



Global
Safety
Evidence
Centre

Learning from innovation: Case study synthesis of Safetytech Accelerator projects



Safe
Work



Evidence
Review



Technical
Report



Preface

Lloyd's Register Foundation partnered with researchers at Nottingham Trent University Business School to systematically mine the *Safetytech Accelerator* case study bank for evidence of effectiveness and innovation. By synthesizing these case studies, the Foundation aimed to understand the true value of case study-level evidence and develop robust methodologies to support the Global Safety Evidence Centre. This work helps build an evidence pipeline, ensuring that insights are relevant, accessible, and actionable for diverse audiences, and strengthens the Foundation's ability to learn from practice and drive safety innovation. For more information on the Centre, please visit:

gsec.lrfoundation.org.uk

About the Lloyd's Register Foundation Global Safety Evidence Centre

The Lloyd's Register Foundation Global Safety Evidence Centre is a hub for anyone who needs to know 'what works' to make people safer. The Centre collates, creates and communicates the best available safety evidence from the Foundation, our partners and other sources on both the nature and scale of global safety challenges, and what works to address them. It works with partners to identify and fill gaps in the evidence, and to use the evidence for action.

To find out more about the Global Safety Evidence Centre,

visit gsec.lrfoundation.org.uk

About Lloyd's Register Foundation

Lloyd's Register Foundation is an independent global safety charity that supports research, innovation, and education to make the world a safer place. Its mission is to use the best evidence and insight to help the global community focus on tackling the world's most pressing safety and risk challenges.

To find out more about Lloyd's Register Foundation, visit lrfoundation.org.uk

Lloyd's Register Foundation, 71 Fenchurch Street, London, EC3M 4BS, United Kingdom

Lloyd's Register Foundation is a Registered Charity (Reg. no. 1145988) and limited company. (Reg. no. 7905861) registered in England and Wales, and owner of Lloyd's Register Group Limited.

Copyright © Lloyd's Register Foundation, 2025.

This work is licensed under [CC BY-SA 4.0](https://creativecommons.org/licenses/by-sa/4.0/)

doi.org/10.60743/Q9HX-9G38

About Nottingham Trent University Business School

This work was conducted by researchers at the Centre for People, Work and Organizational Practice (*CPWOP*) at Nottingham Business School (*NBS*), Nottingham Trent University (NTU). NBS is distinguished by its *EQUIS*, *AACSB*, and *AMBA* accreditations, reflecting its excellence in business and community engagement. CPWOP has a proven track record of delivering employer- and industry-focused research, supported by a robust infrastructure.

NTU has collaborated with Lloyd's Register Foundation on a series of reports addressing psychological wellbeing and safety in employment contexts. These include a rapid review of evidence on psychological wellbeing and safety in a global context; an exploration of the wellbeing agenda in relation to safety issues in the wake of COVID-19; and a rapid evidence assessment of the core literature around seafarer wellbeing.

This report was authored by: **Brown, S.D., Dahill, D., Smith, S., Abreu Scherer, I. & King, D.**



Executive summary

Case studies are a valuable form of evidence within innovation programmes, particularly in the context of small-scale projects, pilot initiatives, and early-stage product development. Yet there is a tendency for completed case studies to be archived without further evaluation, meaning that they are often underutilised as a body of knowledge in the medium and longer term. This report details the application of Case Study Synthesis (CSS), a methodology for comparing and analysing case study materials, to 26 publicly available case studies of innovation projects supported by Safetytech Accelerator (STA).

The method involves a systematic stepwise approach where initial research questions are defined along with a body of relevant case materials. Data is then extracted from each case study and subject to quantitative scoring across a number of quality domains (Integrity, Transparency, Completeness, Responsibility, Format and Learning). This enables comparison across cases, leading to a qualitative thematic analysis of emergent shared features and challenges. Expert interviews with 4 stakeholders along with meetings with the Safetytech Accelerator team were held to clarify and augment contextual details.

The quantitative analysis of the cases indicates that the majority of case studies provided a reasonable number of details covered by each of the quality domains, with a clustering of 7 highly scoring cases and a small number of low scoring cases, mainly due to missing details.

The qualitative analysis drew out a series of emergent themes. There were differences in the ways that innovation challenges were set and evolved across the projects. The cases demonstrate a number of ways of addressing the safety problem, including solving, managing, preventing and reinventing. The routes by which safety outcomes were reached can be conceptualised in terms of relative positions of the projects either upstream or downstream of the ultimate safety problems. Across the cases, the focus varied between ensuring compliance with safety standards or supporting professional competencies, which had implications for how professional knowledge was reconfigured. Different approaches were taken to evaluation within the projects and, where relevant, to addressing assurance for safety intervention.

The analysis has shown that 'safety' is not simply an initial condition or standard that determines how projects are implemented but is rather an evolving criterion that in some cases is transformed through the innovation process. The Hierarchy of Controls approach is useful in describing how innovation displaces problems, but Actor-Network Theory (ANT) provides a more rounded understanding of how safety problems are translated across the Safetytech Accelerator projects. In particular, ANT draws attention to the networks of stakeholders, practices and technologies which are developed across projects and to the longer-term durability of these networks.

Case study synthesis is then able to generate actionable knowledge from the retrospective analysis of existing case studies. It can show contextual details and shared features which are not wholly contained within any single case study, and which can inform future funding decisions. We make 13 recommendations about the future presentation of case studies for external audiences and offer a series of reflections on how case study synthesis can assist in transforming professional knowledge, recognising the status of case studies as legitimate forms of evidence, supporting longitudinal and iterative learning, tailoring evaluation metrics to innovation stages and fostering cross-sectional and interdisciplinary collaboration.



Contents

Preface	1
Executive summary	2
Figures, tables, and boxes	3
1. Introduction	4
2. Context	4
3. Method	4
3.1 Stepwise approach	5
3.2 Expert interviews	6
3.3 Research questions	6
4. Quantitative analysis	7
4.1 Sectoral distribution	7
4.2 Safety challenges addressed	8
4.3 Commissioning clients and location	8
4.4 Technology developers	8
4.5 Quality appraisal	9
4.6 Limitations	9
5. Qualitative analysis of case studies	9
5.1 Defining a safety problem	9
5.2 Addressing a safety problem and defining an intervention	10
5.3 Routes to safety outcomes	12
5.4 Fostering competence and ensuring compliance	12
5.5 Evaluation and assurance of safety interventions	13
6. Discussion	15
7. Conclusion and guidance	17

Figures, tables, and boxes

Figure 1: Example of scoring across domains	6
Figure 2: Sectoral distribution of Safetytech Accelerator case studies	7
Figure 3: Distribution of quality appraisal scores	9
Figure 4: Case studies by Safetytech Accelerator challenge programme	9
Table 1. Research questions	7
Table 2. Safety challenge types	8



1. Introduction

Case studies are a valuable form of evidence within innovation programmes, particularly in the context of small-scale projects, pilot initiatives, and early-stage product development. Yet there is a tendency for completed case studies to be archived without further evaluation, meaning that they are often underutilised as a body of knowledge in the medium and longer term. Case studies can offer rich, real-world insights, and, when synthesised effectively, can illuminate the enablers and barriers to implementation. Case Study Synthesis (CSS) is a method designed to provide insights through comparison of existing case study materials. It aims to highlight unexpected outcomes, and surface practical considerations that have emerged in dynamic, real-life settings. The potential of this method is to act as an effective tool for generating transferable learning on project development, adaptation, implementation, and scaling – insights that are highly relevant to both funders and delivery organisations.

The key aim of this report is to advance the methodology for creating and synthesising case studies, with a focus on innovation contexts where conventional evidence can be limited or still emerging. The objectives include exploring how case studies can complement other forms of evidence, exploring what insights can be generated through retrospective synthesis of publicly available case, and assessing how these insights can inform future decision-making. The report also considers some of the practical challenges of data access and evidence review, especially in less mature fields. The case study synthesis method offers an opportunity to better understand how user needs may be aligned with the real-world requirements of funders, researchers, and practitioners. The report aims to build trust in case study synthesis as a credible and useful approach, while identifying opportunities for its broader application across innovation and early-stage development landscapes.

2. Context

Innovation programmes often operate in environments where conventional forms of evidence – such as formal evaluations or peer-reviewed research – are limited, fragmented, or still emerging. This creates a challenge for funders, practitioners, and researchers seeking to understand the effectiveness, scalability, and broader impact of such initiatives. Case studies are a standard way of reporting the outcomes of innovations and interventions and can provide powerful narratives and real-world insights about the implementation of funded work. However, their potential is often constrained by inconsistent formats, variable depth of reporting, and a lack of standardised synthesis methodologies. As a result, the learning embedded within individual case studies may remain siloed, anecdotal, or difficult to translate into actionable knowledge.

This report addresses the need for a more systematic and credible approach to case study synthesis within innovation settings. It defines the problem as twofold: first, the under-recognition of case studies as a legitimate and complementary form of evidence; and

second, the methodological challenges associated with synthesising diverse case materials to generate transferable insights. By adapting and applying an established synthesis methodology to the Safetytech Accelerator (STA) programmes, this work aims to demonstrate how case studies can be used more effectively to inform decision-making, reflect the full value of funder investment, and support future innovation strategy.

Safetytech Accelerator supports research and development work in safety-critical industries and around safety infrastructure. It does this by supporting innovation projects commissioned in response to challenges set by industry partners. Safetytech Accelerator was established in 2018 by Lloyd's Register and became an autonomous entity in 2021. Safetytech Accelerator supports feasibility and pilot studies across a range of safety-relevant areas, with particular expertise around safety, operational risk, performance and resilience and sustainability. The accelerator aims to support the growth and development of safetytech as a market.

The publicly available material for analysis on Safetytech Accelerator programmes comprises 26 individual case studies (along with some shorter summary materials) published on their website. Through discussion with Safetytech Accelerator, it was made clear that these case studies have been produced over several years through a structured approach using defined sections describing the formulation of the challenge, the problem definition and approach, results and industry implications. Although the published case studies are intended to provide good overviews of each project, Safetytech Accelerator regard their primary use as marketing materials, with further technical details and reports existing within the organization. There are differences in the level of detail and description across the cases, and where possible we sought to elicit information through meetings and interviews with Safetytech Accelerator staff and challenge stakeholders. The aim of the work described in this report was to synthesise the publicly available materials across the Safetytech Accelerator case studies to generate insights and further value above and beyond that already accumulated through the delivery of the individual funded projects.

3. Method

CSS is a systematic method for integrating findings from multiple case studies to develop broader, more robust conclusions. This form of evidence synthesis aims to extract meaningful patterns, themes, and insights from individual case studies, enabling researchers to draw more generalisable conclusions or identify key factors across different contexts (Thomas and Harden, 2008; Mills et al., 2010).

The What Works for Wellbeing Centre Guide to Synthesising Case Studies (Hardoon et al., 2021) outlines a step-by-step approach to synthesising practice-based case studies. Key steps include defining the research question, selecting relevant case studies, extracting data, synthesising findings, and interpreting results. Building on this strategy, South et al. (2024) emphasise the development of a conceptual framework to define and categorise interventions. They provide detailed guidance on identifying websites and case study



collections and define more structured templates for organising data, adding a layer of systematic organisation. Furthermore, in cross-case analysis, the use of matrices offers a more structured approach to identifying patterns and differences. This method favours writing a narrative report enriched with contextual information and illustrative quotations.

3.1 Stepwise approach

Our approach for case study synthesis combined and analysed multiple case studies, drawing broader conclusions and insights. This project followed a systematic, transparent, and pragmatic approach to synthesising case study evidence via the following steps:

A. Developed research questions and conceptual framework

- Defined review questions: Research questions were collaboratively developed with stakeholders to ensure relevance and focus.
- Conceptual framework: A framework was created to develop research questions and further refined during the analysis stage.

B. Identified and searched for evidence

- Search strategy and selection criteria: A comprehensive and clear selection criteria were developed.
- Evidence identification: Keyword searches in online databases and repositories were used to identify relevant websites and case study collections.

C. Selected studies

- Inclusion and exclusion criteria: Criteria remained open, as the number of case studies provided by Safetytech Accelerator was fixed prior to data extraction. No sifting of cases was conducted either by Safetytech Accelerator or as part of the case study synthesis.

D. Extracted and organised data

- Data extraction template: A structured template was used to systematically gather relevant data fields from each case study.
- Data organisation: Data were organised using common fields/domains to display extracted data.

E. Assessed quality

Quality evaluation

The quality of case studies was assessed based on the domains of integrity, transparency, completeness, responsibility, format, and learning reported. The scoring criteria for each domain is described below.

Quality assessment tool

To ensure the credibility and utility of the case study synthesis, each case study was subject to a structured quality assessment. The tool used for this purpose was a bespoke quality appraisal framework, adapted from established guidance (e.g. Hardoon et al., 2021; South et al., 2024) and tailored to the specific context of innovation case studies.

Quality assessment tool domains

The assessment framework comprised five domains, each reflecting a key dimension of quality relevant to practice-based case studies. Each domain was scored on a scale from 0 to 4, with higher scores indicating stronger performance in that domain. The maximum possible score for each case study was 20 points (5 domains × 4 points each).

Domains and scoring criteria

- Integrity (0–4 points)
 - Assesses the accuracy, honesty, and reliability of the case study.
 - Criteria include: clear description of context, transparency about methods, and avoidance of selective reporting.
- Transparency (0–4 points)
 - Evaluates the openness with which the case study reports its processes and findings.
 - Criteria include: explicitness about data sources, clarity in reporting outcomes, and disclosure of limitations.
- Completeness (0–4 points)
 - Measures the extent to which the case study provides a full account of the intervention or innovation.
 - Criteria include: coverage of background, implementation, outcomes, and lessons learned.
- Responsibility (0–4 points)
 - Examines the ethical and social responsibility demonstrated in the case study.
 - Criteria include: attention to participant consent, safeguarding, and consideration of wider impacts.
- Format and Learning Reported (0–4 points)
 - Assesses the accessibility and usefulness of the case study for learning and future application.
 - Criteria include clarity of writing, use of illustrative examples, and articulation of transferable lessons.



Scoring

Each domain was scored independently by two reviewers, using a rubric that specified what constituted a score of 0 (absent), 1 (limited), 2 (adequate), 3 (good), or 4 (excellent) for each criterion. The total quality score for each case study was calculated out of a possible 20 points. Reviews were then compared and moderated

Figure 1. Example of scoring across domains

DIMENSION	SCORING Completely (2) Partially (1) Not at all/CAN'T tell (0) N/A									
Study ID		1	2	3	4	5	6	7	8	9
Integrity	Is an evaluation or learning method described?	1	1	1	1	1	1	1	1	1
	Is the writing accurate, balanced and objective?	1	0	0	0	0	0	0	0	0
	Is the evidence base used? (Published papers & work of relevance; evidence informed theory)	0	0	0	0	0	0	0	0	0
	Is there attribution of authorship and contribution?	0	0	0	0	0	0	0	0	0
Completeness	Is the setting (i.e. organisation or sector) described?	1	1	2	1	2	1	1	1	1
	Is the population described?	0	1	2	2	2	1	1	1	1
	Is the geography defined?	1	1	1	1	2	1	0	1	1
	Is the intervention described?	1	1	1	1	2	2	2	1	0
Transparency	Are the aims/objectives of the intervention clear?	1	1	1	1	2	1	2	1	0
	Are the outcomes of the intervention clear?	1	1	1	1	2	1	2	1	0
	Is there thick description?	0	0	0	0	0	0	2	0	0
	Are all results published regardless of outcome?	0	0	0	0	0	0	0	0	0
Responsibility	Is it clear why this case study was written?	1	1	1	1	1	1	1	1	1
	Are the funding sources and sponsors described?	1	1	1	1	1	1	1	1	1
	Are any potential conflicts of interest disclosed?	0	0	0	0	0	0	0	0	0
	Are the data collected made accessible? (e.g. link to empirical data)?	0	0	0	0	0	0	0	0	0
Format	Is there discussion of any limitations of the intervention?	1	0	0	0	0	0	0	0	0
	Is there discussion of any limitations of the evaluation / research?	1	0	0	0	0	0	0	0	0
	Are the dates of when the project took place provided?	0	0	0	0	0	1	1	0	0
	Is there a clear statement that peer-review or evaluation of the case study has been undertaken?	0	0	0	0	0	0	0	0	0
Key learnings/recommendations	Is there an external evaluator?	0	0	0	0	0	0	0	0	0
	Is the content in a suitable format for other practitioners?	1	1	1	1	1	1	1	1	2
	Is there a clear structure?	1	1	1	1	1	1	1	1	2
	Do the case documents report key learning and/or make recommendations based on learning?	0	1	0	1	1	0	0	0	1
Total		10	11	10	10	14	10	12	8	8

F. Cross-Case analysis and synthesis

- Framework analysis: Thematic coding was applied, and matrices were developed to identify patterns and themes. This involved:
 - Familiarisation: Researchers immersed themselves in the data by reading and re-reading the case studies.
 - Identifying a thematic framework: Key themes and concepts were identified based on research questions and objectives.
 - Indexing: The data were systematically coded according to identified themes.
 - Charting: Coded data were organised into charts or matrices, summarising data under relevant themes.
 - Mapping and interpretation: Charts were used to identify patterns, relationships, and key findings, drawing conclusions to answer research questions.

G. Developed an overarching framework

- Explanatory framework: A framework that fitted with the data in the sample was produced.

H. Reported findings

- Narrative report: A report was written, grouping results around higher-order themes or categories, including contextual information and illustrative quotations.
- Synthesis results: A narrative account of synthesis results was provided, including an overview of case studies, major themes, summary tables, and quality appraisal results.

3.2 Expert interviews

To complement the case study synthesis, four expert interviews were conducted with representatives of technology developers involved in different case studies to gain deeper insights into underreported areas and validate findings. This involved:

A. Interview design

- Question development: Interview questions were developed based on gaps identified in the case study data.
- Participant selection: Representatives of technology developers were identified and contacted by Safetytech Accelerator.

B. Data collection

- Conducted interviews: Open discussions and semi-structured interviews were used to allow for in-depth exploration of topics.
- Recording and transcription: Interviews were recorded and transcribed for accurate data capture.

C. Data analysis

- Thematic analysis: Interview transcripts were analysed using thematic coding to identify key themes and insights.
- Integration with case study data: Interview findings were integrated with case study data to enhance the overall analysis.

3.3 Research questions

The research questions that guided this synthesis were:

1. What were the settings, sectors, technologies, and outcomes of the pilots?
2. What patterns, trends, and gaps could be identified across these dimensions?
3. What were the enablers, barriers, and implementation challenges of the pilots?
4. What was the specific role of the Safetytech Accelerator in influencing implementation and outcomes?
5. How suitable were the case studies for synthesis, and what was the overall robustness of this methodological approach?

These questions were partly informed by the Hierarchy of Controls approach (see Burk, 2016). Together they ensured a focused, practical, and methodologically sound synthesis process.

**Table 1. Research questions**

#	Research Question	Feasibility/Initial assessment	Next steps/Suggested Approach
1.a	What were the settings, workplaces and sectors in which the Safetytech pilots took place? Who were the people involved?	We are confident we can draw a good picture of these with the data and show some trends and gaps.	Proceed with analysis
1.b	Who were the commissioning clients and tech suppliers involved?		
1.c	What were the safety challenges which the pilots intended to address?		
1.d	What were the interventions involved in the pilots? Which types of Safetytech were involved in the pilots and how were they applied?	We feel that we should be able to answer this from the case studies.	We could usefully do a typology of technologies used, range, trends, whether they were transferred from other sectors, etc.
2	What were the outcomes of the pilots in terms of safety? What other outcomes were reported? How were these measured or assessed?	Reporting is patchy and variable.	Extract what is available. Note gaps. Flag need for stronger reporting guidance in future.
3	What were the challenges involved in carrying out the pilots? What were the enablers and barriers to the success of the pilots?	May not be reported in depth enough in the case studies.	Could be answered with follow-up interviews with a smaller set (clients/suppliers/STA staff) but would require access to appropriate interviewees.
4	What was the role of the Safetytech Accelerator in the implementation and outcomes of the pilots?	Unlikely to have data from these case studies.	Could be answered with follow up interviews with a smaller set (clients/suppliers/STA staff) but would require access to appropriate interviewees.
5	What was the suitability of the case studies for synthesis and to answer the research questions?	Can be assessed through analysis.	Feed into development of synthesis methodology and guidance. Review structure and content gaps.

4. Quantitative analysis

A total of 26 case studies were included in the synthesis, each representing a pilot or feasibility study supported by Safetytech Accelerator (see appendix A – please note that CS10 is a shortened description of CS7 and was treated as a single case study). Most of the studies were classified as pilot studies (24 out of 26), with only two described as feasibility studies, reflecting emphasis on testing early-stage technologies in real-world operational settings. There do not appear to be any substantive differences between the two kinds of studies. The case studies were published between 2019 and 2024, with a notable increase in activity, with 17 case studies from 2022 onwards. This trend may reflect a growing institutional interest

in documenting innovation pilots, particularly in response to post-pandemic recovery and digital transformation agendas. These studies span a diverse range of sectors, technologies, and safety challenges, which we have outlined below, with the analysis focusing on study characteristics – i.e. the types of safety challenges addressed, the technologies trialled, and so on.

4.1 Sectoral distribution

Unsurprisingly, given Safetytech Accelerator's origins, the sector most represented was shipping, which featured in over half of the case studies. Other sectors included construction, food production, education, and energy.

Figure 2. Sectoral distribution of Safetytech Accelerator case studies



4.2 Safety challenges addressed

The most frequently cited safety challenge was faster fire detection in containers on ships, underscoring a critical concern in maritime safety, however the case studies addressed a wide range of safety challenges (listed below). While some were succinctly defined (e.g. “data security” or “crew fitness for duty”), others were described in more narrative terms. Each challenge was coded and counted to identify patterns across the dataset.

Table 2. Safety challenge types

Safety challenge type	Number of case studies
Faster fire detection in containers on ships	3
Manual data collection and operational inefficiencies	2
Working at height and fall prevention	2
Mental wellbeing and stress detection	2
Food safety and allergen transparency	2
Predictive maintenance and digital twins	1
Compliance automation and document analysis	2
Remote inspection and defect detection	2
Construction site safety and AI-based risk detection	2
Safety in ports and terminals	2
Data security and anonymisation	1
Crew alerting and positioning in emergencies	1
Training and simulation for maritime operations	1
Listeria detection in food production	1
Pipeline defect identification	1
Robotic hull cleaning and biofouling prevention	1
Electrical cabinet inspection and error detection	1
Human error and behavioural risk	1

Note: Some challenges were described in long-form narrative text. These were manually reviewed and categorised to ensure consistency.

4.3 Commissioning clients and location

While some case studies lacked precise geographic detail, those that did were spread across Europe, North America, Asia, and Australia. The most frequently cited locations included Singapore, Germany, and the US, reflecting the global scope of the programme.

The case studies featured a wide range of commissioning organisations, including global corporations, public sector bodies, and industry consortia. Notable names include:

- Cargill, Seaspan, PepsiCo, and Phillips 66 from the private sector
- Sellafield National Nuclear Laboratory, Health and Safety Executive (HSE), and University of Manchester from the public and research sectors
- Lloyd's Register Foundation and Safetytech Accelerator as recurring funders and conveners
- Anchor partners such as Evergreen, Maersk, and Ocean Network Express (ONE) involved in collaborative maritime initiatives

Some organisations appeared multiple times, due to repeated involvement across different challenges and pilots. Plug and Play, a separate technology accelerator, also collaborated on some of the earlier Safetytech Accelerator projects.

4.4 Technology developers

The pilots employed a wide range of technologies, with software-based solutions being the most prevalent. These included AI-driven analytics platforms, computer vision systems, and digital twins. Hardware innovations such as robotic devices and sensor networks were also well represented. A diverse set of technology providers were engaged, ranging from start-ups to established firms, with expertise spanning AI, robotics, sensor networks, and digital platforms. Examples include:

- Alicia Bots (robotic hull cleaning)
- Everactive (energy-harvesting sensors)
- SnapDNA (rapid pathogen detection)
- Clarifai (AI-based image and text analytics)
- MonoLets (wireless mesh networks)
- audEERING (voice-based emotion analytics)
- SmartNanotubes Technologies (SNT) (e-nose sensors)


This diversity raises interesting questions about sector-specific vs. cross-sector applicability of safety technologies. For instance, is fire detection fundamentally similar across maritime and industrial contexts, or do sectoral nuances demand tailored solutions? The synthesis suggests that while some technologies were adapted from other domains, others were developed specifically for the sector in question.



4.5 Quality appraisal

Each case study was assessed using a structured quality appraisal framework using the 5 domains outlined above (Integrity, Transparency, Completeness, Responsibility, Format and Learning) on a scale of 1–4, giving a maximum possible overall score of 20. The average quality score across all studies was 10.3, with a median score of 10, indicating moderate consistency in reporting. However, variability in completeness and transparency suggests opportunities for improving future case study documentation.

Figure 3. Distribution of quality appraisal scores



The aggregating of quantitative scores in this way provides only a general sense of the distribution of the quality domains across the cases. The clustering around the mid-score of 10 demonstrates that the majority of cases included a reasonable number of details which afforded opportunities for future learning. The lower scores in the is also worth noting that there 7 cases in 13–15 range, who score consistently across a number of domains.

4.6 Limitations

The review team of two reviewers were responsible for most assessments, which were undertaken jointly or with sense checking meetings to ensure consistency. Whilst steps were taken to limit issues such as reviewer bias (including regular check-ins and update reports with the PI and methods consultant) the team acknowledges that this never fully eradicates implications around consistency and inter-rater reliability in quality appraisal. Each case study was assessed using a structured quality appraisal framework.

5. Qualitative analysis of case studies

Thematic analysis was applied to both case study materials and interview transcripts. Recurrent themes were identified along with contextual differences and illustrative insights. This process involved familiarisation, indexing, charting, and interpretation, which culminating in a rich, narrative synthesis of the data.

5.1 Defining a safety problem

The case studies supported by Safetytech Accelerator operate through a challenge-led process where problems are defined by stakeholders. Technology developers are then recruited, in most cases through a competitive process. Of the 26 distinct case studies reported, the majority (12) arose through individual challenge setting, 5 through the Waypoint programme, 2 through the Cargo Fire and Loss Innovation Initiative (CFLII), 3 in collaboration with the HSE Discovering Safety programme, 1 through a partnership with the National Safety Council and the remaining 3 from challenges set by Lloyd's Register.

Figure 4. Case studies by Safetytech Accelerator challenge programme





The published case studies do not provide full details of the process of formulating and publishing the individual challenges. However, there does appear to be a relationship between the number of partners involved in a challenge and the longer-term acceptability and implementation of the safety innovation. For example, CS13 ('Developing ship handling skills using virtual reality') lists six partners, including the technology developer (Kilo). A version of the VR ship handling platform which was configured within the project was already under development by Kilo and has been further refined since across a number of other subsequent projects with different partners as 'VASCO'. By contrast, CS5 ('Anonymising and desensitizing health and safety data') lists four partners, including the technology developer (Ohalo). The developer configured their existing system 'Data X-Ray' as a solution to the challenge problem of anonymising data held by the Health and Safety Executive (HSE) and were ultimately successful in contracting the solution to the challenge partner for a number of years following on from the project itself. This suggests that a smaller number of project stakeholders is associated with a higher degree of 'stabilisation' of the innovation at the time of project and a quicker route to its implementation in that form. But it is important to note that with CS13, Kilo have continued to develop their presence within the Safetytech market, yet following CS5, Ohalo are now predominantly positioned in relation to security and AI.

Some of the problems formulated within the challenges are versions of longstanding problems. For instance, the CFLII programme addresses the problem of large-scale cargo loss at sea. Whilst this is a problem that is deeply historical in its nature, it is clearly given new impetus by the scale and complexity of modern cargo vessels. The case studies in this area might then be seen as 'new solutions to old problems'. For instance, CS20 ('Pioneering WiFi-based fire detection technology') presents an innovative solution to the early detection of fires in cargo holds by switching away from the existing technique of detecting smoke particles towards detecting temperature changes through variations in radio signals within a wi-fi based sensor grid. The technology itself was already at a high level of readiness, with the innovation here being its use in a very different context.

A slightly different approach is taken in CS4 ('A pilot to explore robotic hull cleaning'). Here the technology – the robotic device 'Roverclean' – was already designed to meet the challenge problem of cleaning the hulls of vessels without the use of human divers or dry-docking. The project then evaluated the efficacy of the new technology and demonstrated the further development required for it to meet the operational requirements for its implementation at sea, rather than as maritime port-based practice. Inevitably, this kind of new solution to an old problem then creates further issues to be addressed, such as the need to train crews in the use of Roverclean whilst at sea, defining the maritime conditions for its safe use and the need for spare parts and maintenance. This is arguably a displacement of one safety problem for another, although the two problems differ in terms of their overall shape and potential significance to the stakeholders involved (i.e. risk to human life versus operational safety of the robotic technology).

There are a number of case studies where the innovative nature of the technology has the potential to either restructure the nature of the initial safety problem or else to reveal the existence of previously unknown safety issues. CS3 ('A pilot to explore automated data collection from ship machinery'), for instance, developed non-invasive means of collecting real-time data on the engine and navigational performance of vessels at sea along with analytics to remotely process the data. Whilst the safety intervention here to some extent replaces an existing practice of manually reporting information, it does so in a way that allows for different kinds of safety-relevant relationships to be established in real-time data on vessel performance. This can be considered as a transformation rather than a displacement of the safety problem.

The definition of the safety problems involves not only the network of stakeholders around the challenge definition, but also those involved in the potential implementation of the safety intervention. For example, CS1 ('A study into using energy harvesting sensors to detect fires in cargo holds') describes a simulation pilot study in the use of sensors for the early detection of fires on board a vessel. Whilst the data from the simulation indicates the potential benefits of developing the technology further, the study itself suggests that further work to investigate the viability of deploying the sensors in a dense steel below-deck structure, involving routing cables throughout the vessel will be required. This additional work inevitably draws in crews, managers and ship-owners into the network of stakeholders that will need to be aligned to implement the safety intervention.

This kind of issue around the potential stability of the stakeholder network is also apparent in CS12 ('Communicating accurate ingredient data in school canteens'). This project reports on a safety intervention which might be implemented within schools to manage issues around children selecting foods containing either allergens or dispreferred food ingredients. The study reports on the overall acceptability of the system with staff and parents but does not indicate whether there were any issues with children as the potential primary end-users of the system. The success of the implementation would then depend upon maintaining the relationship between school staff, parents and children around the safe and efficient use of the system, including, for instance, issues around how potential disagreements between children and parents around how children actually act upon the information provided by the system in the food choices might then impact upon its use.

5.2 Addressing a safety problem and defining an intervention

The case studies cover a wide range of safety challenges which raises questions surrounding how safety is conceptualised and safety problems operationalised and addressed. Across the different interventions described in the 26 case studies, there are four different ways of working to address safety issues that are present: solving, managing, preventing, reinventing. While these areas have been identified as distinct ways of working with safety challenges, some of the interventions presented in the case studies approach the problem using combined ways of working.



Solving

In projects where the problem was solved, it ceased to be a problem both currently and in the foreseeable future. For example, CS5 ('Anonymising and desensitizing health and safety data') details an intervention to remove identifying data from health and safety reports so they could be used for the prevention of further accidents. The study used AI software and was able to anonymise a large data set in a few hours compared the predicted 12 years that it would have taken to do this manually. In this sense, the initial problem was solved in that the historical documents could be used without identifying individuals, and furthermore, while the case study only contains details of the specific intervention, it would be reasonable to assume that any documents produced later could be anonymised and added to the data set.

This raises questions around what types of safety issues can be considered to be 'solved'. While CS5 solves a tangible problem, there are current limitations in how many of the problems addressed in the case studies (or wider safety issues) are able to be resolved with a single intervention, both at the point of project implementation and going forward. Arguably, the problem in CS5 was sufficiently well defined that the key actions required – ensuring that the data for use was anonymised – could be clearly conceived and addressed. However, problems that involved multiple interacting factors and those where the safety intervention itself creates further consequential safety considerations are unlikely to be amenable to single 'solution' approaches.

Managing

The improved management of safety risks implies that the nature of the risks is sufficiently well-known and what is required is either better techniques to implement a version of existing practices or improved co-ordination of information through enhanced data monitoring and/or processing. CS19 ('Pilot to speed up detection of listeria') is a good example of the first approach. The safety challenge for the food production industry is clearly defined as the detection of listeria monocytogenes in food products, with the secondary problem of distinguishing between live and dead cells. There are existing testing procedures, but these are based off-site and slow to process. The on-site intervention demonstrated improved speed and along with reliable accuracy of diagnosis, allowing for potentially more efficient management of the problem.

An example of the second approach is CS22 ('Reducing risk in ports using AI'). Here the overall safety problem of protecting safety hazards to pedestrians on port facilities is well-established. The intervention then aims to support better management of the risk through improved data analytics around specific behaviours that are likely to increase exposure to hazards and the possibility of the automatic detection through live video fields and alerts. The system was trained to detect four distinct risky behavioural patterns, with the possibility of further configurations. This kind of intervention then offers an improved technical means to reach known safety goals through a different process which potentially offers increased economies in terms of human and financial resources.

Preventing

Several of the case studies developed interventions aimed at preventing a known safety issue from occurring. For instance, CS4 ('A Pilot to Explore Robotic Hull Cleaning and Inspection Technology') and CS16 ('Exploring Robotic Solutions to Remove the Need for Humans to Conduct Inspections in Confined Spaces') both used robotic devices, one to clean the hull of a ship and another to inspect small spaces. In CS4, the device ensures that a ship's hull can be cleaned while at sea (something that can usually only be done in port and often manually) which has the benefit of clearing any hazards, such as sea debris as it attaches to the ship rather than waiting for the ship to dock for it to be cleared (which is also not permitted in many European ports). Additionally, it ensures that dock workers or ship's crew do not have to enter the water in order to clean the ship. In the second instance the device ensures that workers do not have to enter potentially dangerous small spaces and also ensures that problems with visibility that can occur in those small spaces can be dealt with.

In both these cases, the safety intervention is an alternative which removes the need for a current, potentially hazardous practice. As such, the intervention does not solve the current problem entirely but rather opens a pathway to developing a different safety practice. Inevitably, this pathway creates further safety issues to be addressed. For example, in CS4 the RoverClean technology suffered from a variety of technical issues during testing related to the level of fouling on the ship's hulls. Similarly, in CS16 the current technical limitations around the flight time of the drones used for inspection of pressure vessels was felt to potentially compromise the thoroughness of the procedure. In this sense the prevention approach does not entirely remove the safety problem but rather translates it into a different kind of problem. However, this translation itself can be considered to have clear benefits, despite the degree of success of the intervention. As one of the anonymous testimonies in CS16 notes "Every inspection by drone or robotic inside a vessel which prevents entry by a human is a win".

Reinventing

Some of the case studies went beyond the existing legal and/or practical definition of a safety problem to define a new context in which the intervention could be demonstrated. For example, CS12 ('Communicating accurate ingredient data in school canteens') addressed a legal requirement for 14 different known allergens to be clearly reported. However, the challenge proposed that childhood food allergies are both increasing and diversifying. The intervention then went beyond the current legal definition of the safety problem by making information about a wider range of ingredients and their potential 'nesting' within one another available in advance. This redefinition or reinvention of a safety problem could be thought as either a further enhancement of safety or as an extension of what constitutes safety beyond its immediate domain. As one parental piece of feedback notes "my son has several 'dislikes' which he claims are intolerances, but it [the system] helps to filter out things that may cause a fuss and direct him which options to choose". Safety is then arguably redefined as the removal of the consequences that flow from lack of awareness around dispreferred behavioural options.



Reinventing a safety problem can also involve shifting the locus of a series of known risks into a new domain. For instance, CS14 (Discover the Safetytech Solution that can Assess if a Crew is Fit for Duty') addressed the well-established safety problem of the psychological fitness of seafarers by trialling an eye-tracking system to "quantify previously inaccessible cognitive insights". Some aspects of this system might be seen as a technical improvement to an already established practice – for example, physically observing a crew member for signs or indicators of impairment from drugs or alcohol. But the broader ambitions of the system to recognise the presence of stress or depression through eye movements is arguably a reinvention of the safety problem away from holistic behavioural observation to technical measurement of a narrow range of physical activities. The reliability and acceptability of this reinvented safety practice then depends upon the quality of underpinning evidence that eye movements are valid indicators of transient and enduring psychological states.

5.3 Routes to safety outcomes

Across the case studies there is a continuum between interventions which directly address a safety problem through to those which address either the conditions through which a safety problem might emerge or a secondary issue whose resolution may have safety-relevant consequences.

A prime example of a case study which directly addressed a known safety problem is CS8 ('An open innovation challenge to transform industrial inspection'). The problem is the legal requirement to conduct regular inspection of ventilation ducts, which is currently mostly done through manual means with high degrees of human error. The intervention was conducted with Nuclear Industry partners, with implied additional emphasis on reducing human exposure to risk. The acoustic sensor technology deployed was able to demonstrate high levels of accuracy in detecting both normal functioning and abnormalities. In this sense, the intervention appears to directly meet the task requirements in a way that removes human safety errors and exposures.

In similar vein, CS15 ('Early fire detection on container ships using E-Nose technology') addresses a longstanding maritime safety problem which is becoming 'more frequent and severe'. The existing technology is conceived as being no longer fit for purpose on modern cargo vessels, whose size and complexity require more rapid alarm systems. The intervention is then a new means of detecting fires by analyzing changing patterns of air molecules which might indicate heating precursors before the outbreak of fires themselves. Unlike CS8, this case study shifts the locus of the problem slightly away from fire detection as such to the detection of the precursors of fire yet still addresses the problem directly.

CS27 ('Understanding decision making on a ship through sensing') also moves away from safety hazards and towards their possible precursors. The intervention sought to monitor and analyse the behaviours of a crew in a ship's wheelhouse, with a particular focus on signs of stress and fatigue. The continuously recorded data was

then retrospectively correlated with the vessel's performance and safety audit records to establish the relationship between behaviours and specific operational incidents. This approach differs in that the relationship between the behavioural precursors and safety-relevant events was not firmly established in advance. The approach then offers the promise of a 'real time' solution once the technology has become sufficiently embedded.

The intervention in CS21 ('Providing Construction Safety Guidance Using Visual Observations and Historical Reports') shifts the problem to be addressed even further temporally through offering a means to exploit already captured data assets that can then inform contemporary safety practices. The pilot study used an AI platform to mine existing historical data collected by the HSE which were then used to teach a model to detect safety hazards in new images taken at construction sites. The intervention is then split between exploiting an existing resource and connecting the resulting outputs to the current safety problem.

Finally, CS5 ('Anonymising and desensitizing health and safety data') moves the furthest away from the current safety problem in its intervention around anonymising existing HSE records to form large GDPR compliant datasets which might then be subsequently exploited to develop improved safety guidance. This case study was solely focused on a technical solution to a secondary problem which might in the future then inform interventions around primary safety issues.

The continuum present across the case studies can be conceived as degrees of being either 'upstream' or 'downstream' in relation to a specific safety problem. The case studies which are more downstream and directly connected to the problem and the success criterion of the technology implementation can be directly qualified, at least in terms of possible pilot study outcomes. Those which are progressively more upstream may retain a direct connection to the downstream safety problem, as in CS21, or they may offer the promise of that connection becoming developed or strengthened in due course, as with CS27 and CS5. In these latter cases, further and perhaps distinct forms of technical innovation may be required to realise those promises.

5.4 Fostering competence and ensuring compliance

A number of case studies are concerned with building and supporting the competencies of workers within safety critical occupations. CS13 ('Developing Ship-Handling Skills Using Virtual Reality Data'), for instance, developed an untethered virtual reality training environment for mariners to develop skills in ship handling and navigation. The promise of the system is to deliver this kind of training remotely, without the need to visit dedicated simulation training facilities, offering both cost savings and more routine opportunities for developing ship handling skills. The intervention then offers the opportunity for embedding support for mariner safety-relevant skills and competencies beyond the usual training schedules.



The approach taken in CS26 ('Transforming Operational Risk Assessments to Better Inform Personnel') is slightly different. It is also concerned with professional competencies, in this case the ability of engineers and other personnel to identify safety risks effectively. The case starts from the working assumption that knowledge and experience around risk detection tends to be built up contextually by professionals over the course of their careers, rather than shared. The intervention was a machine learning and analytics platform that could be trained on existing data around safety incidents, and which could then provide advice and guidance to safety professionals around the planning of tasks. The intervention then supports the development of competencies to assist professionals to "make informed decisions that will lead to a reduction in safety related incidents" by extending, automating and analysing information around relevant prior cases. So rather than foster competencies directly, the intervention is a resource to support autonomous professional learning.

From a different direction, some cases aimed directly at ensuring compliance with standards and practices on the part of workers. CS9 ('Automating safety compliance in construction') details a pilot using a machine learning approach to detect likely safety compromising shortcuts taken by workers, based on automated analysis of safety reports in the construction. The system could then generate likely scenarios where different shortcut might occur. Whilst the pilot intervention was able to demonstrate an effective means of detecting likely areas of non-compliance, it was not designed to provide assurance of compliance. This latter step is the explicit goal CS7 ('A pilot to test whether technology can recognize safety hazards within the workplace'). The pilot developed a computer vision system to automatically detect whether airbags were deflated or missing alongside trucks in a shipping and handling bay. Non-compliance in the use of airbags whilst workers were operating at height was considered a major safety risk. The system was able to provide real-time alerts to Environment Health and Safety Teams to intervene in operations to ensure airbag compliance.

There is a clear difference in the assumptions about the nature of the safety problems between the two approaches. A competency fostering approach assumes that professionals making safety-relevant decisions would do better when offered either better opportunities for training or exposure to a wider range of relevant case-specific information. The compliance approach contrastingly assumes the inevitability of safety incidents through either shortcuts taken around tasks or non-compliance with standard operating procedures. To some extent the assumptions follow differences in sectors and roles, with the competency-approach cases based around maritime navigation and safety inspection, and the compliance-approach cases developed within construction and transportation/handling. However, CS24 ('Startup improves safety using deep learning-based computer vision') stands somewhat between the two approaches. It addresses the problem of electrical engineers making erroneous repairs to or leaving tools/waste behind when conducting work on electrical cabinets. Then pilot developed a computer vision-based system utilising deep learning models to automate error detection in photographs of completed work uploaded by engineers. The system then provided real-time

advice to engineers before they left the site. The intervention then appears to be a mixture of compliance assurance, in that errors are automatically detected and logged, and supporting competency, in that the professional judgment of the engineer is being augmented through external technical means rather than circumvented.

5.5 Evaluation and assurance of safety interventions

The Safetytech Accelerator cases are by their very nature all descriptions of either feasibility or pilot studies. As such, there is typically no data around the longer-term reliability or validity of the intervention. The promise of the technology is usually described in each case study with a short account of the potential industry-wide impact if the intervention were rolled-out and adopted. In many cases this is accompanied by testimony from challenge partners, technology developers and LRF as challenge sponsor.

Few pilots were able to demonstrate 'in situ' reduction in the number of safety incidents. CS7 ('A Pilot to Test Whether Technology Can Recognise Safety Hazards Within the Workplace') is an exception in that field testing was able to generate an overwhelming reduction in potential safety incidents. In the absence of long-term data on the safety enhancement, the pilots tend to present the success of the intervention in terms of metrics relating to the operation of the system itself. This often relates to the accuracy of detection rates of the safety-relevant phenomenon – e.g. levels of live listeria in food products (CS19 'Pilot to Speed Up Detection of Listeria in Food Production'); number of anomalies in ventilation ducts (CS8 'An Open Innovation Challenge to Transform Industrial Inspection and Improve Safety of Workers'). It can also reflect the relative speed of detection when compared to other methods. For instance, CS20 ('Pioneering WiFi-Based Fire Detection Technology') reported the detection of temperature changes with cargo containers within 'seconds' as opposed to the much longer time required with technologies based around detecting smoke particles. Where the intervention sat upstream from the safety problem itself, success was reported in measures such as the speed of processing for the anonymisation of documents compared to manual measures (CS5 'Anonymising and Desensitising Health and Safety Data') or the quality of the scenarios produced on the basis of machine-learning for datasets (CS21 'Providing Construction Safety Guidance Using Visual Observations and Historical Reports'). Finally, CS23 ('Start up deploys wireless sensors across ship's cargo hold to predict fire') did install the technology on a vessel, but only tested the reliability of its wireless capabilities and economic savings as a function of spread and cost of individual sensor units.

The challenges to demonstrating the success of the intervention are then based on the extent of the journey from the kind of pilot testing adopted to full implementation. There are case studies that describe the results of laboratory testing as a stage in a planned journey to implementation, such as CS19 ('Pilot to Speed Up Detection of Listeria in Food Production'), although in this case the testing was originally planned in-situ but proved impossible due to COVID restrictions. CS27 ('Understanding Decision-Making on a Ship Through Sensing')



describes work done by placing sensors on the bridge of a vessel which effectively transformed it into a 'mobile laboratory'. In CS15 ('Early Fire Detection on Container Ships using E-Nose Technology') and CS20 the testing moved from laboratory conditions to testing in operational settings, using both actual and simulated conditions, including onboard vessels. This form of pilot testing involves applying laboratory type experimental procedures and conditions in settings that are progressively closer to the intended end-user sites. Effectiveness is then demonstrated by producing similar results at each progressive stage.

A different approach is taken in the pilot tests which developed the intervention in-situ and present data from either part or full implementation at the testing site with the potential for a wider roll out. CS6 ('A pilot to optimise port visits') configured an AI platform to work with data received at a port which would assist in adopting a 'just-in-time' approach to estimating the arrival of vessels. The pilot testing for CS2 ('Onboard positioning to enhance fire response times at sea') moved from a partial to a full implementation of a wearable device to locate seafarers during potentially dangerous situations (e.g. fires, overboard, falls). This is described as a step towards a decision on full rollout of the intervention across the fleet of the stakeholder partner. In both cases, the pilot involves both the initial implementation of the technology under test conditions, and its further configuration to the specific needs of the partners responsible for the test sites. Safety outcomes then depend on the reliability of the technology as it is gradually scaled up.

The route from pilot to full implementation can be conceived as a series of links in a chain of demonstrating safety-relevant results. This typically starts with findings at a remote site or a restricted part of an operational site and is then translated across the different phases of implementation. But at each link there may be subtly different external criterion, benchmark or guarantee against which the outcomes are different. These include: quality of initial training materials, professional experience, sector standards and/or industry norms; feasibility and accessibility with a range of stakeholders.

- The kind of professional experience found in CS26 is central to the development of the intervention in CS13 ('Developing Ship-Handling Skills Using Virtual Reality Data'). The virtual reality environment was developed by drawing on the knowledge held by highly experienced mariners, who were able to advise on both how far the environment captured the relevant dynamics of real-world navigation and on the acceptability of the feature of the overall simulation itself. CS18 ('Improving defect identification in pipelines') piloted an automated machine learning solution to interpret the data produced by in-line inspection tools (ILI's or 'smart pigs') scanning within pipelines. The intervention was demonstrably more sensitive in detecting pipeline abnormalities than the existing ILI analytics. However, this outcome depended upon a comparison with professional judgment by engineers of excavated pipelines, referred to as 'the truth' in the study.
- In some safety challenges, there are clear sector standards which the intervention needs to meet to demonstrate efficacy. For instance, in CS19 the identification of live listeria in the food chain with the highest possible degree of accuracy is a clear requirement for the intervention. Similarly, in CS2 ('Onboard Positioning to Enhance Fire Response Times at Sea'), it is critical that the wearable intervention must be sufficiently well-integrated with crew working practices to enhance rather than impeded their ability to reach muster points during emergencies. In CS7 the efficacy of the intervention is measured against industry standards (supported by the NSC) to push toward the overall elimination of specific safety hazards, such as injuries from falls at height. CS12 ('Communicating Accurate Ingredient Data in School Canteens') is interesting in that it seeks to go beyond existing standards and has the potential to develop new norms in terms of safety practices by catering for dispreferred dietary options as well as allergies in school canteen food consumption.
- Finally, there are case studies where the feasibility and acceptability of the intervention are critical pre-conditions to being able to demonstrate possible safety outcomes. For example, CS11 ('Using human voice to uncover mental wellbeing insights in maritime') piloted the analysis of voice recordings made by seafarers during the course of their work on a vessel as a way of detecting stress and emotion, offering the promise of reliable measures of wellbeing that might be used to design future mental health interventions. The pilot then depended upon both the consent of the seafarers to forgo their privacy concerns and the extent to which fleet managers and owners could be convinced that the voice recordings could provide meaningful indicators of wellbeing. Relatedly, CS2 also depended on the project team being able to convince seafarers as to the acceptability wearing the device 'at all times' whilst on board the vessel.

- Many of the interventions based around machine-learning models and AI systems depend on the quality of initial training materials in order to demonstrate their potential impact on safety. For example, the dataset used in CS21 was derived from a large number of historical reports on safety contraventions supplied by the HSE, which might be considered equivalent to a 'gold standard' for safety data. Similarly, in CS26 ('Transforming Operational Risk Assessments to Better Inform Personnel'), the model used was trained on large dataset of historical incident reports and training records provided by Lloyds Register. In both cases, there is an institutional guarantee implied by the data provider that underpins the initial development of the intervention. In CS26 this is also explicitly made clear through the further participation of "LR subject matter experts" who supplied further HSES domain knowledge.



6. Discussion

The mission of Safetytech Accelerator is to 'make the world safer, more efficient and sustainable through the wider adoption of tech in safety-critical industries'. Safetytech Accelerator pursues this mission by creating relationships between 'industrial problem-owners' and technology developers. The definition and implementation of 'safety' is clearly central to what Safetytech Accelerator does. In some of the Safetytech Accelerator case studies, there are clearly defined standards and regulations which define safety. CS19 ('Pilot to Speed Up Detection of Listeria in Food Production') is the clearest example, with detecting listeria within products in the food chain providing very definite parameters in which safety interventions can be evaluated. At the other end of the spectrum, there are examples of projects where what constitutes safety is being explored and redefined as part of the intervention itself. The work to develop models of the relationship between crew behaviours and vessel performance in CS27 ('Understanding Decision-Making on a Ship Through Sensing') and the provision of food product information in CS12 ('Communicating Accurate Ingredient Data in School Canteens') fall towards this end.

In the vast majority of the cases, what constitutes 'safety' is not necessarily a completely known criterion but is rather a construct that evolves during the course of the development of the intervention. This begins with the initial challenge formulation but continues to be redefined as the technology is configured in response to the challenge. For example, in CS7 ('A Pilot to Test Whether Technology Can Recognise Safety Hazards Within the Workplace') the initial challenge was to reduce workplace fatalities associated with working at height. The project then focused on ensuring that airbags were inflated as the key safety practice and configured a computer vision technology to remotely monitor and alert to incidents of deflation. Whilst the end outcome remained stable, the scope of what constituted safety shifted to include a new network of technologies and actors, including AI models and Environmental Safety team members. Moreover, the case study indicates that once implemented, the system would potentially be reconfigured to detect material defects, malfunctioning furnaces and inventory management, which augment the initially defined outcome.

Treating safety as an evolving construct within the project rather than an initial standard to be addressed allows for a richer understanding of what was achieved within each case study. For instance, the case studies that responded to the challenge of reducing cargo hold fires in vessels might be considered as finding improved technological solutions to a very old problem. But these projects also shifted the nature of the problem in ways that redefine what constitutes safety. In CS1 ('A Feasibility Study into Using Energy Harvesting Sensors to Detect Fires'), for instance, the sensors were designed to detect heat rather than smoke particles, which potentially allows for a range of other safety-relevant ways of monitoring the ambient environment of the vessel. Similarly, in CS15 ('Early Fire Detection on Container Ships using E-Nose Technology'), the 'E-nose' technology made it possible to monitor a range of smells which have a safety dimension, including chemical spills. The problem is then expanded from ensuring safety

in a defined physical environment to registering safety threats in a temperature field or an olfactory environment that was not previously routinely accessible to crews.

Some of the most interesting outcomes across the case studies are those where projects offered the possibility of defining new 'norms'. The system in CS12 ('Communicating Accurate Ingredient Data in School Canteens'), for instance, arguably does not so much solve an existing problem as help to define a new space that goes beyond existing regulatory standards in food service and consumption and provides insights on how technology can link social relationships between school, parents/carers and children in ways that might require different normative safety standards. CS11 ('Using Human Voice to Uncover Mental Wellbeing Insights in Maritime') can be considered a parallel example of exploring new norms, but in this case the relationship between voice and the 'inner world' of a seafarer's personal wellbeing. In both cases the projects offer the opportunity to develop safety-relevant criterion and standards in a space that was not previously subject to clear mapping or evaluation.

The definition of the problem that underpins the challenge has a differential impact on the projects. Framing the challenge as something that might potentially be 'solved' (as with CS5 'Anonymising and Desensitising Health and Safety Data') creates a very different problem space for developers than those which required better 'management' (as in CS22 'Reducing Risk in Ports Using AI-Based Analytics'). Similarly, the 'prevention' approach seen in CS4 ('A Pilot to Explore Robotic Hull Cleaning and Inspection Technology') establishes different operational requirements to the 'reinvention' approach taken in CS14 ('Discover the Safetytech Solution that can Assess if a Crew is Fit for Duty'). These differences can be mapped onto a Hierarchy of Controls methodology. Solutions here suggest either the 'elimination' of the hazard or its isolation from employees through 'engineering controls'. The two case studies exploring robotic inspection solutions (CS4 and CS16 'Exploring Robotic Solutions to Remove the Need for Humans to Conduct Inspections in Confined Spaces') are good examples of the latter. Many of the case studies can be considered as developing 'substitution' (changing processes) or 'administrative controls' (changing ways of working) safeguards. For instance, CS18 ('Improving Defect Identification in Pipelines') and CS19 both offer new ways to improve on existing practices (i.e. tunnel inspection, listeria detection), whilst CS9 ('Automating Safety Compliance in Construction') and CS24 ('Startup Improves Safety Using Deep Learning-Based Computer Vision') develop a technological route to new ways of conducting safety inspections and evaluating compliance.

The Hierarchy of Controls approach (see Burk, 2016) is useful in relation to a number of the case studies, but as with the conceptualisation of safety, it does not necessarily afford a broader understanding of the nature of the innovations that are described. Actor-Network Theory (ANT) is more useful in this regard (Latour, 2005, 1987; Brown, 2011). A key principle of ANT is that innovation typically displaces the nature of a problem from one domain to another. For example, shifting the problem of safety on construction sites away from the sites themselves and towards historical records



forms the basis of CS21 ('Providing Construction Safety Guidance Using Visual Observations and Historical Reports'). We can think of displacement in terms of being upstream or downstream from the ultimate safety problem. Several of the cases (e.g. CS5 'Anonymising and Desensitising Health and Safety Data', CS15; CS27 'Understanding Decision-Making on a Ship Through Sensing') demonstrate that shifting the focus to precursors of the safety problem, or to an entirely different domain (such as existing records) offers a new route to the desired safety outcomes. One of the potential pieces of learning here is that displacing the intervention upstream of a given safety issue may create a more tractable and soluble problem space (although it may require further innovation to establish downstream links).

The transformation of one problem into another and the reinvention of a new kind of safety issue can also be understood from an ANT perspective. Transformation occurs when the 'interests' of different actors – which can be social, economic, relational etc – are aligned in such a way that actors see that the common adoption a particular practice or technology is the best means of address. For example, in CS6 ('A Pilot to Optimise Port Visits for Ships'), the port authorities wanted to maintain an optimal system for the arrival and servicing of vessels, whilst ship-owners and crews wanted to avoid delays. These two sets of interests were not necessarily aligned through the previous 'first come, first serve' approach, which increased the safety risks involved in vessels racing towards the port. The AI 'just in time system' then potentially align the interests of actors, despite the costs of adoption and compliance. Similarly, with C12, the actors have very different interests. Schools want to provide food in canteens that is compliant with regulations, parents want to ensure that their children's food preferences and nutritional needs are met with minimum disruption, and children want to eat food that they like. The system then aligns these needs by obliging all actors to adopt its functions (this is sometime called creating an 'obligatory point of passage' in ANT, see Callon, 1986). CS12 ('Communicating Accurate Ingredient Data in School Canteens') is also an example of reinventing safety problem. As with CS14 ('Discover the Safetytech Solution that can Assess if a Crew is Fit for Duty'), the project creates a new domain for safety norms and practices.

Displacement, transformation and reinvention are all commonly referred to within ANT as 'translation'. This term is generally used to refer to shifting meaning across two or more languages. In ANT it is used to analyse how interests and activities are drawn into a common 'network' of relationships, including social and technical (Latour, 1987). The safeguards within the Hierarchy of Controls methodology can be understood as instances of different translation mechanisms. One of the advantages of the ANT approach is that is very concerned with what happens when translation 'fails', particularly once the networks has been established. For example, in CS2 ('Onboard Positioning to Enhance Fire Response Times at Sea'), the crew of the vessel did express some privacy concerns around the wearables that were being tested, which were successfully addressed. But it might be the case that during further implementation these concerns might return for unexpected reasons, or that there is a failure of the technology which undermines

its acceptability to the crew. Whether or not this might occur, the 'weak point' in the network is the ongoing preparedness of seafarers to wear the device (or to have it activated). If this translation mechanism fails for some reason, the whole network of relations created by the safety intervention is threatened.

The case studies generally focus on the technological translations that are required for the safety intervention to work. These include ensuring interoperability between systems, the reliable flow of information and accuracy of detection rates and other analytics. But these technological translations are embedded within social translations and it is important to understand how these mechanisms work together. For example, in CS24 ('Startup Improves Safety Using Deep Learning-Based Computer Vision'), the technical translation is the AI automation of detecting safety violations in images of the work engineers have completed on electrical cabinets. The social translation is the preparedness of engineers to supply the images correctly on the basis that doing so may make it easier and faster to depart the site without having to wait for a manual inspection. The two are clearly linked. If instances of poor detection by the AI model start occurring that delay the process (e.g. false positives in safety violation) then engineers might seek 'workarounds' with the system such as manipulating images or claiming system or network errors. The social and technical network through which an intervention is delivered is then not necessarily stable over time following implementation, but rather subject to failures in translation which might occur through either internal failings or unexpected changes in wider conditions.

Ensuring the long-term stability of the safety intervention network then requires an understanding of the interests of the actors involved. This is particularly apparent in those case studies which tend towards a compliance approach. For instance, in CS9 ('Automating Safety Compliance in Construction') it is assumed that construction workers will take shortcuts in various tasks that might result in safety incidents. However, it appears to be out of scope for the project to consider why these shortcuts might occur in the first place. Construction workers may have reasons for making these decisions other than simplifying their approach to tasks. They may, for example, feel unequipped to complete the task in the specified way, or they may feel that the approach does not fit with their experience of how things should be done, or that the approach has been imposed without a sufficient rationale and offends their professional judgment. Understanding the relationship between individual interests and competences with systems of compliance seems relevant to the problem definition in these cases.

The question of how far the safety intervention should be in the support of professional competences rather than establishing compliance can be asked across all of the case studies. Several of the case studies embed professional knowledge and experience in the configuration of the technology (e.g. CS13 'Developing Ship-Handling Skills Using Virtual Reality Data'), whilst others make a knowledge base available in a way that can enhance professional judgement (e.g. CS21 'Providing Construction Safety Guidance Using Visual Observations and Historical Reports'). A related question here is around whose professionalism is being supported. For instance,



the systems developed in CS11 ('Using Human Voice to Uncover Mental Wellbeing Insights in Maritime') and CS14 to exploit acoustic and visual analytics to indicate and measure seafarer wellbeing arguably help support the judgment and practices of wellbeing professionals who may be tasked with effectively delivering mental health and wellbeing interventions. But they do so in a way that shifts the locus of judgment away from the subjects of the intervention itself – i.e. the ability of seafarers to judge and report their own lived experiences – and towards a less well-defined series of external actors.

Finally, the differences in the sectors covered by the safety interventions across the case studies would bear further analysis. There does appear to be a greater tendency to approach construction and logistics/handling from a compliance approach, with the exception of those case studies that are framed around the work of safety inspectors (CS21, CS24). The three case studies which address food production and service are remarkably varied in their approach, which is potentially a function of where they sit within the food chain (i.e. CS19 at the production end is the most solution/management focused and CS12 at the service/consumption end takes a reinvention approach). For the engineering sector case studies, the interventions tend towards the top end of the Hierarchy of Controls and as such the case studies are mostly concerned with the operational readiness and fitness for purpose of the technologies involved. Arguably, these cases might consider further the social relations that are required as part of the implementation, such as the skills required for the deployment and maintenance of the Roverclean technology in CS4 ('A Pilot to Explore Robotic Hull Cleaning and Inspection Technology') or the different inspection robots in CS16 ('Exploring Robotic Solutions to Remove the Need for Humans to Conduct Inspections in Confined Spaces').

The maritime sector continues to offer a complex set of operational conditions and safety problems. The case studies here are divided between those which have emerged from specific Safetytech Accelerator programmes and individually set challenges. Many of the case studies in this sector appear to have been conducted around initial feasibility testing or early pilots of technology rather than work close to the point of implementation. It is notable that many of the case studies involve a large number of varied stakeholders, including operating companies, vessel owners, insurers and oversight bodies, along with onboard crews and offshore staff. The outcomes of the case studies may then be related to both the complexity of the stakeholder network and the range of different interests requiring translation through the safety intervention.

7. Conclusion and guidance

The application of the case study synthesis to the 26 Safetytech Accelerator (STA) published case studies shows that there is considerable variability in terms of the details provided in the cases, but that it is nevertheless possible to generate a series of emergent themes that provide additional insights and ways of re-evaluating the legacy of the Safetytech Accelerator innovation projects. The research questions described in section 3.3 gap analysis of evidence which informed the subsequent analysis of the cases. The quantitative analysis demonstrates that the majority of the case studies had a reasonable spread of information relative to the 5 quality domains (Integrity, Transparency, Completeness, Responsibility, Format and Learning), with a cluster of 7 cases scoring relatively highly. But it is the qualitative analysis which was made possible by the comparison across the quality domains which provides the richest details. These include different ways in which the safety problem is defined across the challenge setting and the response by the technology developers; the variety of modes through which the problem is then addressed; the distinct routes to safety outcomes and how they may be positioned upstream and downstream of the ultimate safety problem; the treatment of compliance and competence; and finally the range of ways in which evaluation and assurance is handled.

On the basis of the analysis, we can see that what 'safety' is and the ways that is translated into projects follows a range of different pathways, depending how challenge responses and upstream / downstream positioning of the project. In all cases, there are insights into the cases which add value to what can be gained from considering each one individually. Follow-up work with interviewees and the Safetytech Accelerator team has been of considerable benefit in adding to our understanding of the contextual details around the projects.

It should be borne in mind that what constitutes 'quality' in relation to case studies is a matter for debate and reflection. The quality assessment process we have followed was necessarily constrained by the range of different objectives that Safetytech Accelerator sought in constructing the original case study materials. Our focus has therefore been on extracting evidence and value that arises from the comparison of the cases, rather than on assessing them individually on their own merits.

Based on the case study synthesis, we have identified the following as guidance and recommendations around the external presentation of case study reports and on the gathering and curation of evidence from funded research to support future value generation.

We note that the Safetytech Accelerator case study materials were produced for an external audience and not as complete reports on the innovation cases. The recommendations here are made to indicate areas where additional information would be useful for further understanding the projects, both individually and collectively. We have also included a summary graphic in the appendices which captures some of these points.



7.1 External presentation of case studies

1. The process for challenge formulation can be further described, where commercial interests allow. Specifying the contribution and requirements which each partner brings to the process would enable better understanding of how the challenge evolved through negotiation. It would also provide a basis for developing a narrative around the extent to which the proposed innovation solutions and the outcomes of the feasibility or pilot studies helped to clarify or shape the needs of the stakeholders.
2. Providing the recent background to the problem and the why the safety challenge has persisted despite previous efforts to address it – i.e. the current state of play in the sector in question – would help to better understand why particular approaches have been taken to the problem. In particular, it would be useful in assessing whether the innovation is evolutionary or revolutionary.
3. Problems may be addressed through different modes. A typology for the different potential ways of addressing the problem (e.g. solving, managing, preventing, reinventing) may help in comparisons across a body of related innovation projects.
4. The level of resourcing provided is clearly a crucial variable for calculating immediate return on investment. But value may also be estimated in relation to the opportunities provided to a project for iterations in the overall innovation of a product/technology and route to market (i.e. proof of concept, initial feasibility).
5. Specification of the underpinning theory is required, particularly in cases where the intervention is exploratory or where the technology operates in a diagnostic way. The conceptual relationship between the measurable indicator and the state to be estimated should be sufficiently well described.
6. Where possible, links to full technical details on trials and pilot work could be provided in order to provide the background to the case studies. If there are commercial interests at stake in the project, sufficient details to allow for a clear understanding of main effects and outcomes should be included in the case study narrative.
7. Specification of the criterion used to measure the 'safety outcome' or impact on safety practices is required. This may take a variety of different forms, including predicted estimates based on the outcomes of the project trials and pilots. An indication of the extent to which the impact is direct or 'upstream' in relation to the intended safety problem would also assist in comparing case studies.

8. In cases where there were negotiations around the acceptability of the technology within the field site, a brief narrative around how these negotiations were concluded/ issues resolved would allow for better estimates around likelihood of successful implementation of the technology.
9. Feedback from all challenger partners should be gathered and reported to enable contextual understanding of the innovation testing. This ideally should include feedback from participants within the trials themselves, particularly in cases where assessment of the feasibility and acceptability of the intervention is critical.
10. Consideration should be given in case study narratives to the challenges and broader opportunities in 'scaling up' as part of describing the potential sustainability of project outcomes.
11. Knowledge of what happened following conclusion of the feasibility / pilot work would help to establish what role the project played in the evolutions/iterations of the technology. Where possible, updating case studies with further information about implementation would add value to the project.
12. Follow-up interviews to establish the impact of the project on the safetytech (or other) markets would help to judge the long-term feasibility of interventions. It would also assist in clarifying areas of good practice around the commissioning and support of projects.
13. Potential broader impacts on the network of stakeholders also worth considering and detailing, particularly in case of disruptive innovation. These impacts have an economic dimension but also value in relation to ways of working and social/cultural safety-relevant dimensions.

7.2 Building on case study synthesis

Transforming professional knowledge

Funded research and innovation work may be intended to directly address challenges in a given sector. But it can also help to redefine how problems are conceptualised and formulated. For example, in relation to safety infrastructures, research and innovation projects may contribute to redefining what 'safety' is – what it has been, what it might be, the various external norms and standards in play. Case study reports can be analysed in terms of how far their outcomes have enhanced or supported existing professional knowledge or challenged and transformed such knowledge. The ways in which professional knowledge was engaged with and embedded within the structure of the project itself may also be considered, particularly in relation to work that uses AI to formalise or automate knowledge practices.



Recognising case studies as a legitimate form of evidence

Funders and programme designers should treat case studies as a complementary form of evidence, particularly in early-stage innovation contexts where formal evaluations may be premature or impractical. This requires developing frameworks for assessing narrative credibility, contextual relevance, and implementation trajectory. This includes reflective reporting that captures not just outcomes but processes, challenges, and adaptations. Case studies have the potential to capture how problems become transformed from the perspective of multiple stakeholders. However, commercial sensitivity and broader ongoing developments in sectors and fields may complicate how these perspectives are represented.

Support longitudinal and iterative learning

Innovation rarely unfolds in a single phase. It may be an iterative process that results in an evolution of the intervention, or a radical / disruptive approach that restructures the nature of the underlying problem and value chains. Funders should enable follow-on studies that revisit case study interventions over time. Furthermore, documentation of implementation pathways, including stakeholder engagement, technical refinement, and contextual adaptation, should be encouraged.

Tailor evaluation metrics to innovation stage

Rather than applying uniform success markers, evaluators may adapt more targeted and appropriate criterion for each case study. For example:

- Use proximal metrics (detection accuracy, stakeholder acceptability and so on) for upstream interventions.
- Reserve impact metrics (e.g. incident reduction) for downstream or scaled implementations.
- Consider process indicators (such as stakeholder alignment, feasibility testing) as valid signs of progress.

The question of 'what does success look like for the projects and for the programmes they sit within?' might be asked relative to different points within innovation cycle. Is short-term impact with specified stakeholder group the priority or is longer-term shaping of safety practices and markets also accomplished?

Foster cross-sectoral and interdisciplinary collaboration

Many safety challenges are complex and multifaceted. Innovation programmes should facilitate collaboration between technical developers, domain experts, and end-users. The role and experience of the funder in brokering challenges can be made central to the narrative around addressing challenges, notably in relation to mediating between diverse groups of professionals and other stakeholders. The crosscutting aspects of projects with the broader range of initiatives and programmes supported by funder may usefully be highlighted. This might suggest a greater role for 'collaboration and support activities' alongside directly funded research and innovation work.



References

- Brown, S.D. (2011). Actor-Network Theory. In Key Concepts in Critical Management Studies, eds. M. Tadjewski, P. Maclaran, E. Parsons and M. Parker. London: Sage, pp. 24–28.
- Callon, M. (1986). Some Elements of a Sociology of Translation: Domestication of the Scallops and the Fishermen of St Brieuc Bay. *The Sociological Review*, 32(1_suppl), 196–233. <https://doi.org/10.1111/j.1467-954X.1984.tb00113.x>
- Hardoon, D., South, J., Southby, K., Freeman, C., Bagnall, A.-M., Pennington, A., & Corcoran, R. (2021). A guide to synthesising case studies. What Works Centre for Wellbeing. <https://whatworkswellbeing.org/wp-content/uploads/2021/04/Guide-to-synthesising-case-studies-2021-FINAL-1.pdf>
- Latour, B. (1987). *Science in Action: How to follow scientists and engineers through society*. Cambridge, Mass: Harvard University Press.
- Latour, B. (2005). *Reassembling the Social: An introduction to Actor-Network Theory*. Oxford: Oxford University Press.
- Mills, A.J., Durepos, G., & Wiebe, E. (2010). Case-to-case synthesis. In *Encyclopaedia of case study research*. London: Sage, pp. 130–134. <https://doi.org/10.4135/9781412957397.n45>
- South, J., Southby, K., Freeman, C., Bagnall, A.-M., Pennington, A., & Corcoran, R. (2024). Synthesising Practice-Based Case Study Evidence from Community Interventions: Development of a Method. *International Journal of Qualitative Methods*, 23. <https://doi.org/10.1177/16094069241276964>
- Thomas, J., & Harden, A. (2008). Methods for the thematic synthesis of qualitative research in systematic reviews. *BMC Medical Research Methodology*, 8(45). <https://doi.org/10.1186/1471-2288-8-45>



Appendix A. List of Safetytech Accelerator case studies

#	Name
CS1*	A Feasibility Study into Using Energy Harvesting Sensors to Detect Fires
CS2	Case Study: Onboard Positioning to Enhance Fire Response Times at Sea
CS3	A Pilot to Explore Automated Data Collection From Ship Machinery
CS4	A Pilot to Explore Robotic Hull Cleaning and Inspection Technology
CS5	Anonymising and Desensitising Health and Safety Data
CS6	Case Study: A Pilot to Optimise Port Visits for Ships
CS7**	Case Study: A Pilot to Test Whether Technology Can Recognise Safety Hazards Within the Workplace
CS8	Case Study: An Open Innovation Challenge to Transform Industrial Inspection and Improve Safety of Workers
CS9	Case Study: Automating Safety Compliance in Construction
CS10**	Case Study: Mitigating Work at Height Risk with Computer Vision
CS11	Case Study: Using Human Voice to Uncover Mental Wellbeing Insights in Maritime
CS12	Communicating Accurate Ingredient Data in School Canteens
CS13	Developing Ship-Handling Skills Using Virtual Reality Data
CS14	Discover the Safetytech Solution that can Assess if a Crew is Fit for Duty
CS15	Case Study: Early Fire Detection on Container Ships using E-Nose Technology
CS16	Exploring Robotic Solutions to Remove the Need for Humans to Conduct Inspections in Confined Spaces
CS17	How Can We Enable Engineers Working with Safety Critical Equipment to Use Predictive Analytics?
CS18	Improving Defect Identification in Pipelines
CS19	Pilot to Speed Up Detection of Listeria in Food Production
CS20*	Case Study: Pioneering WiFi-Based Fire Detection Technology with University of New South Wales (UNSW) and Ginigai (formerly Envision)
CS21	Providing Construction Safety Guidance Using Visual Observations and Historical Reports
CS22	Reducing Risk in Ports Using AI-Based Analytics
CS23	Startup Deploys Wireless Sensors Across Ship's Cargo Hold to Predict Fire
CS24	Startup Improves Safety Using Deep Learning-Based Computer Vision
CS25	Transforming Food Safety, Assurance and Compliance Using AI Technology
CS26	Transforming Operational Risk Assessments to Better Inform Personnel
CS27	Understanding Decision-Making on a Ship Through Sensing

* Denotes feasibility study

** Denotes a short and longer version of the same project



Appendix B. Five steps to better case studies

