ASSURING AUTONOMY
INTERNATIONAL PROGRAMME

DEMONSTRATOR PROJECT

FINAL REPORT

CSI:Cobot Confident Safety Integration for Cobots

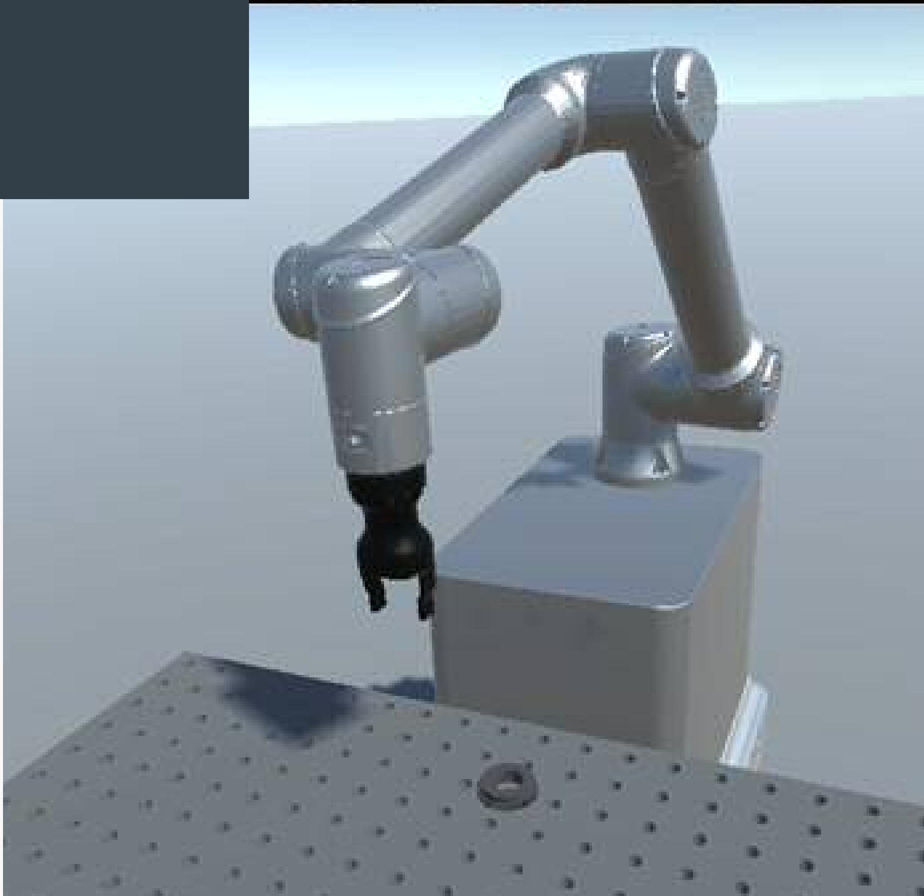

JULY 2022

Confident Safety Integration for Cobots (CSI:Cobot)

Technical Report

22th July 2022

Abstract

A major opportunity for industrial manufacturing is increased collaboration between humans and robots, and the resulting combination of automation and human skill. However, the combination of humans and robots working in close proximity raises complex safety issues, which are a critical barrier to adoption. The CSI:Cobot project aims to demonstrate how research in the fields of safety engineering, machine learning, and cybersecurity can be applied to human-robot collaborative processes to address safety concerns and improve confidence. Employing two case studies, specific emphasis is placed on digital twins for safety analysis, machine learning for vision-based proximity detection, synthesis of safety controllers, testing approaches for analysis of hazards, and security policy, user authentication, and intrusion detection. This report summarises the approaches and findings of the project, which we have undertaken in collaboration with regulators, standards communities, and industrial stakeholders, to ensure relevance to future industrial application.

Contents

Executive Summary

The emergence of ‘collaborative robots’ promises to transform the manufacturing sector, enabling humans and robots to work together in shared spaces and physically interact to maximise the benefits of both manual and robotic processes. Whereas traditional, non-collaborative, processes rely on segregation of robots and workers to ensure safety, collaborative working introduces complex challenges around the monitoring and control of systems and processes; where people and robots operate in share environments, and where physical interaction is a possibility, it becomes much harder to guard against potential hazards. Additional safety considerations are therefore required before robots can be deployed alongside people in industrial processes.

This project stemmed from a series of engagements (workshops, interviews, site visits) with industrial organisations with interests in collaborative robots, which identified safety as the greatest barrier to adoption. The lack of best practice examples in collaborative safety resulting in companies falling back on traditional segregational approaches, including physical barriers and cages. The following were identified as specific challenges, which form the objectives of this project:

1. To understand the requirements, risks, and safety needs of current collaborative robotics processes in industry, and identify the safety barriers to further adoption.
2. To provide new volumetric sensing and tracking techniques able to predict potential collisions between humans and robots before they occur
3. To provide a detailed security threat analysis, and identify intrusion detection methods, for collaborative robotic systems. (Insufficient security has distinct safety implications, yet this is an area currently overlooked in industrial applications)
4. To provide rigorous safety hazard analysis and testing approaches for cobot systems in order to identify hazards and determine whether they can occur in the implemented robot software
5. To provide evidence-based modelling, analysis and controller synthesis techniques for dynamic operation-mode switching, in response to available sensing, security, and situational data (or lack thereof)
6. To provide best-practice demonstrators and boost confidence in safe operation of collaborative robots in manufacturing processes
7. To engage with the regulatory community to ensure these methods are valid in terms of the changing safety standards, and to help disseminate the project outputs among potential users

To address these challenges, we undertook to apply research methods from the areas of sensing (obj. 2), security (obj. 3), analysis (obj. 4), and planning (obj. 5), to industrial demonstrators (obj. 6), under the guidance and feedback of industrial and regulatory partners (obj. 1 & 7). The work was undertaken in two phases, in relation to two case studies: in Phase 1 we consider an existing, constrained, industrial process involving a single robot mounted in a fixed location, operating in a confined space by a single user, on an isolated network; in Phase 2 we consider the more complex case of a future (planned) process involving a mobile robot operating in a less confined environment, on a more integrated network involving wireless communications, and where environmental conditions are more subject to change. In response to Covid-19 restrictions, which prevented physical collaboration, access to facilities and stakeholders, work was largely conducted in the digital domain. Bespoke digital twinning tools were developed to support research activities to enable the design, test, deployment, and validation of methods in both virtual and physical (when available) worlds. These activities have highlighted the value of digital twins for safety applications, and offer an exciting avenue for further research.

The key technical outcomes of the project can be summarised as:

- Development of a modular digital twinning framework for robotics, including tools for safety analysis and visualisation
- An approach to visual detection and tracking of humans and robots using Region Based Convolutional Neural Networks
- Stochastic modelling and controller synthesis methods to enable dynamic switching between safe robot operation modes
- Application of manual (STPA) and semi-automated (SASSI) techniques to analyse hazards in the system and validate safety controllers
- Threat modelling and identification of security policy/requirements for collaborative systems. Development of methods for intrusion detection and continuous user authentication.
- Virtual and physical demonstrations of collaborative robot safety techniques

Additionally, the project delivered the following benefits:

- Interviews and workshops with industrial partners to identify safety challenges, understand attitudes, co-design research, and build confidence in novel safety approaches
- Engagement with the Health and Safety Executive (the regulator for industrial robot safety) on issues of collaborative robot safety. This resulted in increased understanding of research practice within the HSE, sharing of information to support future regulatory decision making, and shaping of research developments toward regulatory approval.
- A 3-day Manufacturing Robotics Challenge, including training in safety assurance and robotics, for 38 early-career researchers based in 11 countries.
- Input to international standards via feedback through the ADRA standardisation community
- 13+ peer-reviewed international publications (4 journal papers, 8 conference papers, 1 book chapter)
- 12+ conference, workshop, & seminar presentations
- 11+ funded projects (~£4.9M)
- 7 AAIP Body of Knowledge contributions
- 2 public machine vision datasets and an annotation tool
- Content for future taught courses on safety and security

The remainder of this report provides further details on the case studies and research activities undertaken in the project.

Case Studies

The following outline the two case studies referred to throughout the remainder of this report.

Phase 1

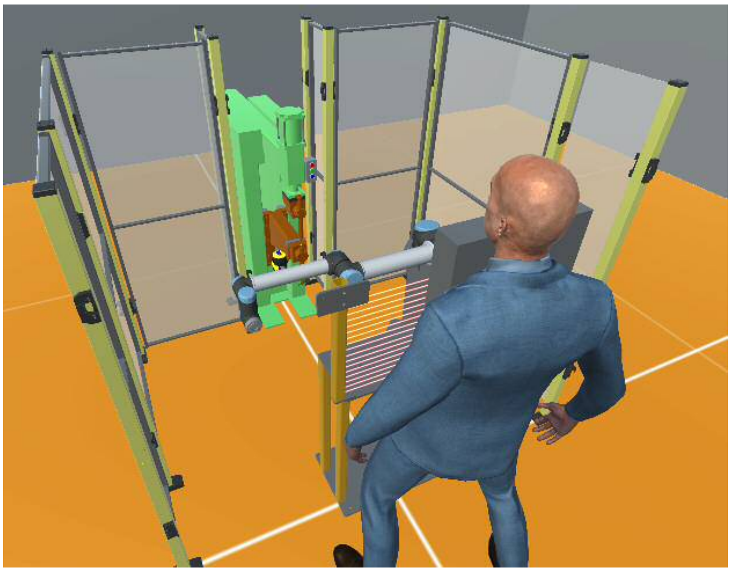

Figure 1: digital twin visualisation of the spot-welding case study

In this process the robot is separated from the operator by a workbench and a small light barrier (seen just in front of the worker in Fig 1.). The human assembles a component and passes it through the light barrier, placing it on a shelf within reach of the robot. Once the operator has withdrawn, and the light barrier is no longer broken, the robot picks up the assembly, performs a weld, returns the assembly to the shelf, and withdraws. The operator then reaches through the light barrier to retrieve the welded assembly, and replaces it with another. The process then repeats.

In addition to the light barrier there is a laser rangefinder at shin-height by the base of the welding machine that detects anyone who has entered the workcell.

The process includes several safety modes, triggered by different sensors:

1. A 'safety-rated monitored stop' is put in place at the point the worker breaks the light barrier to place or retrieve parts from within the robot's workspace. This immediately pauses robot activity. The light barrier is muted whilst the robot is at the spot welder, enabling the operator to place/retrieve parts without disrupting the robot when it is at a safe distance.
2. 'Speed and separation monitoring' is activated if someone enters the portion of the workcell away from the robot and welding machine. This limits the speed of the robot.
3. A 'protective stop' (non-collaborative and requiring a full system reset) is activated if someone approaches too closely to the robot or welder

Phase 2

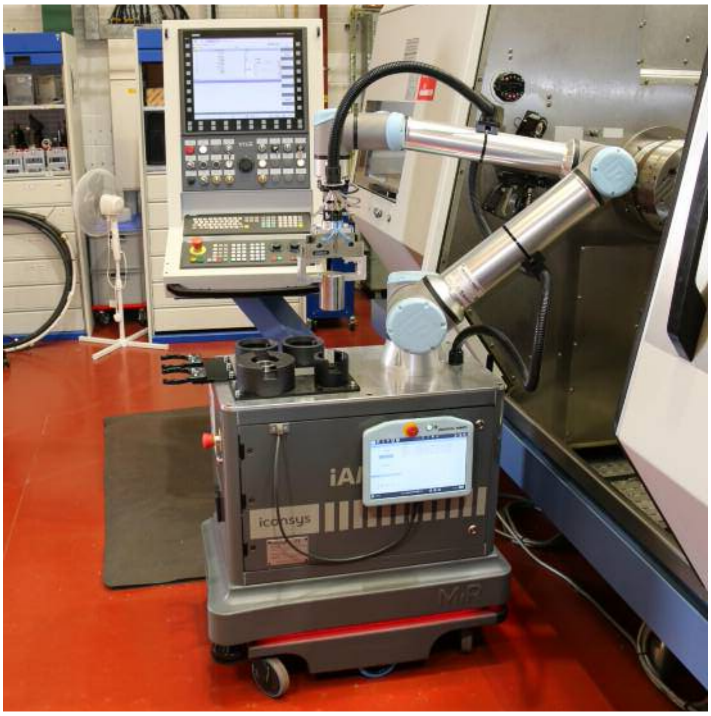

Figure 2: the iAM-R performing machine tending

In this process an iAM-R mobile collaborative robot tends to several machines for the process of manufacturing gears. The robot is tasked with placing materials and components at stages of their manufacture into milling and tumbling machines, and retrieving the processed pieces. As machines are located at different locations, the robot is required to traverse the shop floor during operation[1].

Safety is predominantly ensured by laser rangefinders on opposite corners of the mobile base. These detect obstacles in close proximity to the robot, and cause the robot to stop if an unexpected obstacle appears within a proximity threshold. To comply with existing regulations, the arm must be stowed before movement of the mobile base.

The process raises additional complexities over the Phase 1 study, including:

- Mounting of the collaborative robot on a mobile base, thereby expanding the range of motion, and allowing the reference frame of the arm to move relative to the fixed infrastructure of the environment
- Removing all physical barriers from the workspace
- Much greater possibility for changes in the operation environment (lightning, obstacles, etc.)
- More open operating environment leading to greater potential for people to enter the workspace of the robot, and from unknown directions, including people not trained to work with/alongside the robot
- Inclusion of wifi connectivity, increasing the potential for security threats
- Interaction with multiple machines

[1] A video of the process can be found at https://youtu.be/ng7VCN2Bj24

Digital Twinning

Digital twins provide levels of insight, access, and transparency to processes that simulation and process monitoring tools alone cannot. Beginning with a simple twin of an industrial process in Phase 1, we have developed a sophisticated *digital twin framework* (DTF) to support the increasing complexity of our research needs which have required the incorporation of additional tools. This DTF has provided an invaluable tool for research and collaboration, whilst also opening up new avenues to explore digital twins as tools for safety assurance.

Phase 1

The initial objective of this task was to develop a simple twin of the Phase 1 case study that would enable collaborators to develop methods remotely, and then deploy them on the physical assets for testing. This would include the proposed safety controller (see Planning) and exploratory and semi-automated analysis regimes (see Analysis). The digital twin (shown in Fig 1.) was initially developed in Unity-3D (a commercially available game-engine) as a preexisting tool for physics simulation, with some existing hardware and software integrations.

As a result of COVID-19, which prevented access to physical systems, work on the digital twin was expanded to support the more sophisticated requirements of longer-term remote research. This required the implementation of a library of sensors, robotics and human models to provide a basis for independent, simulation based, analysis. These models had to be carefully curated to present the same interface as the physical system in future demonstrations. As a result, a greater emphasis was placed on the system for communication between models, data logging, ground truth data generation, database storage and reconfiguration from external analysis tools. Further development of the software focused on containerising the deployment of our API/tools implementing existing protocols (i.e SQL,TCP, HTTP, JSON) into ‘services’ to allow them to be interacted with more portably. This involved close collaboration with the other work-packages to ensure a testable process interface was maintained as new tools were introduced.

This emphasis on creating a more sophisticated process interface, better able to represent the spot welding process and its components as individual digital twins, provided a solid basis for the deployment of a synthesised safety controller (see Planning), as a means to validate its design. Simultaneously, the digital twin provided a scalable data interface that influenced the development of the SASSI method for semi-automated coverage analysis (see Analysis). Together, this exercise validated the use of digital twins as a means for controller deployment and testing, as well as a valuable asset in the field of configuration-lead exploratory testing. This was shown by successfully deploying the synthesised safety controller and validating its interactions with devices and operators within the digital twin[2]. Using this controller, we have been able to simulate a range of use cases whilst providing the necessary data infrastructure for our experimental validation tools to quantitatively analyse its safety.

[2] See https://doi.org/10.1016/j.scico.2022.102809

Phase 2

The increased complexity of the Phase 2 case study prompted the re-development of our twinning tools into a more flexible DTF, which could more scalably represent unique combinations of robots, machines, operators and processes with a greater emphasis on modularity and tool independence. In this framework, each entity is represented as a single digital twin, with collections of entities connected together to represent a process. The resulting representation of individual digital twins as *modules* and *service modules* can be seen in Fig 3.

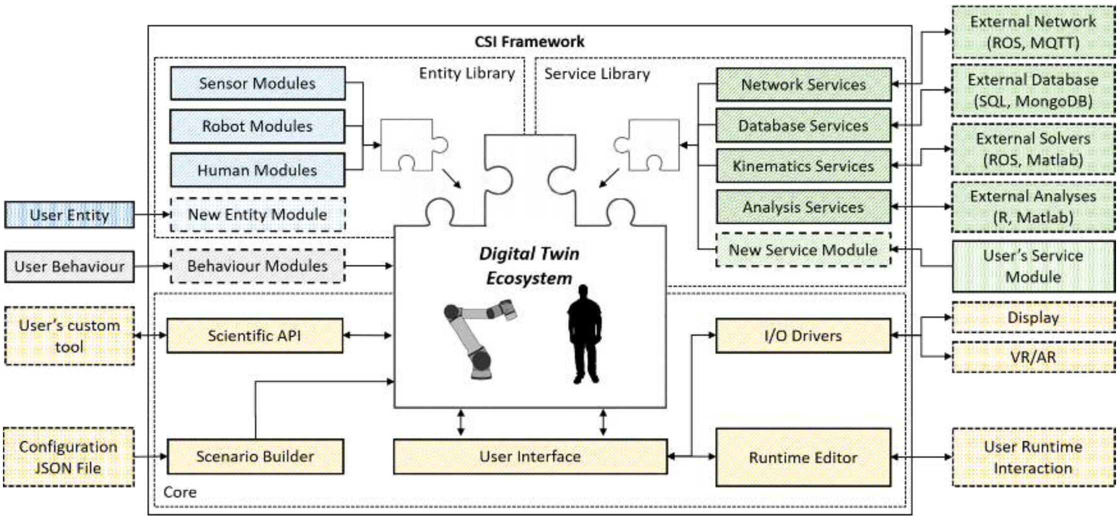

Figure 3: modular digital twinning framework highlighting the relationship between entities, services, and core functions

The framework enabled remote hosting of the UK-RAS Manufacturing Robotics Challenge in 2021, whereby an interactive scenario was used to teach 38 international participants about robotics, safety and digital-twinning technologies.

In collaboration with the Health and Safety Executive, we have developed a more formalised approach to process descriptions. Specifically, we have developed a representation of abstract safety concepts within the framework to support better communication and understanding to human operators and regulators. This lays the groundwork for more sophisticated digital-twin based safety assurance techniques. This same, formalised representation of digital twins and the services, also allowed for the development of a more sophisticated analysis pipeline as shown below, developed in collaboration with the Analysis and Planning activities.

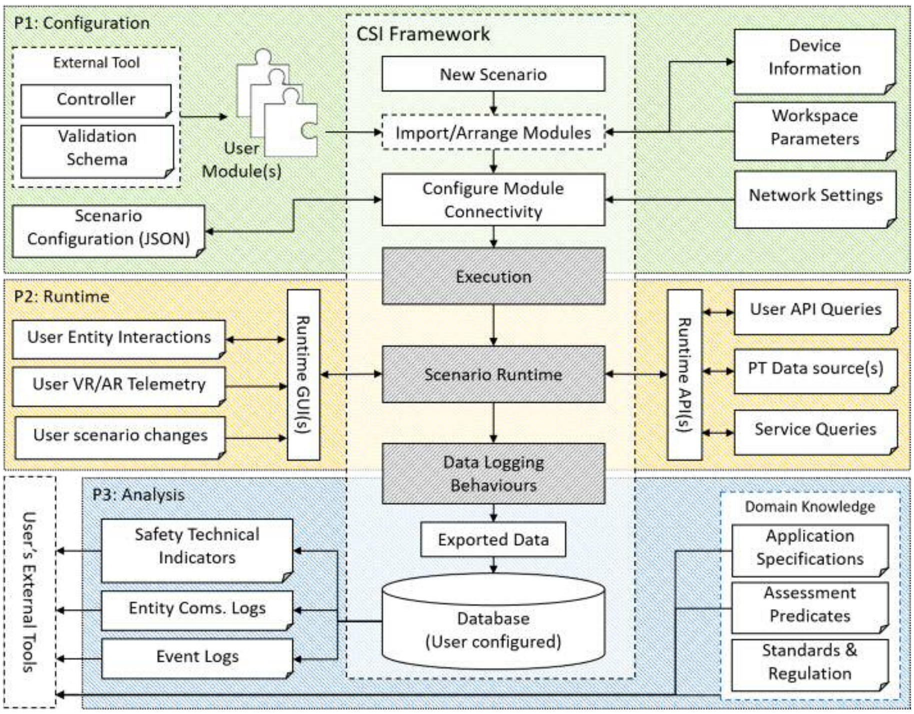

Figure 4: Digital twinning framework analysis pipeline

Here, the Phase 2 case-study was assembled as a novel arrangement of digital-twin modules and services in order to analyse the deployment of a second synthesised safety controller designed to observe the mobile manufacturing process (see Planning).

Sensing

The rapid development of machine learning techniques and advanced sensing systems in the past decade has offered great potential for assuring safe human-robot interactions in manufacturing towards a high level of autonomy and flexibility. The traditional sensing systems majorly deploy physical safety barriers, light gates or proximity sensors to avoid collisions between humans and robots, causing efficiency loss. The *Sensing* strand aims at providing intelligent sensing and autonomous decision making systems by detecting the presence of humans as well as robots and identifying their interactions in shared working cells.

Phase 1

The tasks in Phase 1 of the *Sensing* strand include the aspects of sensor selection, framework structure, methodology determination and safety criterion definition. A Kinect V2 camera was installed above the cell, looking down over the workspace. Eye2Hand calibration was carried out to obtain the relative position of the camera with respect to the robot base coordinate system. In our example, the camera recorded images of size 1920 x 1080 at 30 fps on a UR10 platform.

We processed the real data being collected from Kinect V2 to train a deep learning model. The output of the model shows the contour information of robots and humans with masks and bounding boxes. The bounding box defines an area enclosing the robot arm and operator for use in identifying potential collisions with operators; we use the intersection of bounding boxes of the robot arm, the operator, and dynamic objects in the cell as a criterion to make safety related decisions, i.e. whether to switch safety modes.

Decisions are made according to the separation of the robot and operator. Suppose one has obtained the bounding boxes of the operator and the robot, and the overlap (indicated by red rectangles in Fig. 5) between the two bounding boxes that is denoted as OVERLAP. The coordinates of these bounding boxes are transformed from the image space to the physical world and the area of the bounding boxes and the OVERLAP can therefore be calculated. The safety criterion is defined as follows and the safe threshold can be defined by the user in specific applications:

- **Safe:** If the area of OVERLAP is below the safe threshold, as shown in (a);
- **Potential:** If the area of OVERLAP is between the safe threshold and the dangerous threshold , as shown in (b);
- **Dangerous:** If the area of OVERLAP is over the dangerous threshold , as shown in (c)

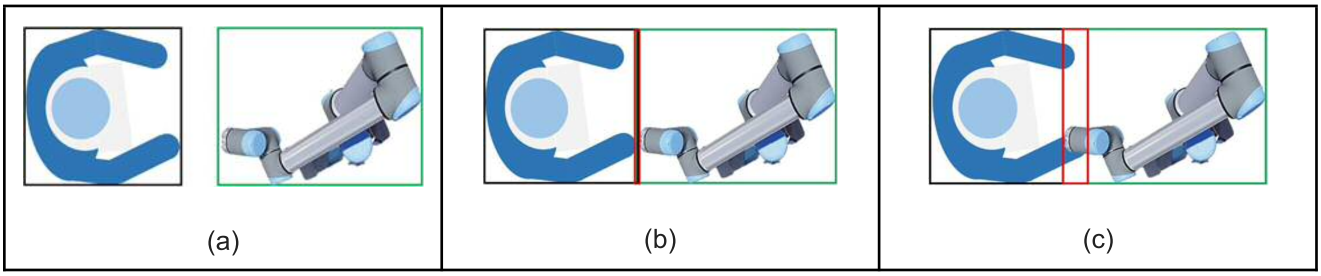

Figure 5: Safety criteria: (a) Safe; (b) Potential; (c) Dangerous. The red rectangles indicate the overlapping areas[3]

[3] from P. Wang *et al*, ‘’2.6.1 – Monitoring RAS operation’’, available online: https://www.york.ac.uk/assuring-autonomy/guidance/body-of-knowledge/implementation/2-6/2-6-1/cobots/

The proposed framework is demonstrated under a scenario consisting of a single operator and a single robot. In this example, a Kinect V2 was utilised for real data collection, the backbone of the mask R-CNN model is ResNet101-RPN. Fig. 6 provides the object detection and classification results with classification confidence.

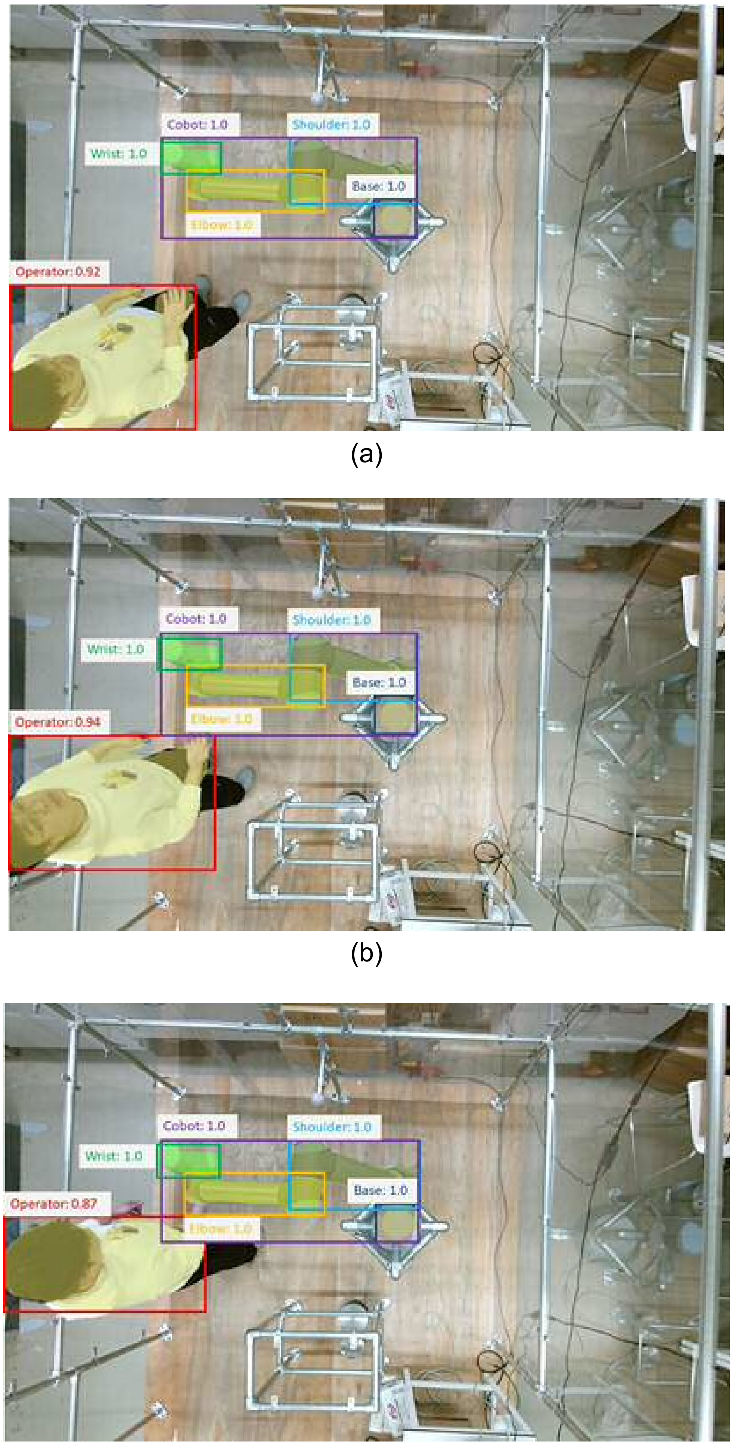

(a)

(b)

(c)

Figure 6: The detection and classification results for monitoring human-robot interactions: (a) Safe; (b) Potential; (c) Dangerous[4]

[4] From P. Wang *et al*, ''2.6.1 – Monitoring RAS operation'', available online: https://www.york.ac.uk/assuring-autonomy/guidance/body-of-knowledge/implementation/2-6/2-6-1/cobots/

Phase 2

In Phase 2, the major objectives were to improve the accuracy of the trained deep learning model, to mitigate the sim2real (simulation to real) gap, to evaluate the model performance under variable environmental conditions (e.g. different lighting levels), to identify the potential influence of environmental changes and link the developed Digital Twin with the regulatory requirements. A distinctive aspect of Phase 2 was to investigate how the developed framework and machine learning approaches can influence the regulation standards for safe human-robot interaction in manufacturing.

First, A Digital Twin was built based on Unreal Engine 4 (UE4) as a digital representation of human-robot interactions. With the help of the Digital Twin, diverse synthetic data with accurate annotation information was obtained for model training so as to bridge the gap between simulated and real systems. Unlike the physical camera mounted on the top of the robot in the real system, its virtual counterpart has more flexibility in position, orientations, and views. Hence, the Digital Twin offers the benefits of data augmentation, adjustable light conditions, randomised background, etc. The communication system was built with the Robotics Operating System (ROS) to achieve a transmission of robot poses between the virtual and real robot. Consequently, the digital robot can match the pose of the real robot. Operators can also control the poses of the robot and the movement of the digital camera through ROS. In addition to various synthetic datasets, different lighting levels were created for experimental validation of the model performance.

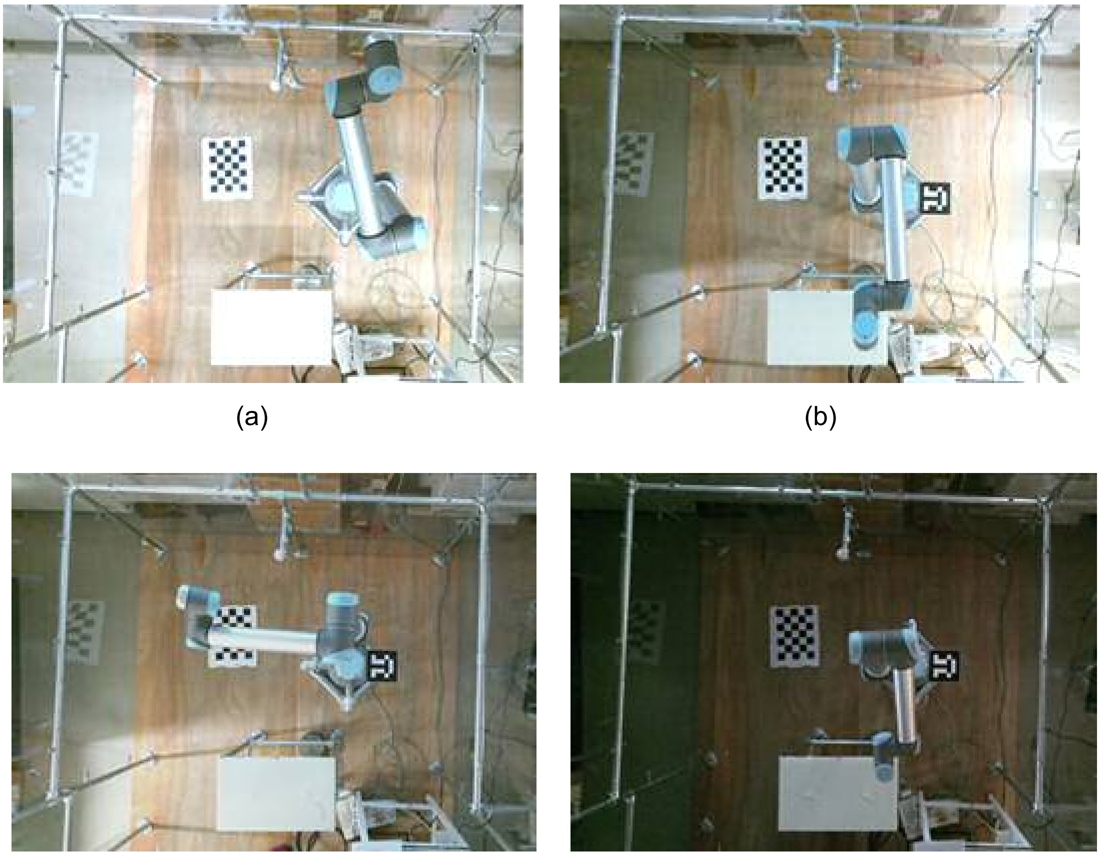
(a)

(b)

(c)

(d)

Figure 7: Different environmental conditions. (a) Full-light; (b) Semi-light; (c) Semi-dark; (d) Dark[5]

[5] From P. Wang *et al*, ‘’2.6.1 – Monitoring RAS operation’’, available online: https://www.york.ac.uk/assuring-autonomy/guidance/body-of-knowledge/implementation/2-6/2-6-1/cobots/

The results demonstrate that when the lighting condition is satisfactory (i.e. full-light, semi-light), the trained model performs constantly with high accuracy of no less than 90% in AP_{75} (Average Precision at IoU = 0.75, here IoU means Intersection over Union) indicating a good potential for ensuring safe human-robot collaborations in dynamic industrial environments. In the meantime, it should be noted that the performance of the trained model can be slightly negatively impacted by the illumination level, for example, when the illumination level drops to semi-dark or dark. This indicates that the poor lighting condition or sudden changes in the environmental conditions may result in less accurate sensing and thus unsafe operation.

Planning

The aim of the *Planning* strand during phase 1 of the CSI:Cobot project was twofold. First, devising a stochastic modelling method to achieve dynamic switching between safe robot operation modes. Second, developing a controller synthesis method using the stochastic model of the previous task and providing evidence that the controller operates as intended. The second phase of the project introduced the element of mobility to the robot, requiring the robot's as well as the environment's modelling extension. An extension of the controller synthesis method was also necessary in order to generate controllers achieving user-defined trade-offs between risk and performance of the robot (in terms of progress with the manufacturing process it contributes to).

The following sections provide brief overviews on how the above objectives were achieved during the two project phases, the lessons learned from this work, and further information on the Planning strand's impact.

Phase 1

Phase 1 of Planning provided guidance on the analysis, design, synthesis, and assurance of safety controllers for use in human-robot collaboration (HRC) settings. An HRC setting typically comprises a mobile or stationary robot (the Cobot) collaborating with one or more human operators on shared repetitive tasks (the Process). The goal is to combine the capabilities of humans and machines in order to improve quality and reduce cost. An automatic safety controller is used to improve occupational safety for the tasks performed in this work cell, and is responsible for handling critical events.

Our proposed integrated synthesis, verification, and test approach (depicted Fig. 8) is informed by the process, risk analysis, and relevant safety regulations for the target application. Controllers are selected from a design space of feasible controllers according to a set of optimality criteria, are formally verified against correctness criteria, and are translated into executable code and tested in a digital twin. The resulting controller can detect the occurrence of hazards, move the process into a safe state, and, under certain circumstances, return the process to an operational state from which it can resume its original task.

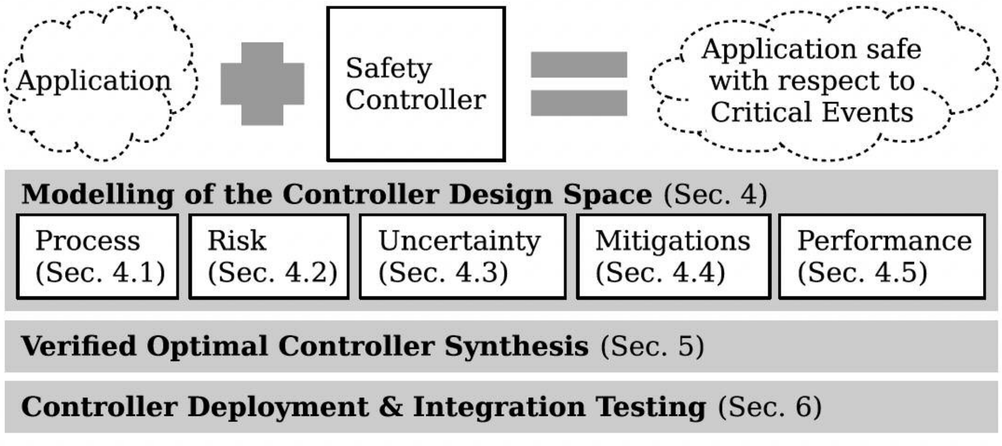
Application
Safety Controller
Application safe with respect to Critical Events
Modelling of the Controller Design Space (Sec. 4)
Process (Sec. 4.1)
Risk (Sec. 4.2)
Uncertainty (Sec. 4.3)
Mitigations (Sec. 4.4)
Performance (Sec. 4.5)
Verified Optimal Controller Synthesis (Sec. 5)
Controller Deployment & Integration Testing (Sec. 6)

Figure 8: Stages of the proposed safety controller design method[6]

[6] From M Gleirscher et al, Verified synthesis of optimal safety controllers for human-robot collaboration, *Science of Computer Programming* **218**, 2022)

The effectiveness of our software engineering approach is demonstrated through a case study involving the development of a safety controller for a manufacturing work cell equipped with a collaborative robot. More details regarding the evaluation, and our approach in general, can be found in our *Science of Computer Programming* article '*Verified synthesis of optimal safety controllers for human-robot collaboration'* [7].

Phase 2

The goal of phase 2 was to extend phase 1 of the project to a more complex industrial case study, which introduces the element of mobility to the collaborative manipulator. The resulting techniques, that ensured the safety of the human operators in the new case study, have been developed in the context of the manufacturing domain; however, it is envisaged that they can be applied to other domains with similar characteristics.

The new element of cobot mobility introduced additional safety concerns that increased the complexity of developing safety controllers. Avoiding collisions caused by cobot movements inside the work cell was of primary focus, and has been encapsulated in the updated safety controller development technique.

Our proposed approach (shown in the diagram below) is divided into three stages: a) hazard identification, b) stochastic modelling, and c) synthesis. During the first stage, an UML activity diagram and a floorplan graph are used to identify potential hazards in the process. This leads to the annotation of the UML diagram with risks and mitigation strategies for the identified hazards (e.g., slowing down the cobot's speed to avoid collision with the human operator. The second stage synthesises a discrete-time Markov chain (DTMC) model based on the annotated UML diagram of the previous stage, and to formalise the requirements of the process. Finally, the third stage employs probabilistic model checking and applies an exhaustive search over the discretised parameter space of the model to synthesise safety controllers for the industrial process.

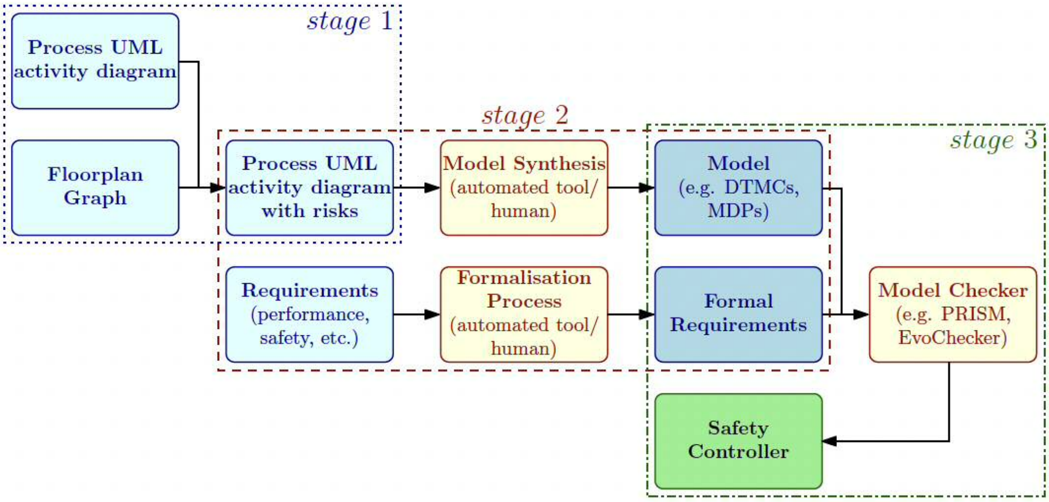
stage 1
Process UML activity diagram
Floorplan Graph
stage 2
Process UML activity diagram with risks
Model Synthesis (automated tool/ human)
Requirements (performance, safety, etc.)
Formalisation Process (automated tool/ human)
stage 3
Model (e.g. DTMCs, MDPs)
Formal Requirements
Model Checker (e.g. PRISM, EvoChecker)
Safety Controller

Figure 9: Stages of the approach for the synthesis of mobile-cobot safety controllers[8]

[7] Available at https://doi.org/10.1016/j.scico.2022.102809

[8] From I Stefanakos et al., (2022) Safety Controller Synthesis for a Mobile Manufacturing Cobot. To appear in proceedings Software Engineering and Formal Methods, 28-30 Sep 2022, Humboldt University.

We evaluated the safety-controller synthesis approach in a case study by employing a two-pronged evaluation methodology. First, we used our approach to synthesise a set of safety controller instantiations that meet strict safety constraints and achieve optimal trade-offs between the efficiency of the manufacturing process and the level of residual risk. Second, we developed a digital twin of the manufacturing process and used it to trial one of these safety controller instantiations[9].

[9] More information on our approach, case study and evaluation methodology can be found here.

Analysis

A key element of the project was to build confidence in system safety, both in existing approaches but also in any techniques developed through the project. Specifically, there was a need to analyse safety controllers developed in the Planning strand to ensure that it worked as expected, and did not introduce any new risks. The goal of this strand, therefore, was to evaluate a promising existing method, and develop a further new method, in order to analyse collaborative systems and gain confidence in their safety.

Phase 1

The goal of the analysis strand in Phase 1 was to apply a manual technique (STPA) and a novel semi-automated technique (SASSI) to hazard analysis of the Cobot system based on the spot welding case study. It built on cobot integration, digital twin, and safety controller work undertaken elsewhere within the project.

Task1
In Task 1 for this phase, we carried out an STPA analysis of the case study system based on the spot welding case study. This proved to be tractable (as expected) and identified a range of plausible hazards. In the process, we produced useful intermediate artefacts e.g. a control structure diagram. We refined this based on the feedback from the various stakeholders. A description of the work, and its benefits and limitations, are available online[10].

Task 2
In Task 2 for this phase, we used the digital twin to systematically test the case study system and the controller developed by the Planning strand across a diverse range of situations. “Situations” include a variety of workplace configurations and user actions. Loss events (e.g. harm to the user) and related near-misses were detected and logged. Beyond identifying faults, we showed how this method can provide assurance evidence of the systems’ safe behaviour by targeting systematic situation coverage — we defined a space model with multiple axes of situation type and generated tests using a coverage strategy (e.g. of the form “all combinations of ...”). Our choice of situations to consider was informed by those that appeared significant in the STPA results. We used procedural situation generation to create situations, only executing those tests that improve our coverage.

In order to achieve the above, we worked closely with colleagues on development of the digital twin, establishing the requirements on the twin for it to be useful for the strand and on helping develop the twin to meet those requirements. This involved work in terms of the format of the available information, and the nature of said information. This notably includes defining a set of messages, available as a ground truth in the twin environment, to capture the occurrence of relevant safety events. Those messages allowed us to implement “safety monitors” that detected violations of our safety properties. Throughout the period of the task we held meetings with the Digital Twin team to ensure that these monitors could be implemented.

We selected from the artefacts produced through STPA those situations of interest which could occur and be accounted for in the context of the digital twin. Said artefacts, such as the identified hazards, were formalised to allow their occurrences in the digital environment to be monitored. Further efforts

[10] https://www.york.ac.uk/assuring-autonomy/guidance/body-of-knowledge/required-behaviour/1-2/1-2-1/cobots/

focused on a library of tools to generate and to process data as required, or produced by, the digital twinning environment.

We further drafted, using the same artefacts as a baseline, a set of environmental parameters and their configuration domains which might influence the occurrence of hazardous situations in the system. Those then provided the basis for the exploration of reasonable parameters in the system through automated search techniques.

The simulated environment includes safety conditions which capture both the occurrences of hazards, as well as situations which may lead to such occurrences. E.g. a collision hazard may stem from a Cobot moving while its path is obstructed, and preventing the latter will decrease the likelihood of the former. The formalisation of safety conditions using temporal logic expressions provides the baseline for the verification system traces. The safety conditions may reflect properties only available from the simulated environment, e.g. the position of a cobot in logical regions. This provides for the definition of a ground truth within the context of the simulation. One challenge we identified is ensuring the implementations of these conditions are in line with the related safety analysis artefacts.

To control the analysis, we defined metrics to guide and assess search-based testing heuristics and provided a framework for search algorithms to interact with the Digital Twin. This supported an initial evaluation of different search strategies, Genetic and Quality Diversity algorithms, to evaluate safety aspects of the Demonstrator. This work also supported the infrastructure of the Digital Twin, and the integration of the controller developed in the planning work package.

We defined coverage and fitness metrics based on situation components to guide automated testing strategies and assess their coverage of the situation space. The metrics drove an initial evaluation of automated testing techniques for safety, i.e. Genetic and Quality Diversity algorithms, on a subset of hazards identified through the safety analysis of the demonstrator. Abstraction of the demonstrator led to a highly hazardous configuration space, with no clear benefits of situation-based metrics for the search heuristic. The metrics however helped highlight scenarios of interest amongst the generated ones, and capture the hazardous situations leading to them. Initial results still highlight potential improvements for the metrics and the evaluated demonstrator.

A project deliverable[11] describes our approach to simulation based-testing using the Digital Twin and automated search techniques. It describes the situation space defined by our industrial use case, the prototype simulation based on the digital twin, the analysis framework implementation, and a preliminary evaluation of automated testing techniques.

Phase 2

In Phase 2, the analysis started by performing additional experiments on the spot welding cell use-case using the integrated safety controller. The objective was to assess the safety of the overall system with a safer baseline environment. The work led to a number of refinements to the digital twin, and the validation toolchain. Effort was also spent supporting the Manufacturing Robotics Challenge, preparing the digital twin environment and standalone builds for use by the participants.

[11] Deliverable A2: Situation space description/model, prototype testing simulation, and preliminary evaluation of testing

The industrial cell model had to be refreshed in line with API and core changes in the digital twin to allow for any experimentation. An experimental behaviour module was introduced to ease the definition of complex processes in the environment. The operator model was revised accordingly to support actions such as hand gestures, waypoint following, and interactions with other entities. Collisions on the operator are further tracked with increased accuracy to allow for the identification of the collision region and strength, e.g. to track unsafe collisions with an operator's head. All industrial assets and sensors were similarly reviewed.

Further work on the sensor definition in the digital twin allowed for the modelling of sensor blind spots. Further improvements to the digital twin targeted improving the performance and stability of the simulation to reduce the loss of data points in the digital twin. Refinements on the analysis toolchain focused on the formal safety monitors to reduce the occurrence of false positives in the system, i.e. hazards wrongly identified as occurring. Additional improvements allow for faster and more usable coverage metric computation and reporting.

We released the py-csi-cobotics library (https://github.com/Gaudeval/py-csi-cobotics), which provides building blocks to interact with the digital twin, and analyse the behaviour of a system. The analysis relies on monitoring situations, temporal logic expressions which assess the occurrence of specific system configurations at runtime.

This supported the development of safety concepts, to enable users to highlight areas of concern in the system by describing its actors and attaching safety-relevant labels (e.g. fragile or temperature-sensitive) within the digital twin. Situation definition and monitoring help identify hazardous occurrences, while a rule-based system automatically identifies some safety-relevant situations from the system description. Work focused on refining concept definition, defining an analysis method for practitioners, and prototyping to ensure the applicability of the proposed concepts and rules system.

Further improvements focused on the operator behaviour and model in the Digital Twin, to improve collision detection and introduce static obstacles' avoidance for the operator. A prototype safety controller and the related setup were further introduced using the mobile robot base.

Security

Security of robots, and cobots in particular, is set to be a major issue. Security issues are both general (the state of industrial robotics security is poor or even unaddressed in some cases, witness for example the work by Alias Robotics) and specific, e.g. authentication is regarded as a major problem. In the case of cobots, authenticating the user to the robot, a fundamental security process, seems almost unaddressed. The security work in CSI:Cobot has sought to address fundamental general cybersecurity aspects such as threat modelling and identification of security policy/requirements for cobotic systems, and intrusion detection. We have also sought to investigate suitable user authentication methods to support the general secure operation of cobots.

Phase 1

Threat Modelling and Security Requirements/Policy
We addressed early stage issues for cobot system development. The identification of threats to systems is crucial. In addition, where threats are identified it is important to assess them and develop appropriate countermeasures.

Method
We have adapted the STRIDE threat modelling approach and applied it to our exemplar cobotic system - the UR10 spot-welding system. The approach encourages the identification of actors, assets, entry points, an adversarial model and preliminary security requirements. Threats are identified under important general categories: spoofing, tampering, repudiation, information disclosure, denial of service and elevation of privilege. The overall approach drives the definition of new security requirements to mitigate the identified threats.

Results
A simple STRIDE-based approach has been developed and applied to the exemplar arc-welding system.

Attacks and Attack Detection

Method - Effects Based Attack Simulation and Detection via Digital Twin
We originally envisaged implementing controlled attacks on the actual spot-welding system. The goal was to simultaneously collect data and investigate how well we could identify attacks on the system. Due to Covid constraints we were unable to build the relevant prototype for the actual system. However, the Digital Twin prototype of this system became available in a suitable state of development for us to implement a related and important idea. Namely, that we could *simulate the effects of an attack* and then determine whether intrusion detection approaches. A major benefit of this approach is that we can experiment with the detection of the symptoms even if we do not know how these symptoms might be caused in practice.

Much intrusion detection either detects attacks directly, e.g. by recognising a malicious payload in an incoming packet for some application protocol (which it then blocks or deals with in some way). But also, a great deal of intrusion detection raises alarms based on the *symptoms* or *effects* of an attack, e.g. repeated authentication failures or abnormal numbers of file changes (symptomatic of ransomware, for example). In our work we gave a specific cobot focus to this idea, seeking to detect intrusions that lead to ***integrity*** problems in the operation of the cobot system.

The principal target of our investigation was the physical ***trajectory*** of the spot welding end-effector. The trajectory matters a great deal in many industrial robotic systems. Imagine if the positioning of the artefact to be spot welded was out by 1mm. This would have a significant effect on the integrity of the end products. Additionally, even if the final position at the spot welder is fine, the way it gets there matters. Some trajectories are more efficient than others. Others may place more stress on the robotic system, causing wearout. The Stuxnet malware is an excellent example of this type. However, anything which interferes with the physical trajectory could/should be detectable by *suitable* analysis of the trajectory information available via the digital twin's information gathering capabilities.

Our initial method simulated normal trajectories of the end-effector and monitored their properties (position, velocity etc). Malicious action was then simulated by perturbing such information directly in profiled datapoints. We therefore obtained normal trajectory point information and generated representatives of abnormal trajectory points from it. Machine learning algorithms were deployed to classify normal and malicious points.

Results

For the implemented perturbation strategy the results (shown in TABLE 1) are clearly excellent. Here ***recall*** is the fraction (percentage) of actual malicious points that are deemed anomalous/suspicious by the system. ***Precision*** is the fraction (percentage) of all alarms raised by the detection system that are actually malicious. Thus recall and precision focus on the alarm signalling performance of the intrusion detection system. ***F1*** is an equal balance of recall and precision (the harmonic mean in fact). ***Accuracy*** is a more general measure of what fraction of decisions (alarms or non-alarms) are correct. Note that classification here is at the level of individual points. To escape detection completely all dynamically monitored points would have to escape detection, which is highly unlikely.

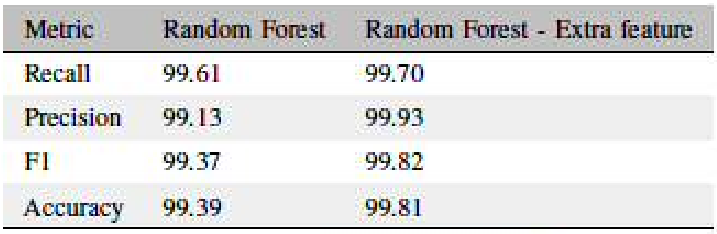

Metric	Random Forest	Random Forest - Extra feature
Recall	99.61	99.70
Precision	99.13	99.93
F1	99.37	99.82
Accuracy	99.39	99.81

TABLE I: Detection results

Authentication of Users to Cobots

Authentication has been identified as a major problem in robotics. There would appear to be very little work carried out on user-to-cobot authentication. PhD work driven by the CSI:Cobot agenda has addressed this issue with considerable success. One issue is that we must protect against outsiders masquerading as legitimate users but also protect against recognised users doing tasks they are simply not allowed to (e.g. when a colleague steps in to do a worker's task even though they are not authorised to do so). There is a need to *continually* assess who is accessing the cobot. Our principal observation is that in industrial settings engaging in the usual user authentication mechanisms may be impossible or highly inconvenient. Imagine for example, having to repeatedly take off gloves to type in a password. Hence we must develop mechanisms that are both high performing but as ***unobtrusive*** as possible.

Method - Leveraging Available Cobot information

A clear and original way to provide user-to-cobot authentication is to adopt a continuous behavioural biometrics approach. Loosely speaking, we measure and classify characteristics of the user's working that are distinctive.

Experiment 1) Initial work was a simple maze guiding task using a commercial UR10 Robot in compliance mode (this is the same robot that is used in the arc-welding system). As the user guides the end-effector round the maze, we monitor available information at the end effector and the joints, including position and force and torque in x, y, and z directions. This requires **NO EXTRA FUNCTIONALITY** in the cobot since the UR10 provides access to this info directly. Similar information will be available from other commercial robotic arm robots.

Experiment 2) Under Covid we could not perform additional physical user experiments and so we leveraged an available sensorised glove data set where the user carried out a variety of manufacturing related tasks (screw high, screw, medium, screw low etc.). Again we showed that the information coming from the sensorised glove could act as a means for authentication. This has longer term implications for any exo-skeletal assistive cobotics that could be adopted in manufacturing settings. Once again, we leverage information that is already maintained or determined by the cobot. We anticipate similar state or dynamics information to be available with assistive robotics technology. Again, **NO EXTRA FUNCTIONALITY** was required of the glove. We made use of the information it already generated.

Experiment 3) We used a simulator for the maze task and monitored the mouse actions used to control its operation. Under Covid restrictions this allowed experiments to proceed with subjects at home running our framework software. **But it is also a good proxy for remote operation**. Remote operation is set to be of major interest in industrial operation. We monitored simple measurements such as end effector position (x,y) and related component velocities as the user guided the end-effector round the maze. We also implemented a more sophisticated intruder model based on statistical profiling.

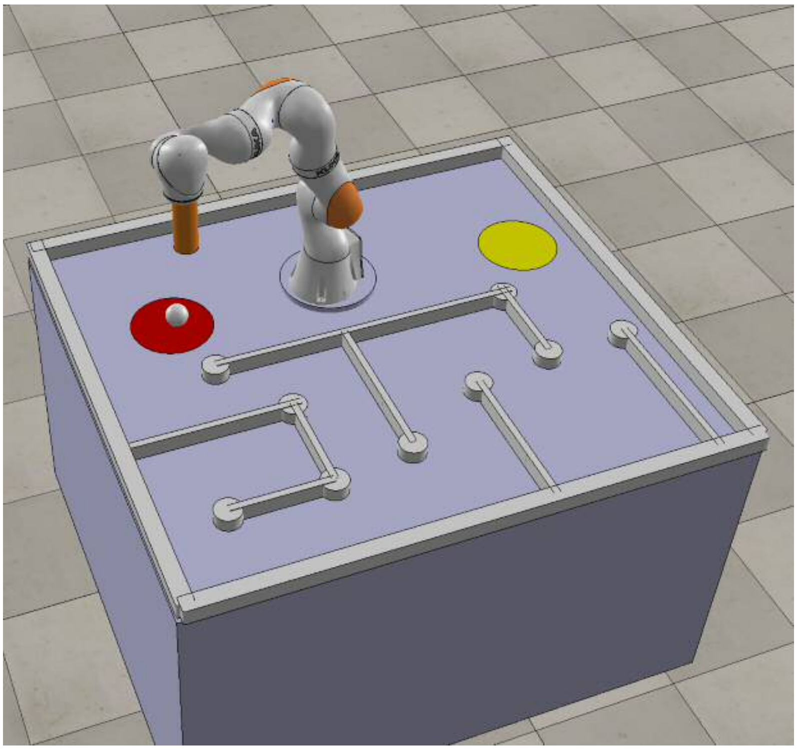

Figure 10: Remote controlled maze

Results

In all cases a continuous biometric approach shows significant promise.

In all cases the user simply goes about their task. The authentication is done without any additional effort on the part of the user.

Phase 2

Two extension aspects have been developed: application of threat modelling and security requirements/policy development approach to a more sophisticated case study involving wifi networking and mobile robots; and development of software to support the profile monitoring for iDS purposes.

A report on the application of the threat modelling and security requirement/policy guidance has been produced. A BoK contribution has also been produced.

Software developed to support IDS work will be further exploited and refined under the Security of Manufacturing Grant.

Existing security standards for robotics are under examination. The application of safety analysis techniques to the security analysis of cobots is currently underway.

Demonstrators

The goal of the demonstrator strand was to provide a software sandbox and digital-twin for the testing and demonstration of the other work-packages. The platform was tasked with emulating the process and performance of the real-world system, whilst providing both an interface for continued analysis, and connectivity with the physical system in place at our industrial partners. As the project progressed, and with the onset of COVID-19, the need for a more sophisticated system arose, able to operate independently of limited equipment access (see Section Digital Twinning). With the introduction of mobile systems, more sensors, and research tools, the resulting demonstrator is now a sophisticated software package that has been demonstrated on three real-world systems.

Phase 1

Task one

We developed a model of the spot welding case-study, utilising information from the industrial partner, generating an interface for a human operator, robotic manipulator, safety lidar and automatic welder, that can effectively represent and participate in a realistic process. The developed demonstrator accurately simulated the communication of these process members such that the process could be controlled and augmented by the addition of a synthesised safety controller, in addition to providing an analysis API to support its validation. Development of an initial connection interface with the Robotic Operating System (ROS) provided a basis for interacting and controlling the physical process, but with limited access demonstration was only initially tested in the laboratory.

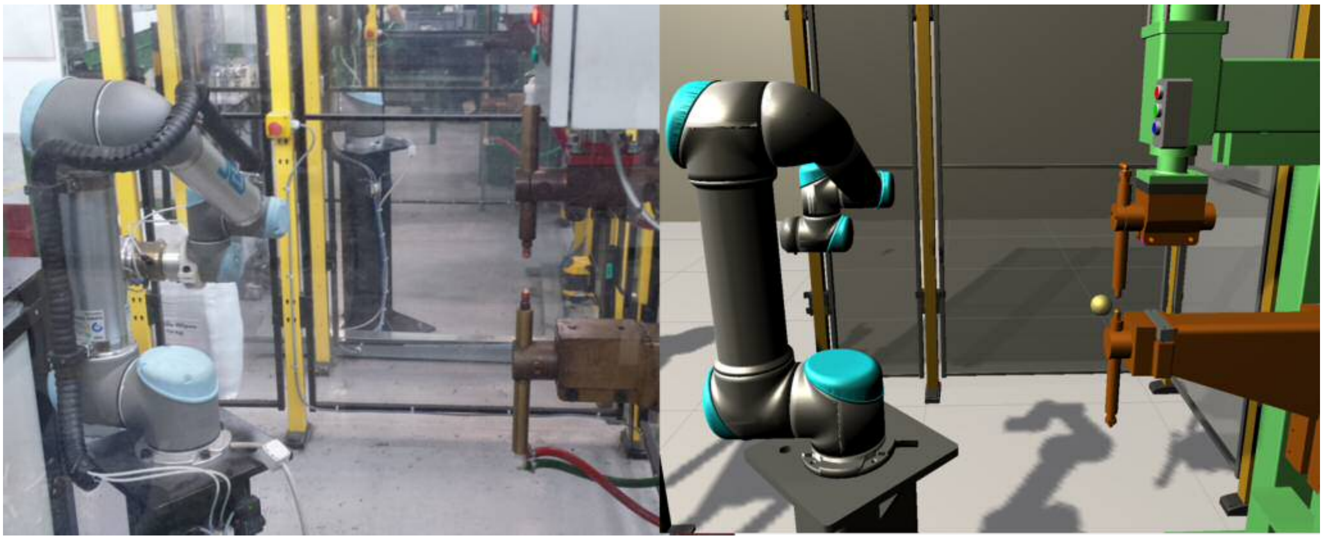

Figure 11: Real (left) and digital twin (right) of the spot welding process

Task two

The second task was to extend the existing testing API to support online automated testing of the spot welding demonstrator to allow systematic coverage analysis of a synthesised safety controller. The demonstrator required simplified but sufficient testing heuristics. With them established, the demonstrator was packaged using docker for parallel deployment with a faster than real-time configuration tool controllable from an externally defined configuration file.

Task Three

The third task set to the demonstrator branch was in the development of a digital twin to facilitate the Manufacturing Robotics Challenge 2019. Hosted at the AMRC on the 20th July 2021, this demonstrator provided access to a COVID-19 sample testing lab digital twin, centred around a real-robotic manipulator. This allowed 38 students from 11 countries to engage with industry 4.0 technologies by learning how to collaboratively program a real robot to respond to the digital-twin.

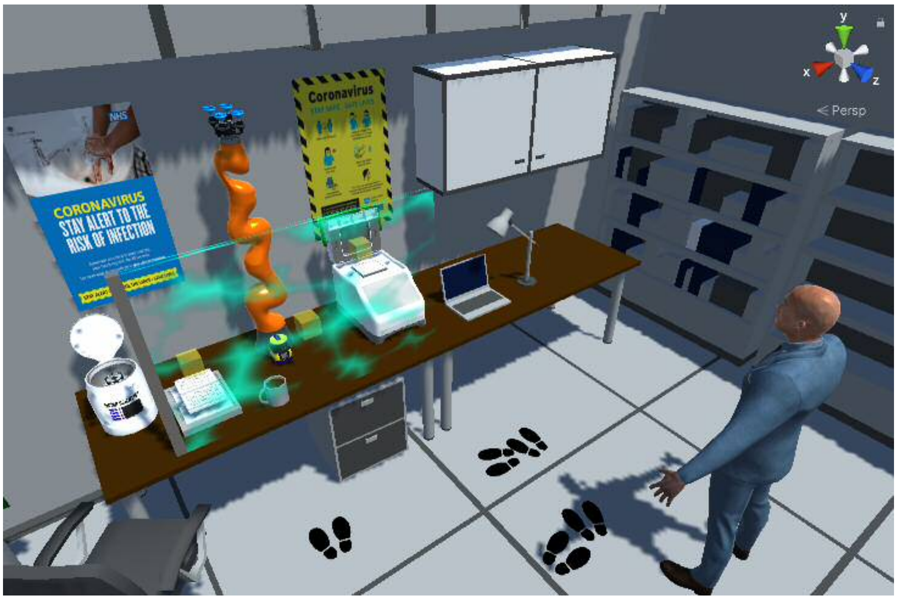

Figure 12: The Covid-19 laboratory digital twin

The demonstrator for this competition took the form of an interactive environment, with individual systems represented as complete digital twins. This required significant restructuring of the underlying architecture to support more accurate human modelling, more generalised process, system and service representations, in addition refining new and existing methods of communication/networking with robotic devices. It is this restructuring which ultimately lead to the formation of the CSI digital twin framework (DTF) by the end of phase one.

Phase 2

The goal of the demonstrator branch in phase two was to produce a digital twin of the collaborative machine tending process at the AMRC Gear Centre (GC), centred around a human operator and the mobile iAM-R platform. The focus of the demonstrator branch can be summarised as follows:

Task 1

We 1) expanded the existing implementation of robotic systems within the DTF to accommodate mobile robots and hybrid manipulators (arms with mobile bases); 2) ensured compatibility with our existing software infrastructure and analysis techniques for scenarios when movement of the operator is unrestricted and the shop floor is open; 3) developed a digital twin from the process description and shop-floor diagrams to match the real-world process.

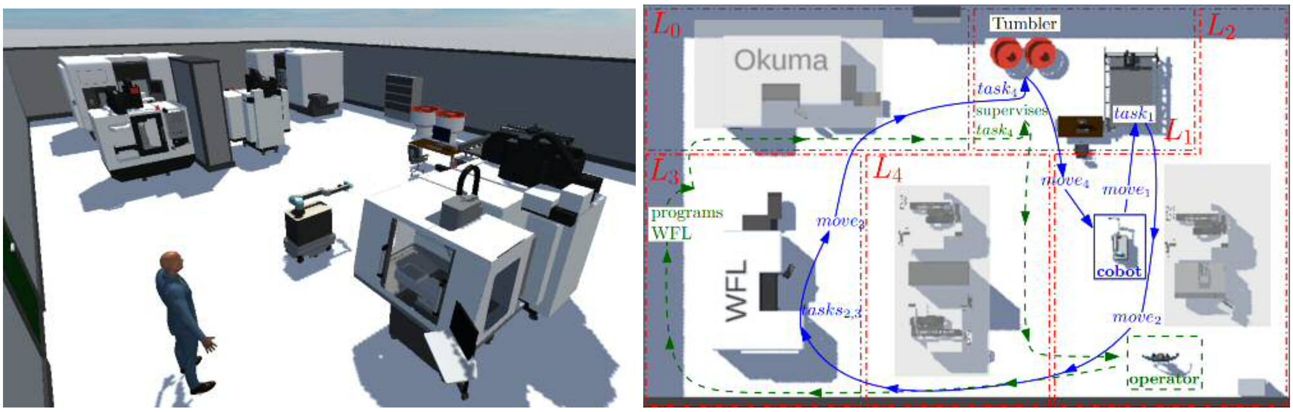

Figure 13: The Gear Centre twin. 3D visualisation (left), floorplan and process sequence (right)

Task 2

We expanded the existing communication infrastructure both with the DTF and in ROS to establish connection and expose the iAM-R's mobile base, arm and gripper for participation in the GC process. Using insights from the previous phase, further ROS development, and collaboration with the AMRC, we successfully developed a unified ROS library combining the UR10 (used in Phase 1, Task 1), the KUKA iiwa (Phase 1, Task 3) and the new iAM-R.

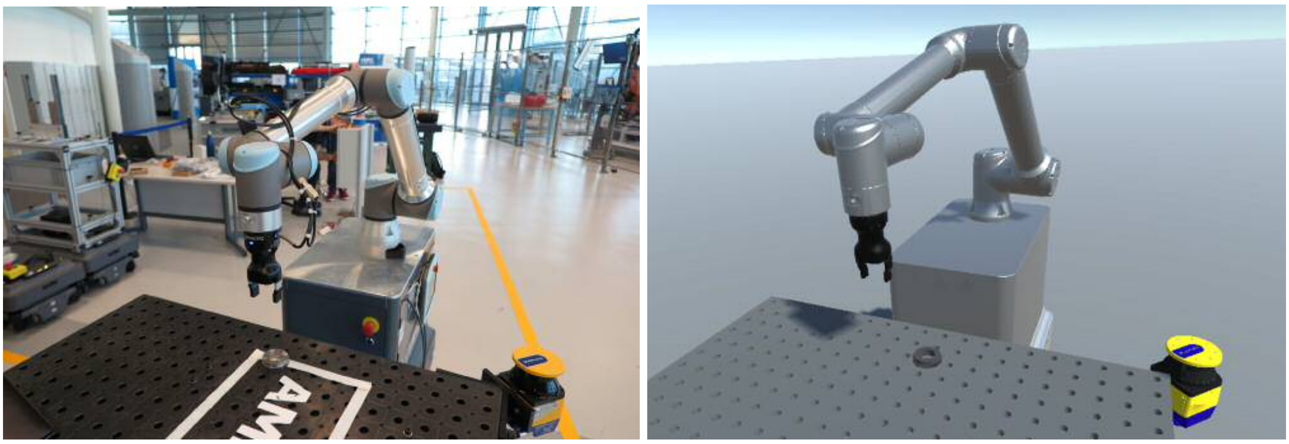

Figure 14: The iAM-R (left) and digital twin (right)

The providers of base game-engine Unity, release a toolbox for the utilisation of ROS and ROS2 with their base platform. It was therefore necessary (due to software conflicts) to update our robot structure to support their ROS plugin and drivers. This had the additional effect of lowering our barrier toward a future transition into ROS2. We have integrated this with our accompanying ROS library in addition to the application of the DTF in AR/VR.

Task 3

In collaboration with ongoing work from the Analysis strand, and our safety policy and regulatory partners, the final task of the demonstrator branch was to examine how the lessons learned from Phase 1 and Phase 2 might be consolidated by refining our approach to safety analysis. We have developed a prototype exchange language for the communication of safety standards using situation-based policies that support the mixture of abstract quantities with practical process decisions enactable through a DTF scenario.

Lessons learnt

The CSI:Cobot project has enabled us to address key challenges relating to the safety and security assurance of collaborative robots, a critical barrier to their further adoption in industrial environments. Working with stakeholders, and through existing and planned process case studies, we have co-created a series of research challenges, each designed to address a requirement of one or more partners. By addressing these challenges, we have demonstrated concepts to help remove constraints to collaborative working in ways which also elicit confidence in users. The close involvement of our regulatory partner has supported an exchange of knowledge to ensure our methods are aligned with current and future regulatory thinking, whilst providing insights into future capabilities emerging from the research community.

The following provide a summary of lessons from individual research themes:

Digital Twins

Conventional approaches to digital twins are dominated by domain-specific, application orientated digital twins with an emphasis on research. Complete frameworks that promote modularity, reconfiguration, especially those that are cross domain, represent a unique and valuable opportunity to generate research impact and increase knowledge transfer between academia and industry. This has already been seen in this project in how heavily we have contributed to the other work-packages; whether in the development of the DTF, our associated works and in supporting collaborative works. Involvement from our regulatory partner has also led to the proposal of a safety framework and exchange format within the DTF. This represents an opportunity to further enhance safety and regulator understanding across a wide range of collaborative robotic processes.

Sensing

The proposed framework based on deep learning and digital twin techniques allows to visually detect the presence of robots and humans and classify their interactions in a collaborative working cell. Safety related decisions can be made according to the separation of robots and humans. The digital twin framework is modular and flexible, it allows addition of other algorithms and can be connected to the control unit of the robot.

Experimental validation indicates satisfactory performance of the trained deep learning model based on a UR10 platform. Using a digital twin to generate diverse synthetic data helps to improve generalisation and robustness of the trained model, ensuring it is less affected by dynamic changes in environmental conditions. To conclude, the experimental outcomes suggest the potential of the proposed framework as a reliable and intelligent sensing solution to the safety requirements.

The developed AI and machine learning algorithms have the potential to enhance the safety of human-robot interactions, but this requires changes in the regulatory standards. However, there is a need for systematic validation and verification of the Digital Twin framework for different manufacturing conditions and different case studies.

Planning

Rigorously developed safety controllers enable cobots to mitigate hazards with optimal performance-risk trade-offs in manufacturing processes.

Augmenting collaborative robots with mobility allows them to support a broad range of manufacturing processes, but introduces additional safety concerns that are not captured by existing cobot standards.

Analysis

STPA can produce hazard descriptions that can be used to help define safety properties and monitors for simulation-based software testing. Non-trivial work is needed, however, to translate these into something concrete enough to monitor in a simulation. There is not obviously any more guarantee of completeness than with any other hazard analysis method.

The requirements of simulation-based testing do not always square easily with the other requirements of digital twins, or with the assumptions made by Unity and ROS.

The approach can reveal faults in robot control software.

Security (Policy)

A STRIDE-based approach to threat modelling and the development of countermeasures can be applied effectively to cobotic systems. Software to support its application would help its use.

Security (Attacks and attack detection)

The notion that a simulator can be used to generate data for intrusion detection purposes is powerful and can be used for many types of intrusion detection work. Monitoring the system at an appropriate granularity can be carried out via the digital twin's information gathering capabilities. The digital twin can be an excellent source of security information.

The focus on physical trajectory is important for industrial cobots. Maliciously affecting the physical operation of the cobot has major integrity implications and so focusing on this aspect is an important advance.

Although this is motivated by actual physical trajectories, it should be possible to generalise the notion of 'trajectory' to any data variable maintained by the digital twin. Thus, there is the possibility for a unified approach to monitoring and raising alarms based on both physically-related data and physically-related data.

Security (User Authentication)

Continuous biometric authentication has considerable promise as a means of authenticating users to cobots.

In actual physical systems the approaches adopted may typically be implemented with **NO ADDITIONAL FUNCTIONALITY IN THE COBOT.**

The approach in all cases was unobtrusive. The user is **NEVER REQUIRED TO CARRY OUT AN EXPLICIT AUTHENTICATION ACTION**. They are simply monitored as they go about their task. In industrial cobotic settings, this unobtrusiveness is highly appealing.

There is a need to determine how more traditional authentication needs to be incorporated. For example, at the start of a shift a user may be required to authenticate in a more traditional way. Also, if the system becomes very suspicious and locks the user out, there needs to be a means of re-initiating the authentication.

Overall the work has been very successful. Its two major features - no extra functionality in the cobot, no explicit authentication action needed - make it an appealing choice in industrial cobotic settings.

Demonstrators

Development of a framework able to support multiple research objectives, tools and demonstrators has been challenging. The movement away from a contained environment towards more flexible collaborative processes in Phase 2 has required software infrastructure that is more modular, and versatile than in Phase 1. The introduction of a modular infrastructural design at the onset of phase 2 required the reevaluation of our existing concepts of digital twins, services and their implementation. In addition our existing analysis tools, centred around those definitions, were reassessed to ensure that the software remained compatible by maintaining the same research/analysis interface. Throughout phase 1 and 2, the demonstrators have faced deployment *for interaction*, and deployment *for analysis*, which has presented further design challenges. This has required that our concept of digital twins, communication and services be malleable enough to be deployed in these different configurations and respond appropriately when executed with suppressed functionality, faster than real-time or without network connectivity.

Towards the end of Phase 2, a shift towards the extension of the *interactive* deployment demonstrated some promising results in establishing an effective, persistent digital twin. This was due to the introduction of a networked VR/AR client, that enabled visualisation and control of data from any source within the framework. Whilst these were only initial results, this immediately presents a solid opportunity for increasing safety (and process) understanding, training and dynamic interactions.

Acceptance of novel safety techniques, regulatory change, and incorporation of new ideas into standards takes time beyond that available within a single project such as this. To ensure new ideas and developments continue to impact discussion requires that they are disseminated widely, and to relevant communities. It also requires that new researchers and safety engineers are trained with the knowledge developed through such work, and that opportunities are provided for them to develop. Through this project we have become more engaged with regulatory bodies, and standards communities, and have published and presented our findings widely. We have provided training to robotic engineers, made our data and tools available for others to use, and are embedding our learnings into new taught courses. Finally, we have secured new funding and projects to address questions arising from the work, and develop our ideas towards implementation.

Impact table

Publications	Eimontaite, I., Cameron, D., Rolph, J., Mokaram, S., Aitken, J.M., Gwilt, I., Law, J. (2022). “Dynamic Graphical Instructions Result in Improved Attitudes and Decreased Task Completion Time in Human–Robot Co-Working: An Experimental Manufacturing Study”. Sustainability, 14,3289. https://www.mdpi.com/2071-1050/14/6/3289 Gleirscher, M., Calinescu, R., Douthwaite, J., Lesage, B., Paterson, C., Aitken, J., Alexander, R., Law, J. (2022).“Verified synthesis of optimal safety controllers for human-robot collaboration”. Science of Computer Programming, Volume 218. https://doi.org/10.1016/j.scico.2022.102809 Gleirscher, M., Johnson, N., Karachristou, P., Calinescu, R., Law, J., & Clark, J. (2022). Challenges in the safety-security co-assurance of collaborative industrial robots. In *The 21st Century Industrial Robot: When Tools Become Collaborators*(pp. 191-214). Springer, Cham. https://arxiv.org/abs/2007.11099 Pisanelli, G., Tymczuk, M., Douthwaite, J. A., Aitken, J. M. & Law, J. (to appear). “ROSIE: A ROS adapter for a Modular Digital Twinning Framework”. In 31st IEEE International Conference on Robot & Human Interactive Communication (RO-MAN), 2022 Stefanakos, I., Calinescu, R., Douthwaite, J., Aitken, J., Law, J. (2022). “Safety Controller Synthesis for a Mobile Manufacturing Cobot” (to appear). In the 20th International Conference on Software Engineering and Formal Methods (SEFM). https://eprints.whiterose.ac.uk/189758/1/Safety_Controller_Synthesis_for_a_Mobile_Manufacturing_Cobot.pdf Almohamade, S. S., Clark, J. A., Law, J. (2021) *Behaviour-Based Biometrics for Continuous User Authentication to Industrial Collaborative Robots*. SECITC 2020. 185-197 https://link.springer.com/chapter/10.1007/978-3-030-69255-1_12 Almohamade, S. S., Clark, J. A., Law, J. (2021). Continuous User Authentication for Human-Robot Collaboration. In The 16th International Conference on Availability, Reliability and Security (ARES 2021). Association for Computing Machinery, New York, NY, USA, Article 115, 1–9. https://dl.acm.org/doi/abs/10.1145/3465481.3470025 Almohamade, S. S., Clark, J. A., Law, J. *Mimicry Attacks Against Behavioural-based User Authentication for Human-Robot Interaction.* 4th International Workshop on Emerging Technologies for Authorization and Authentication (ETAA), October 2021. https://link.springer.com/chapter/10.1007/978-3-030-93747-8_8 Douthwaite, J. A., Lesage, B., Gleirscher, M., Calinescu, R., Aitken, J. M., Alexander, R., & Law, J. (2021). “A Modular Digital Twinning Framework for Safety Assurance of Collaborative Robotics”.

	Frontiers in Robotics and AI, 8(December), 1–17. https://doi.org/10.3389/frobt.2021.758099 Gleirscher, M., Calinescu, R., Woodcock, J. (2021) RiskStructures: A design algebra for risk-aware machines. Formal Aspects of Computing. 2021 May 26:1-40. https://eprints.whiterose.ac.uk/172137/7/Gleirscher2021_Article_RiskStructuresADesignAlgebraFo.pdf Lesage, BMJ-R & Alexander, R (2021). SASSI: Safety Analysis using Simulation-based Situation Coverage for Cobot Systems. in Proceedings of SafeComp 2021. https://pure.york.ac.uk/portal/en/publications/sassi-safety-analysis-using-simulationbased-situation-coverage-for-cobot-systems(310627aa-9b13-44fb-8a9f-6afe2cb3e618).html Gleirscher, M. (2020) Yap: Tool Support for Deriving Safety Controllers from Hazard Analysis and Risk Assessments. Luckuck, M. & Farrell, M. (Eds.), Formal Methods for Autonomous Systems (FMAS), 2nd Workshop, Electronic Proceedings in Theoretical Computer Science, 329, 31-47. Open Publishing Association, 2020. https://eprints.whiterose.ac.uk/168776/1/paper.pdf Gleirscher, M., & Calinescu, R. (2020, October). Safety controller synthesis for collaborative robots. In 2020 25th International Conference on Engineering of Complex Computer Systems (ICECCS) (pp. 83-92). IEEE. https://arxiv.org/pdf/2007.03340.pdf Wang, S., Zhang, J., Wang, P., Law, J., Calinescu, R., Mihaylova, M. (under review). “A Deep Learning-enhanced Digital Twin Framework for Improving Safety and Reliability in Human-Robot Collaborative Manufacturing”. Submitted to Robotics and Computer-Integrated Manufacturing.
Presentations	R. Calinescu, UKRI Trustworthy Autonomous Systems Node in Verifiability seminar, April 2021. R. Calinescu, November 2021 visit at UKAEA/RACE, including presentation and discussion of intermediate project results. R. Calinescu, April 2022 visit at NASA Ames, including presentation and discussion of intermediate project results. Y. Cogurcu, AR/VR technology for interfacing users and digital twins. DTOP Workshop, DigiTwiin, Tuesday 21 June 2022, University of Sheffield. J. Law, J. Douthwaite, “Digital twinning for safety assurance in human-robot collaborative environments”. Presented at DTHIVE seminar, University of Sheffield, Feb 2022. L. Mihaylova, UKRI Trustworthy Autonomous Systems Workshop, July 2021. S. Almohameide, International Workshop on Emerging Technologies for Authorization and Authentication (ETAA), October 2021.

	S. Almohameide, The 16th International Conference on Availability, Reliability and Security (ARES 2021) J. Douthwaite, “Digital twinning for industrial robotics”. Presented at European Robotics Forum 2021. J. Law, “Safety and security assurance using digital twins”. Presented at the first DTHIVE (Digital Twins for High Value Engineering applications) workshop. 09/12/2021. J. Law gave an invited talk at a webinar organised by the Danish Embassy on “The Next Wave of Robotics” on 26/01/2021. M. Gleirscher, “Safety Controller Synthesis for Human-Machine Collaborations”. Presented at European Robotics Forum 2021. B. Lesage, *SASSI: Safety Analysis using Simulation-based Situation Coverage for Cobot Systems*. Presented at SafeComp 2021 conference. Audience: ~40 research and industrial. [link] B. Lesage, *SASSI: Safety Analysis using Simulation-based Situation Coverage for Cobot Systems*. Presented at Autonomy and Safety Seminars (AAIP and University of York HISE seminar series). B. Lesage, “Situation Coverage for the Safety of Collaborative Robots”. Presented at European Robotics Forum 2021. S. Almohameide, International Conference on Information Technology and Communications Security, SecITC 2020 M. Gleirscher, Safety controller synthesis for collaborative robots. Presentation at ICECCS 2020. B. Lesage presented work on Hazard Analysis at an AAIP/HISE seminar in 2020 B. Lesage, Situation Coverage for the Safety of Collaborative Robots, Presentation to AAIP at UoY, Nov. 2020 M. Gleirscher and R. Calinescu. Autonomous Mobile Collectives in Complex Physical Environments”. Dagstuhl Seminar, October 2019. R. Calinescu presented a summary of the stochastic modelling work at one of the regular meetings of the Working Group developing the IEEE Guide for the Verification of Autonomous Systems (https://standards.ieee.org/project/2817.html).
BoK Contributions	https://www.york.ac.uk/assuring-autonomy/guidance/body-of-knowledge/implementation/2-6/2-6-1/cobots/

	https://www.york.ac.uk/assuring-autonomy/guidance/body-of-knowledge/implementation/2-2/2-2-4/2-2-4-3/cross-domain-cobots/ https://www.york.ac.uk/assuring-autonomy/guidance/body-of-knowledge/required-behaviour/1-2/1-2-1/cobots/ https://www.york.ac.uk/assuring-autonomy/guidance/body-of-knowledge/implementation/2-7/cobots/ Threat Modelling and Security Requirements/Policy (submitted) A Risk Assessment Approach for Collaborative Robots (submitted) Gaining approval for operation of RAS (in preparation)
Funding	Trustworthy Autonomous Systems Node in Resilience (REASON). £3M, EPSRC Digital Twins for High-Value Engineering Applications (DTHIVE). £900k, EPSRC/ATI Security of Digital Twins in Manufacturing. £611k, EPSRC LongOps SBRI Challenge 1. £150k, UKAEA/RACE Skills and Education in Robotics. £50k, EPSRC IAA (University of Sheffield) SALSA2d. €49k, EU H2020 FSTP (TRINITY) Towards Improved Safety and Reliability of Cobots. £38k, Research England NPIF QR (University of Sheffield) ROSIE2. £36k, EPSRC HEIF & IAA (University of Sheffield) ROSIE2.5. £25k, Research England NPIF QR (University of Sheffield) Skills and Education for Robotics and IoT. £23k, Research England Pitch-In UK-RAS Manufacturing Robotics Challenge. £5k, EPSRC UK-RAS Many-worlds: multi-instance deployment of digital twins and digital-twin based analyses. £5k, University of Sheffield Research Software Engineering - Tier 2 Support
Other outputs and impact (e.g. tools, datasets)	Examples, sample media and data from our collaborative works can be found on our repository (https://github.com/CSI-Cobot/CSI-artefacts). Dataset: J. Zhang, S. Wang, P. Wang, L. Mihaylova, and J. Law, “A vision data repository for human-ur10 robot interactions in manufacturing." [Online]. Available: https://doi.org/10.15131/shef.data.16669315.v1. [Accessed: 13-Jul-2022]. Dataset: S. Wang, J. Zhang, L. Mihaylova, and J. Law, “ Human-Robot Video Data from a Manufacturing Factory." [Online]. Available: https://doi.org/10.15131/shef.data.19299539.v1. [Accessed: 13-Jul-2022].

	Annotation tool: S. Wang, J. Zhang, P. Wang, and L. Mihaylova, “Semi-automated labelme, a deep learning based annotation tool." [Online]. Available: https://github.com/wongsinglam/Semi_Labelme. [Accessed: 13-Jul-2022]. A Python framework for controlling and processing experiments built upon the CSI:Cobot Digital Twin Framework (DTF) — https://github.com/Gaudeval/py-csi-cobotics Datasets from authentication work and from IDS work will be made publicly available.

ASSURING
AUTONOMY
INTERNATIONAL PROGRAMME