the 2004 International Conference on Dependable Systems and Networks, pp. 390-395, Florence, Italy, June 2004.

**Table 2: Summary of achievements against project aims**

<b>AI Capability Fund objective</b>	<b>Project aim</b>	<b>Project output / outcome</b>
<p>“Support or explore initiatives to develop the regulation of Artificial Intelligence in the UK.”</p>	<p>To develop, inform and enable nuclear regulators’ capability to regulate AI applications proportionately and effectively, by working with recognised experts in the sectors and other regulators on practical test cases in a safe environment (regulatory sandboxing).</p>	<p>In collaboration with other regulators and in dialogue with the industry, the project supported progress in understanding how AI may be regulated in the future in the nuclear sector, with particular advancement for a specific type of AI (narrow ML).</p>
<p>“Support or explore novel approaches to regulation through using Artificial Intelligence to improve outcomes i.e. for public safety, for economic growth.”</p>	<p>Through the test cases, to explore opportunities for nuclear regulators to reduce project risk in determining the regulatory readiness of new AI solutions in specific applications. This will encourage wider adoption of AI within the sector, which can help to improve regulatory services and compliance for the nuclear industry.</p>	<p>The discussions with dutyholders were helpful for regulators to understand opportunities where AI can support efficiencies and growth in the nuclear sector. Based on the information available for this project, no fundamental concerns were raised which could prevent further development and initial trials of the use cases considered.</p>
<p>“Support the initiatives set out in the AI Opportunities Action Plan.”</p>	<p>To support a number of actions from the above action plan, including action 25 (“regulators to scale up their AI capabilities”), action 27 (“regulators to accelerate AI in priority sectors and implement pro-innovation initiatives like regulatory sandboxes”), and 29 (“AI assurance ecosystem to increase trust and adoption by: Investing significantly in the development of new assurance tools”).</p>	<p>With the support of AI experts with experience from other industries, the set of workshops on AI design and assurance and those as part of the regulatory sandboxing were helpful in creating awareness and understanding amongst project partners (including nuclear and non-nuclear regulators) and in developing collaborative working relationships that would facilitate future engagement on other applications.</p>

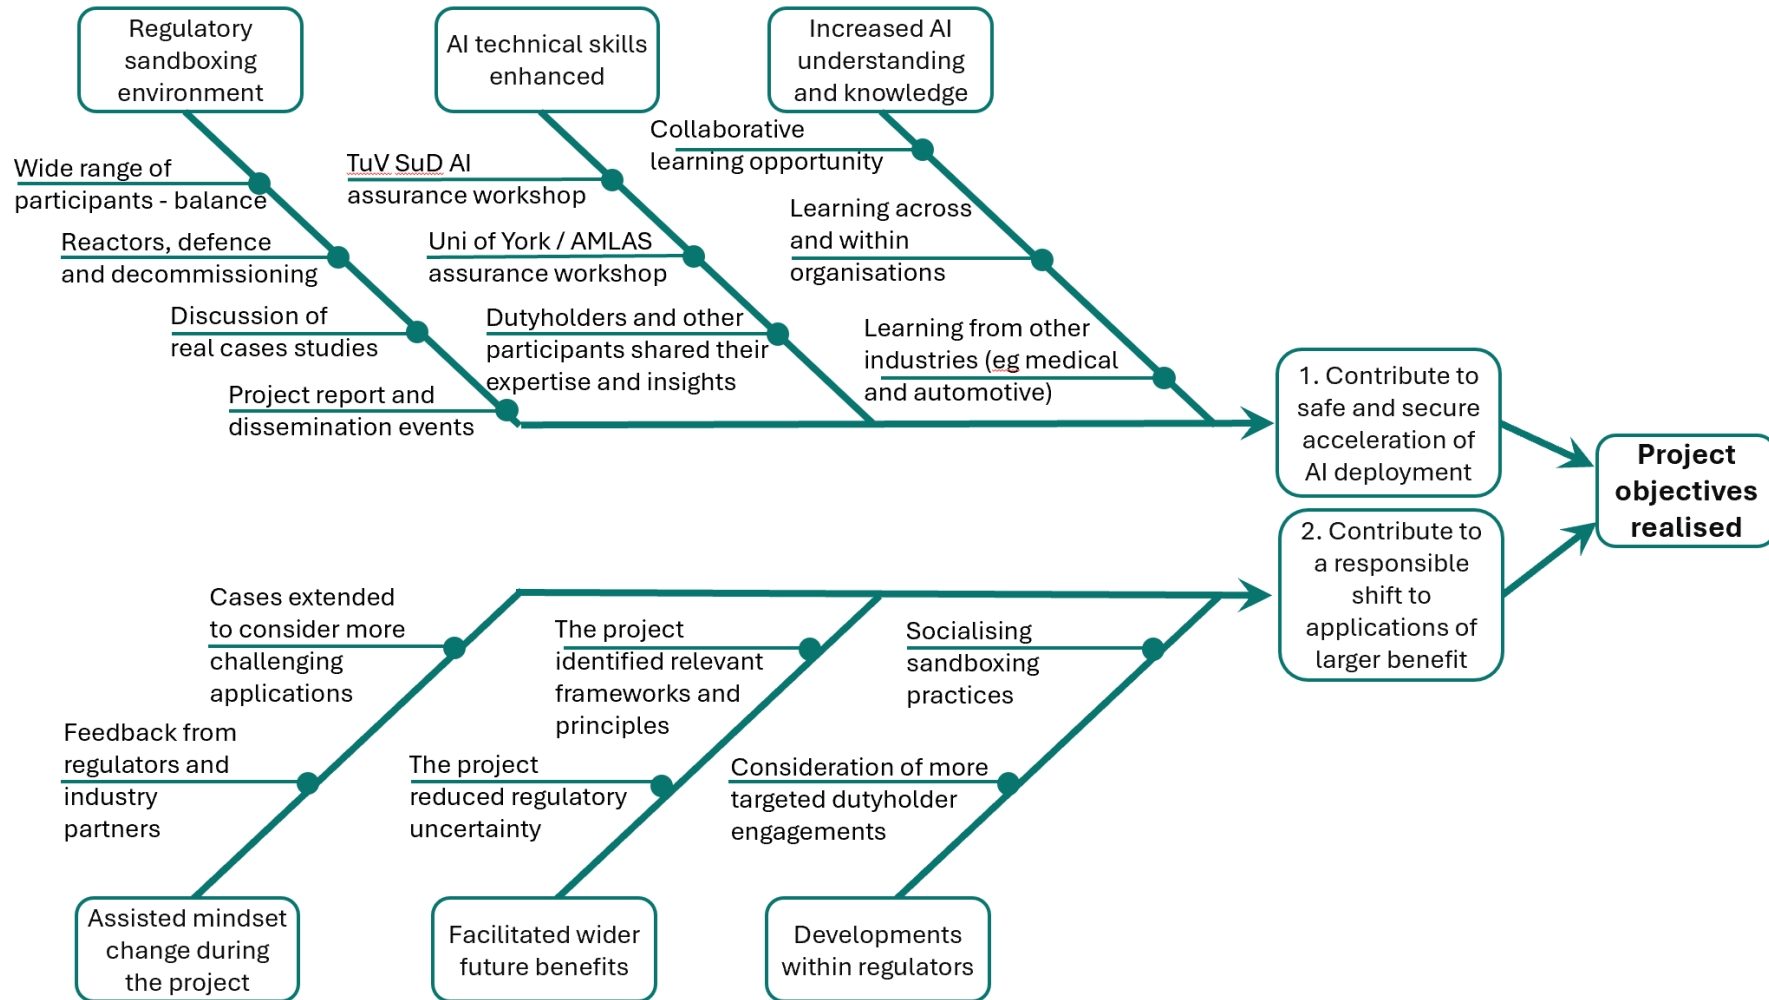


Figure 5: Fishbone diagram to illustrate project impacts