

This is a repository copy of *What's my role? Modelling responsibility for AI-based safety-critical systems*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/206868/>

Version: Submitted Version

---

**Monograph:**

Ryan Conmy, Philippa Mary [orcid.org/0000-0003-1307-5207](https://orcid.org/0000-0003-1307-5207), Habli, Ibrahim [orcid.org/0000-0003-2736-8238](https://orcid.org/0000-0003-2736-8238), Al-Qaddoumi, Joanna et al. (2 more authors) (2023) *What's my role? Modelling responsibility for AI-based safety-critical systems*. Working Paper.

---

**Reuse**

This article is distributed under the terms of the Creative Commons Attribution-NoDerivs (CC BY-ND) licence. This licence allows for redistribution, commercial and non-commercial, as long as it is passed along unchanged and in whole, with credit to the original authors. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

**Takedown**

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

# What's my role? Modelling responsibility for AI-based safety-critical systems

PHILIPPA RYAN, ZOE PORTER, JOHN MCDERMID, and IBRAHIM HABLI, Department Of Computer Science, University of York, UK  
JOANNA AL-QADDOUMI, York Law School, University of York, UK

AI-Based Safety-Critical Systems (AI-SCS) are being increasingly developed and deployed in the real world. These can pose a risk of harm to people and the environment, hence reducing that risk is an overarching priority during development and operation. Many contain Machine Learning (ML) components whose performance is hard to predict. As more AI systems become autonomous, a layer of risk management via human intervention has been removed. Following an accident it will be important to identify causal contributions and the different responsible actors behind those to learn from mistakes and prevent similar future events. Many authors have commented on the "responsibility gap" where it is difficult for developers and manufacturers to be held responsible for behaviour of an AI-SCS which contributes to harm. This is due to the complex development cycle for AI, the uncertainty in black-box AI components performance, and dynamic operating environment. Instead, a human operator of the AI-SCS can become a "liability sink" absorbing blame for the consequences of AI-SCS outputs they weren't responsible for creating, and may not have sufficient understanding of.

This cross-disciplinary paper considers different senses of responsibility (role, moral, legal and causal), and how they apply in the context of AI-SCS safety. We use a core conceptual formulation *Actor(A) is responsible for Occurrence(O)* to create detailed role responsibility models, including related tasks and resources. Our aim is to present a practical method to precisely capture responsibility relationships, and provide clarity on the previously identified responsibility issues. We propose an analysis method to review the models for safety impacts. Our paper demonstrates the utility of the approach with two practical examples. The first is a retrospective analysis of the Tempe Arizona fatal collision involving an autonomous vehicle. The second is a safety focused predictive role-responsibility analysis for an AI-based diabetes co-morbidity prediction system. We show how our notation and analysis can illuminate responsibility issues during the AI development process and identify the impact of uncertainty surrounding how tasks were performed. For both examples, our focus was primarily on safety, with an aim of reducing unfair or disproportionate blame being placed on operators or developers of AI-SCS. We present a discussion and avenues for future research, including the development of a richer causal contribution model.

CCS Concepts: • **Hardware** → **Safety critical systems**; • **Software and its engineering** → **Software safety**; • **Computer systems organization** → **Dependable and fault-tolerant systems and networks**.

Additional Key Words and Phrases: AI-based systems, Safety, Responsibility

## ACM Reference Format:

Philippa Ryan, Zoe Porter, John McDermid, Ibrahim Habli, and Joanna Al-Qaddoumi. 2023. What's my role? Modelling responsibility for AI-based safety-critical systems. 0, 0, Article 0 ( 2023), 22 pages. <https://doi.org/XXXXXXX.XXXXXXX>

## 1 INTRODUCTION

AI-Based Safety-Critical Systems (AI-SCS), such as autonomous vehicles, inspection drones, medical diagnosis systems and others, are increasingly being developed and deployed in the real world. AI-based systems provide many technical and societal challenges, particularly when they include black box Machine Learning (ML) components [4] and are used in a dynamic operating environment. AI-SCS typically have a complex development chain with many actors, including software developers,

Authors' addresses: Philippa Ryan; Zoe Porter; John McDermid; Ibrahim Habli, Department Of Computer Science, University of York, York, YO10 5DD, UK; Joanna Al-Qaddoumi, York Law School, University of York, York, YO10 5GD, UK.

2023. ACM XXXX-XXXX/2023/0-ART0  
<https://doi.org/XXXXXXX.XXXXXXX>

data scientists, regulators, service providers, suppliers [11, 43], whose roles all contribute to safety. This provides many areas of uncertainty, and difficulty in assuring they will behave in an acceptably safe way through-life. Additionally, for many AI-SCS there is limited or no interaction with a human operator [51, 52] and there are gaps understanding who was responsible and/or liable following an incident or accident [9]. Human operators who do compensate for uncertainty in AI behaviour, face an unfair burden of risk, absorbing blame for the consequences of system outputs they weren't responsible for creating [31].

The notion of responsibility has long been studied in a philosophy and law, often with a backwards-looking perspective to understand events that have occurred. However, it has not been studied in depth in safety engineering, which also requires a forwards-looking perspective to predict potential harm and prevent it or reduce its severity. Nor has it been practically applied in the context of AI-SCS, although there is much research into responsibility issues for AI-SCS. This paper considers how a number of different senses of responsibility (role, moral, legal and causal [18]) apply to actors developing and operating AI-SCS. We have an emphasis on role responsibility, and demonstrate how analysis can help to uncover gaps, risk transfer, and burden of risk when developing and operating AI-SCS.

This paper provides the following contributions. First, we extend earlier work on responsibility modelling [43] and task analysis [7, 32] to develop a notation and method for modelling role responsibility and role relationships for AI-SCS. This can be used to describe and link actors, occurrences and resources across AI-SCS lifecycle phases. Second, we apply the notation to two different AI-SCS scenarios, showing how our models illuminate issues (e.g., gaps, conflicts) of responsibility and provide greater clarity on risk distribution and safety impact. We particularly focus on how engineering roles and technical challenges during AI development can lead to burden of risk for operational roles.

In section 2 we review related work in this area, considering problems about the attribution of responsibility for AI-SCS, and motivate the need for responsibility analysis. In section 3 we present our responsibility modelling notation, and a method for analysing and reviewing role responsibility models. Section 4 contains our first example study of responsibility modelling, using the accident findings of the fatal collision of a pedestrian with an autonomous vehicle in Tempe, Arizona. Section 5 contains a second example study of an AI-based diabetes co-morbidity predictor, used to support a clinician during medical diagnoses of Type II diabetes patients. We present a complex forwards-looking model covering both development of the predictor and its operational use. We show how this model can be analysed to determine role responsibility issues. Finally, there is a discussion and considerations for future work in section 6.

## 2 BACKGROUND AND RELATED LITERATURE

### 2.1 AI-based Safety critical systems and responsibility gaps

Safety-critical systems are defined as those whose failure, under certain conditions, can lead to harm to humans or the environment [21]. There are many different examples, such as medical devices, nuclear power plants, defence systems, cars and aircraft. AI-SCS include AI elements to replace some or all of a human function, such as driverless autonomous vehicles, unmanned inspection drones, and medical diagnoses. An AI-SCS is subject to differing legal and regulatory regimes, depending on domain, type, country of use and even degree of autonomy. Rather than attempt to cover all these we briefly summarise some common practices which will broadly apply when developing an AI-SCS.

Typically the manufacturers and operators of an AI-SCS will provide a Safety Assurance Case (SAC) or safety justification that the system is *acceptably safe* [28, 53] for use in a defined context.

This should demonstrate that all reasonable means of risk reduction have been enacted, by following established good engineering practice (e.g., software testing), adding operational controls (e.g., manual intervention), or reducing severity of accidents (e.g., sprinkler systems). Good practice is encapsulated in domain specific standards, such as [25] for medical devices, [24] for automotive electrical/electronic systems, [46] for avionics, and [58] for defence systems (noting that consensus on good practice is not fully established for AI components [4]). The SAC may require external and independent scrutiny, such as by a regulator. Hence, even though there is residual risk associated with operating the AI-SCS, it may not be considered fair to hold developer's and/or manufacturer's responsible for an accident if they exercised due diligence. The SAC may be used as evidence in such a situation (see Section 2.3).

There are established practices for modelling and analysing operational responsibility for tasks performed by humans for conventional critical systems (e.g., [7, 8, 14, 32]). A complication for AI-SCS is that an AI component performs some or all of the decisions and tasks previously assigned to a human operator, including some to reduce the risk of accidents [9, 36]. E.g., an autonomous drone will take actions to avoid collision, rather than a remote operator. Alternatively, an operator may be required to only make interventions when the AI-SCS cannot, increasing cognitive load [6, 60]. When there are accidents, the operator may absorb blame for the consequences of AI-SCS outputs they weren't responsible for creating, and may not have sufficient understanding of. This is termed "liability sink" [31].

Many authors have commented on the "responsibility gap" for AI-SCS e.g. [9, 17, 34, 42], where it is difficult for developers and manufacturers to be held responsible for behaviour of an AI-SCS which contributes to harm. This is due to difficulty in a) ensuring safe behaviour of black-box AI component, and b) defining precise operating domains [4]. It has even been suggested this uncertainty can be used to avoid responsibility (see 2.2). In [62], the authors argue for the AI-based system itself to be programmed to be responsible, although practical methods to do this are unknown.

Munch [37], notes that a responsibility gap could be viewed positively, because it avoids unfairly blaming individuals, and replacing human agency with AI can remove the burden of decision making. We reject the idea of not attempting to clarify and, if possible, reduce responsibility gaps. This is not to punish individuals, but to avoid putting people in the situation where they could face unfair blame<sup>1</sup>, and so that we can learn from mistakes and avoid them in future. Good safety practice is based on a *just culture* as proposed by [12][57][40], in which incident reporting is encouraged within an organisation, promoting transparency without fear of blame. In this atmosphere, taking responsibility can be encouraged.

The recent EU AI Act [1] is intended to close some legal responsibility gaps. The current draft has been criticised for not considering some of the unique problems of AI, including its dynamic nature, its complex lifecycle, understanding end user rights [29], and not tracing responsibility back to AI developers [59]. The EU AI liability directive [2] is intended to provide means of pursuing civil liability, but legal review suggests technical issues of identifying causal responsibility (responsibility gaps) are not addressed [42].

One way to ensure that the nearest human operator isn't held solely responsible for the consequences of AI-SCS behaviour after an accident is to articulate the roles and potential liability of all actors (e.g., safety engineers) with safety-related tasks more systematically earlier on in the system lifecycle. Further, this can protect actors, such as engineers or other decision makers, from unfair blame by demonstrating their tasks were performed with due diligence.

<sup>1</sup>Two aspects of blame, moral and legal liability, are discussed later in the paper. We note that one can have liability without fault, and typically blame is only an appropriate response when someone is at fault.

## 2.2 Responsibility complexity in AI-SCS development and operation

If we are to be clear about the responsibility of actors during the development of AI-SCS, we need understand the complex supply, development and operation chain for AI [11, 29]. To illustrate, the ML development process will typically include some or all of the following (this is not a complete list):

- **Training data** - this could be gathered and curated by multiple organisations and individual actors, such as anonymised health data [50] or publicly available datasets [30]
- **Software tools, libraries and services** - this includes off-the-shelf components produced in an open source environment, such as the Gazebo simulation environment [45] or PyTorch [15], and cloud services for data storage and performing the training.
- **Dynamic system change** - the ML may continue to self-adapt after deployment [29]
- **Informal requirements and incomplete verification criteria** - ML is often used for tasks that are difficult to formalise, e.g., "recognise road signs on UK roads in typical weather and light conditions", which make it difficult to train and verify across all scenarios it will encounter [4, 20]
- **Changing and dynamic deployment environment** - the AI-SCS can be deployed in a continuously changing environment, e.g., in the case of autonomous vehicles [9].

Hence, during development and operation of an AI-SCS there will be multiple developers, engineers, suppliers of components and data, project managers, operators, investors, regulators etc. and it is infeasible to document and link all decisions they make which could contribute to safety. However, if we model key relationships and responsibilities with more clarity, we could understand and resolve some issues such as conflicts, duplicates, non-allocated responsibilities, and prevent safety problems in the future, as well as support accident investigation.

This "problem of many-hands" was introduced by Thompson in [54] and later studied for conventional high-criticality systems [39, 55]. It can be summarised as the problem that, because many individuals and groups of individuals contribute to decisions, activities and outcomes in complex networks and organisations, it is difficult even in principle to determine who is responsible for them. Thompson uses the cases of the Banking crisis in 2007, and Gulf Oil spill in 2010 to illustrate the difficulty in finding accountable persons [55]. Additionally, he notes that some may unfairly avoid blame and others unfairly take blame. To avoid similar incidents, Thompson proposes the use of additional bodies with oversight responsibility and ensuring conflicting responsibilities are assigned to different entities.

The problem of many hands, and the difficulties in identifying responsible agents for machine learning in a non-safety environment is explored in [11] by Cooper et al. and in [10] by Cobbe et al., noting the many different responsible agents. Cooper notes that the inevitability of ML bugs can be used to excuse responsibility for non-safety systems. We note that for SCS reducing bugs is standard practice to reduce risk, and should be evidenced as part of a SAC [19]. However, bugs will still remain despite being reduced systematically. Cooper also argues that responsibility should be designed into organisations while systems are developed. Using a role responsibility model, such as the one we propose, is one way this can be achieved for AI-SCS.

## 2.3 Legal pitfalls and due diligence

One key aspect of responsibility is that of legal responsibility and liability following an accident.

We previously noted engineers follow "good practice" for developing and operating traditional safety-critical systems, e.g., DO-178C for avionics [46], ISO-14971 [25] for medical devices, or ISO-26262 [24] for electronic automotive vehicle components. Not following these could lead to sanction, even if the standards are not legally required, and hence is related to responsibility. One

problem for AI-SCS is that there is limited consensus and little evidence of what constitutes good practice for developing safe ML [4], despite a large number of AI/ML safety standards [23], and huge volume of research in this area. Therefore, alternative means of demonstrating due diligence and good practice will be needed.

In the UK, the Health and Safety at Work etc. Act 1974 [3] and the subsequent Management of Health and Safety at Work Regulations of 1999 lay down the general legal obligations pertinent to ensuring the health and safety of both employees and people who may be affected by the activities of an organisation. The Act gives rise to criminal sanctions for organisational failures and emphasises the importance of proactive measures (i.e., risk assessments, health surveillance, and communication between employers and employees, amongst others) to ensure the overall health and safety of individuals. Based on the guidance issued by the Health and Safety Executive, proactive risk assessments play a vital role in safeguarding individuals by identifying potential harms and implementing measures to mitigate such risks from materialising. This practice is fundamental.

This emphasises the legal importance of due diligence practices. From a legal perspective, one way of describing the purpose of a SAC is a proactive means of justifying and explicitly detailing the roles and responsibilities of the various individuals involved in a system's lifecycle. This clarifies their duties and outlines their subsequent responsibilities. The SAC works in two ways: both forward-looking and backward-looking. The intention is to assess potential risks from the outset and identify the relevant roles responsible for mitigating such harms. However, in the event something goes wrong, the SAC will be used in legal proceedings to assess how and where the risk materialised. In complex negligence-based cases, of specific pertinence to AI-SCS, industry standards and expert evidence may be used to assess the scope of duties, any deviations from the standard, and negligent conduct. Published material or guidance forming the general corpus of knowledge in the field may also be used as a reference and can be the deciding factor in liability claims. Importantly, if a person actively and diligently assesses risks, implements safety procedures, and ensures the relevant guidelines are followed, the person is said to have taken reasonable measures to ensure the safety of their conduct. In the event something undesirable occurs, it can be argued it would be unreasonable to find someone who diligently follows good safety practice guidelines (and the law, more generally) liable – so long as they have evidently carried out their responsibility, as reasonably practicable.

## 2.4 Clarifying types of responsibility for AI-SCS

The previous discussions have touched on a number of different notions of responsibility but has not defined them precisely, nor characterised their relationships. Nor has it identified a practical means to review and analyse responsibility relating to AI-SCS, particularly for stakeholders who may be developing a SAC.

The work of Porter et al. in [43] draws upon Hart's classic taxonomy of the different senses of responsibility [18] to identify and clarify different types of responsibility (role, causal, legal and moral). This breakdown of types can be used to describe different aspects of responsibility relating to AI-SCS development and operation. Role responsibility refers to the tasks, specific duties, and obligations that attach to particular roles. Causal responsibility is another way of referring to causation. If one is causally responsible for something, one is just a cause of it or a salient causal factor in its coming about. Legal liability responsibility includes being required to pay financial compensation, or be subject to a legal order, or face punishment [44]. Moral responsibility (at the heart of much of the debate about responsibility gaps e.g., [34, 56]) concerns whether an actor deserves blame or praise for an outcome. This can overlap with the other senses of responsibility, such as role responsibility. Causal responsibility is generally taken to be necessary



but not sufficient for moral responsibility, but does not necessarily imply legal responsibility. One can be held vicariously responsible without causal contribution. Additionally, there is a distinction between moral accountability (where an actor is held to account) and moral attributability (where an actors contribution can be considered voluntary). This distinction is revisited in our examples (Sections 4 and 5).

In [43], Porter et al. present a formulation for representing responsibility as follows:

*Actor(A) is Responsible for Occurrence(O)*

The *Actor (A)* could be an AI-based system (such as a AI-SCS), an individual involved in development or operation, or an institution (e.g., company or organisation). *Occurrences* are characterised as either *Decisions*, *Actions* or *Omissions*. For each an asterisk (\*) is used to indicate where the *Occurrence (O)* is attributed to an AI-based actor. For example, *Automated Driving System (A) is (task)role-responsible for executing the dynamic driving task whilst activated\*(O)*. We have used this formulation as a basis for modelling responsibility types.

Additionally, we need perform effective analysis of our model. In [7, 32] the authors present a responsibility analysis for high criticality systems. Their analysis is for operational roles, considering duties and tasks of human actors. They also provide a notation for the analysis which includes responsibilities, resources, agents, and different types of relationships. In [32] the authors use guidewords to prompt review of failures across the relationships, such as *early*, *late*, *never*. In [7] (role) responsibility failure types including *unassigned*, *duplicated*, *not communicated* are used. We have adapted their notation and analysis approach (see Section 3), incorporating the types of responsibility described above, reinterpreting and adding to the guidewords and extending to include relationships.

## 2.5 Summary

In summary, research in understanding responsibility for AI-SCS, has identified an number of responsibility related issues. Reviewing and documenting role responsibility during design of AI-SCS could prevent incidents in the future, e.g., when there are conflicting or missing responsibilities that lead to safety risks. Our contribution is to demonstrate a practical means of modelling and analysing responsibility for development and operation of an AI-SCS, to clarify responsibility, and, if possible, reduce responsibility issues. The analysis can be used as part of a SAC to demonstrate due diligence and also potentially as defence following an accident. Further, it could be used to investigate and prevent similar issues in the future, and avoid unfair blame.

## 3 METHOD

In this section we describe our responsibility modelling notation, and give an overview of how to develop and analyse models. Sections 4 and 5 show two different examples of it being used.

### 3.1 Notation

Our notation is an adapted version of the task responsibility notation presented by Lock in [32], incorporating the responsibility concepts provided by Porter et al. in [43]. This combination ensures that the relationships between responsibilities can be modelled, whilst adding the specificity and nuance of the responsibility concepts from [43]. The model elements are shown in Figure 1. The three different actor types identified in [43] are presented (AI-enabled or AI-based system, individuals and institutions). What is referred to as *Occurrence* by Porter et al., and *Responsibility* by Lock is represented by a rectangle. Resources are represented using the standard flow chart symbol for documents. We include resources to allow us to model specific outputs from an *Occurrence*. Actors

may use a resource immediately (for example, a prediction) or later (for example, an Off the Shelf software library used during development).

The relationships between the elements are as follows:

- **(type) responsibility for** - this specifies variations on the four types of responsibility described in section 2.4
- **uses** - we have added this relationship to indicate where a resource or occurrence is used by another actor. This can add insight where problems with the resource (which may be due to prior issues undertaking the task to manage it) may have safety implications. We note that although a resource may be *used* it may not necessarily be *required* for an actor to complete their task.
- **subordinate to** - this represents where there is a power relationship or hierarchy between elements
- **association** - this is used where there is a non-specific relationship between elements, for example, a resource associated with an occurrence
- **acts as** - where one actor or element acts in place of another




Fig. 1. Responsibility model elements

We illustrate the concepts in Figure 2, where an AI-based system has the *(task)role* of monitoring whether there is a pedestrian on a potential collision path. If so it creates a collision warning (resource) which is used by the Safety Driver. The safety driver is responsible for intervening to try and prevent a collision when a warning is issued (if necessary).

The initial model is shown in the top line (A) where the actors and their tasks are described. Considering failures allows us to see the impact on actors, and how they may perform their role. For example, if the warning is late due to poor performance of the AI-based system in its monitoring task, this may mean a late intervention. In the bottom row (B), we show the responsibility contributions following a collision. The AI-based system has a causal link to the insufficient prediction and associated warning, whereas the safety driver may be blamed for late intervention. The AI-based system is not a legal person and is not considered to be a moral agent, hence cannot be held liable or morally responsible [42, 44]. However, the safety driver may be found morally accountable or even liable following the crash, even if the AI-based system causally contributed to their actions. This highlights a responsibility gap between a causal factor and the driver.

We can extend this model and analysis to include the actor(s) responsible for designing the AI-SCS, or who provided training data, etc., and build a chain of actors and relationships. This can help prevent issues of AI technology shortfalls by identifying tasks to reduce risk, as well as



clarifying roles of which manage and contribute to risk. Further, the model can be used to help demonstrate there was no responsibility shortfall within a SAC.

In a forwards-looking analysis there may be many different responsibility models, with different hypothetical outcomes, dependent on the issues uncovered. For example, the *Warning of potential collision\** in Figure 2 could be *Late?* or *Missing?*, with different consequences.

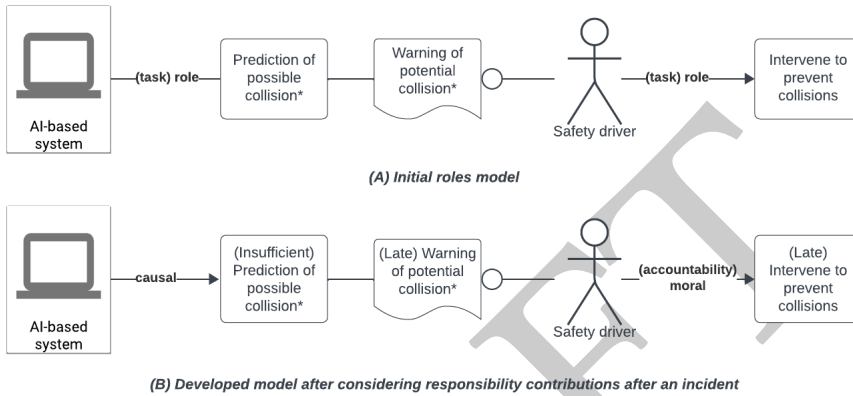


Fig. 2. Development of a simple responsibility model

### 3.2 Creating and analysing the model

There are two main phases for creating and analysing the model. First, information on the elements needed to create the diagram is gathered, concentrating on identifying relevant actors and roles. Second, an iterative analysis of the model is performed either during development or after an incident. These phases should be performed by teams of experts with appropriate specialist knowledge to the AI-SCS (e.g., clinicians, ML software engineers, safety experts, legal experts).

**3.2.1 Gathering information.** The responsibility model is built with the following steps:

- Identify key elements:
  - List key operational and development actors (human, institutions and AI-based)
  - List key resources and occurrences used during operation: such as system status outputs, information, interacting tasks, compliance and<sup>2</sup> legal duties
  - Using the previous data, list additional development roles, resources and occurrences which relate to them
- Characterise the responsibility relationships between the elements
- Arrange elements into graphical model

A few comments on this process. If this is a forwards-looking model, we assume that the main responsibility types will be moral obligations and duties (for example, codes of conduct), legal duties (e.g., duty holders or legislative roles) and task roles (e.g., typical development activities). Backwards-looking models will include different responsibility types such as causal, moral attributions and liability, but we will speculate on these during analysis. As noted, review and analysis will support the model's development and iteration.

<sup>2</sup>Compliance duties relate to the use of safety standards and established "good practice". Whilst not legally mandated, lack of compliance could lead to sanction (see Section 2.3).

Each actor may embody different types of responsibility (leading to occurrences and resources), and each occurrence and resource could potentially have multiple related tasks. For example, a single resource could be managed by more than one actor, and in different ways, such as a car being built by a manufacturer and maintained by a dealer.

**3.2.2 Reviewing and Analysing the model.** The model will require review and analysis to determine potential issues, such as risk reducing tasks with no assigned actors or actors who have overloaded roles and a high burden of risk.

The initial aim of the analysis is to build confidence that all risk related roles have been identified and allocated, as shown in section 5. For a backwards-looking analysis, the model can help determine where role responsibility issues led to an incident or accident. An example is shown in section 4.

Some established safety analyses, such as Hazard and Operability Studies (HAZOPS) [32, 35, 48] use guidewords to suggest failure conditions in SCS, e.g., "Lack of" meaning a communications message is not sent, or "too much" meaning the addition of too many chemicals to an industrial process. As these guidewords are proven, we use these when considering technical issues with resources.

However, no established method exists for our responsibility analysis. Again, taking inspiration from the work of Lock [32] and Baxter [7] we have developed a set of guidewords used to prompt discussion and identification of issues for occurrences. These are now described, with examples of how they might prompt discussion.

- **Insufficient/Partial** - where an Occurrence has happened but it doesn't meet the requirements in full
- **Misassigned** - where it is not appropriate to assign a particular Occurrence to a particular actor
- **Overloaded** - where an actor may have too many tasks, or they are too burdensome
- **Duplicated/Conflict** - this indicates two or more actors performing the same or similar occurrence, noting that they may conflict
- **Missing/Never/Omission** - where an occurrence is not provided
- **Early/late/ordering** - where an occurrence is provided but there is an issue with the sequencing which causes a problem
- **Incorrect/Value** - where there is some other issue with the occurrence, such as it being performed incorrectly

For each of the occurrences and related resources, the safety impact of an issue needs to be assessed. For example, an actor may not be appropriately trained to perform a task (*misassigned*), meaning their outputs would not be sufficient, or it may not be assigned (*missing*). Another scenario is that we cannot provide guarantees in the quality of a resource (*insufficient* e.g., where there are many actors i.e., many hands, producing open source software). The analysis may uncover additional tasks which either alter existing occurrences/resources or create new ones. For example, an additional assurance task needed to assure software should be to be assigned to an appropriate actor, or a duplicated task may need clarification of who should perform it.

## 4 EXAMPLE 1 - FATAL COLLISION WITH AN AUTONOMOUS VEHICLE IN TEMPE, ARIZONA

In this section we present a backwards-looking analysis to illustrate the modelling concepts and methodology in more detail. The example is the fatal collision of an Uber Advanced Technologies Group (ATG) vehicle with an automated driving system (ADS) and a pedestrian pushing a bicycle across a highway in Tempe, Arizona in 2018. The National Transportation Safety Board (NTSB) issued an accident report [38] describing contributory factors, and highlighting a number of areas

where there were safety shortfalls. Uber ATG have since published a report describing improved safety processes to address some of the shortfalls identified [5]. A detailed sociotechnical analysis of the accident can be found in [33] which points at many undesirable organisational factors such as undue pressure to deliver and update a working solution, learning lag when reviewing in service data due to too much information, push for technical capability constraints to prioritise normal performance over safety, and over-reliance on the safety driver's reliability over long periods of time monitoring for hazards. These reports were the main sources of information for our analysis.

Uber ATG were responsible for developing the ADS, which was an adaptation added to a Sport Utility Vehicle (SUV) from a third-party supplier (Volvo). The ADS used a number of ML components to detect and classify objects in the vehicle's path, and also to predict those objects trajectories. The findings of the accident report noted that the classifier failed to consistently categorise the object as a pedestrian with a bicycle. Each time the classification changed, information on the pedestrian's path was lost, and so the ADS could not predict the pedestrian's path correctly to provide a collision warning. The safety driver in the vehicle was trained to disengage the ADS in emergency situations and take avoidance action. She was found to be distracted during the time of the accident. It was noted that there was inadequate regulatory control over operation of an AV.

In 2019, prosecutors said Uber ATG was not criminally liable in the crash. Uber ATG also settled the civil case brought by the victim's family out of court. Uber ATG preferred to use the word "resolved", in place of settled, as they did not admit liability [16], although settlements which compromise a dispute often contain provisions that a party does not accept liability. In our opinion, the out of court settlement represents a missed opportunity to test important questions of civil liability for incidents involving AI-SCS.

#### 4.1 Producing and analysing the collision responsibility model

In previous work [47] we presented an analysis of the Uber Tempe accident looking at each responsibility type and related actor individually. Our work summarised the different safety controls and mitigations which were in place at the time of the accident, their effectiveness, and who was assigned the role of implementing them either for development or operation. Since publishing the legal case against the safety driver has been completed, as she pleaded guilty to "endangerment" [22, 49, 61]. Here we include an updated analysis with the legal outcome, introducing our modelling notation and showing formal relationships between actors. In Table 1 we list some of the report findings, relevant roles, their compliance and legal duties, whether these duties could be a causal factor in the accident and their ultimate liability. We present two models of role responsibility, illustrating some relationships to visualise the relationships between the different roles which led to the accident. The first model shows the actors, and linked role responsibility tasks prior to the incident (Figure 3), and the second model shows the findings after the accident (Figure 4).

The models were developed as per the method outlined in section 3, starting with an initial model that includes the different actors, and their roles and relationships, prior to an analysis. We have modelled their relationships based on our review of accident findings. The initial model shows a subset of the actors with the tasks assigned to them in Figure 3.

A few comments on the initial model (Figure 3). The production of a safety assessment was voluntary in accordance with the regulatory framework, but Uber ATG were *subordinate to* the regulator and so we have presented their relationship as such. Compliance and legal obligations should be considered for every AI-SCS responsibility model due to their importance in development and operation, and specifically due to grey areas around legal liability of AI based systems (see section 2.3).

We have modelled *Ensuring just safety culture* as a (*moral obligation*) role. This reflects the obligation to embed safety concerns within the organisation, in accordance with general societal

Table 1. Example Uber ATG Tempe accident responsibilities following legal proceedings conclusion [38, 47]

Finding	Related Roles	Compliance duties	Legal ties	Du-	Causal Contrb	Liability
Lack of driver attentiveness	Safety driver	Compliance training	to	Yes	Yes	Pleaded guilty to charge of "endangerment" [22, 61]
Limited monitoring of driver attentiveness	Uber ATG	None identified		None	Yes	Not tested
Emergency response	re- Emergency services	N/A		Yes	None	None
Pedestrian behaviour	be- Third party	N/A		Illegally crossed road	Yes	None
Limited risk analysis of experimental systems	Uber ATG	Established good practice for risk assessment	good risk	None known	Yes	Not tested.
Deactivated automated braking system	Uber ATG	Established good practice for risk assessment	good risk	None known	Yes	Not tested
Lack of safety culture	Uber ATG	Established good safety practice	good	None known	Yes	Not tested
Lack of state/ federal oversight for AVs	Regulator/ lawmakers	Not available		Not available	Yes	None

expectations and norms. No individual actor is named as responsible, as it is a collective obligation, although we note that individual degree of responsibility may depend on authority and seniority [12][40]. Other actors have *(task) roles* related to development (e.g., *risk assessment*) and operational occurrences and resources (e.g. *Monitoring safety driver attentiveness* and *Warning of collision* from the ADS). They are not named individually in public materials, so Uber ATG is presented as an institution. The ADS and safety drivers are presented as separate actors as they have distinct functional role during operation.

We have not included the pedestrian as an actor in this model. We could include third parties impacted by the *occurrences*, for example, near misses or collisions with other road users, loss of jobs for drivers due to autonomous taxis. As these actors have no task in this situation, we have not included them. Arguably, other road users do bear some responsibility though, e.g., to follow the rules of the road. A future consideration is how to effectively incorporate impacted third parties where their *(task)role* may not directly interconnect.

Our model has been adapted in Figure 4 to highlight some of the findings of the accident report and other publicly available material. First, a note of a guideword encapsulating a shortfall is included in each of the *occurrences*, e.g., *Warning of collision* becomes *(Late) Warning of collision*. Second, where appropriate we have updated the relationship arrow annotations to reflect the findings, e.g., *(task) role* becoming *causal*.

Again, we have a few comments on this model. In order for someone to be held either morally accountable or attributable to an *occurrence*, it is considered a necessary condition for them to both

have moral agency and some causal contribution [18, 43]. The ADS had a causal contribution, via *(Late) Warning of collision* but has no moral agency.

In accordance with Porter et al. in [43], we can consider "moral responsibility as attributability ... when [actor] has voluntarily performed or voluntarily caused [occurrence]", and wasn't under extreme pressure or duress. With this in mind we consider that there is an *(attributability) moral responsibility* relation with the *(Insufficient) Ensuring just safety culture* for Uber ATG. However, we have annotated a *causal responsibility* link to *(Insufficient) risk analysis of experimental systems*. McCrae [33] notes that there was both pressure to deliver (based on fears that the organisation's existence was at stake) and staff were disempowered or discouraged from speaking about safety concerns. Hence, we consider the insufficient risk analysis was possibly not voluntary, but was arguably due to significant pressure. As noted, Uber ATG are represented as an institution, with no specified individual actors as this information is not available. This is an example of the many hands problem (section 2.2) where we don't have traceability to individual decisions or actions.

We also consider accountability, i.e., where someone faces the reactions and responses of the wider community for their actions. This includes blame or praise, for example, the emergency services were praised as providing "appropriate and efficient emergency medical response" [38]. From the finding of the legal case [22, 49, 61] and the accident report, we show that the safety driver's role can be considered blameworthy as *(accountability) moral* for *(Insufficient) Monitoring for potential collisions* and the late intervention. We have added the *(criminal) liability* occurrence of *Endangerment* to our model to reflect the legal case.

The addition of criteria to our analysis method for determining issues such as accountability and attributability is an avenue for further research, as it may be highly subjective and based on relative importance of causal contributions.

Uber ATG were found not to have monitored the safety driver's awareness, due to underestimation of automation complacency, i.e., where the safety driver trusts an automated system (the ADS) to an extent that they pay less attention. Also, there was a complex chain of related occurrences from problems with risk assessment, impacting on the ADS performance, leading to a late warning from the ADS to the driver. This removed a layer of design mitigation that should have reduced the risk of operating the vehicle, and reduced the cognitive load for the safety driver. A more prompt warning would provide mitigation against automation complacency, given more time for the safety driver to intervene, and also allowed for autonomous action to avoid collision. Arguably, this demonstrates where the burden of risk has increased on the safety driver due to shortfalls in how other actors performed in their roles.

The second model (Figure 4) illustrates many nuances of responsibility, and how the conduct by which some *(task) roles* were performed led to an accident via relationships with other tasks. We note though, that our analysis is based on pre-existing knowledge of the accident. In the next section we explore whether the method can be used identify potential responsibility problems without prior knowledge, and support a SAC justifying they have been adequately managed.

## 5 EXAMPLE 2 - DIABETES CO-MORBIDITY CLINICAL DECISION SUPPORT AI-SCS

Our second example is an AI-based diabetes co-morbidity predictor (DCP) the authors of this paper are developing for treatment of patients with Type II diabetes [41, 48]. The AI predicts a patient's risk of developing a diabetes co-morbidity or having a potentially catastrophic event, such as a heart attack, within the next six months. It is a decision-support system used during a patient consultation by the clinician, and has been built using an ensemble of ML components [41] to process recent patient test data and make a prediction of "High" or "Low" risk of a particular condition. The clinician can choose not to act on predictions made by DCP if they disagree with its assessment or if they have other information, for example, about the patient's preferred treatment,

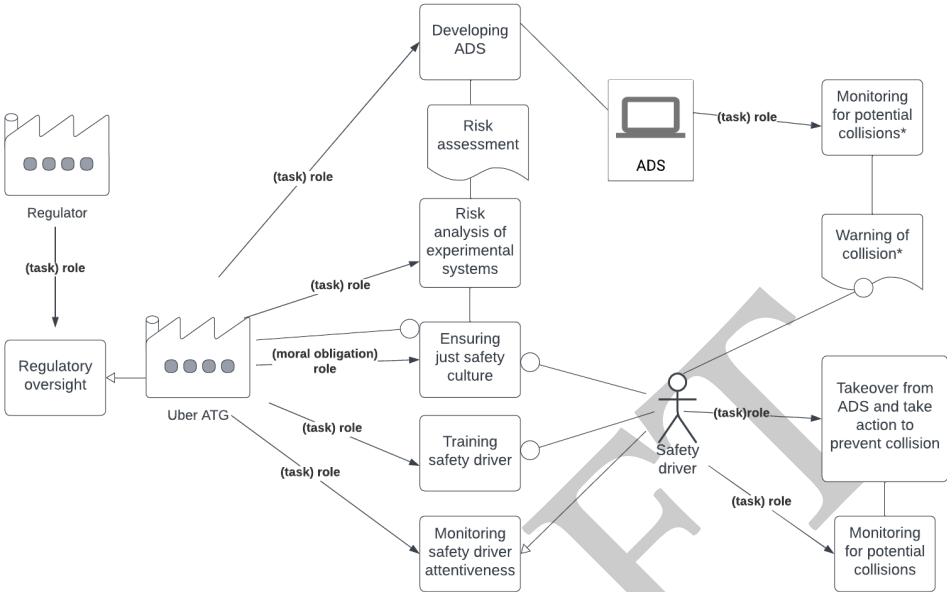


Fig. 3. Initial responsibility model showing actors and roles from Uber ATG crash in Tempe

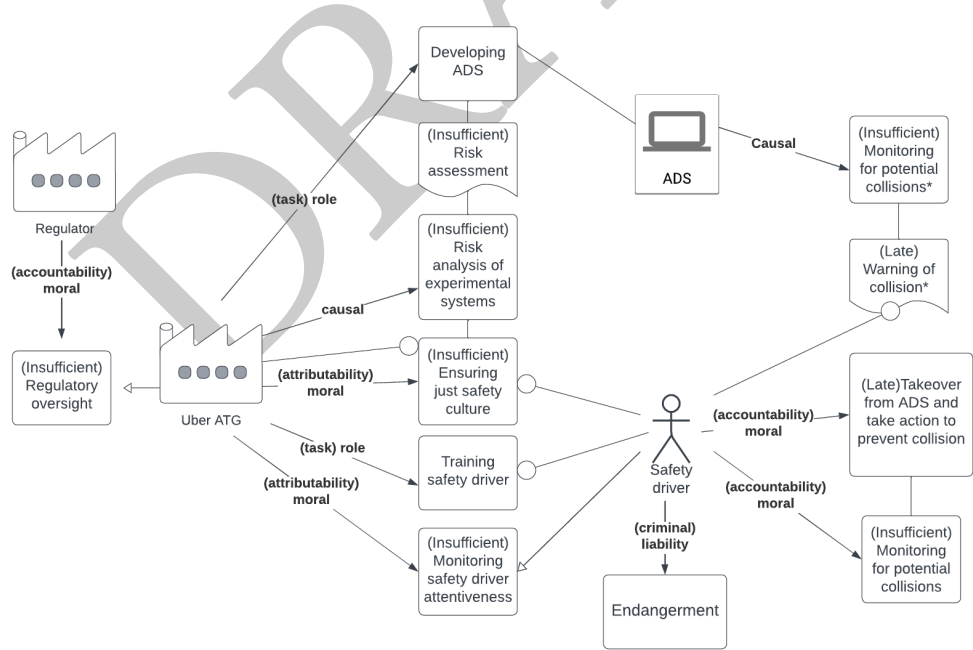


Fig. 4. Extended Uber Tempe responsibility model highlighting some responsibility findings from Uber ATG crash in Tempe [38]






Fig. 5. DCP explainability example showing influence of FOIs [48]

which make alternative courses of action more appropriate. The DCP is intended to function as an independent ‘second opinion’ during consultations. Previously, we identified that there are many complex issues arising from the public data sets used for training [50], which could lead to misleading outputs from the system [48]. These can be partially mitigated by careful curation of the training data by the AI developer, however the impact of some issues (e.g., bias due to lack of data from certain patient demographics) cannot be easily reduced, nor can we exhaustively test the AI as there isn’t a formal specification for required behaviour.

The DCP additionally provides explanation data. This lists the patient data features of importance (FOIs) for a particular prediction with values indicating their influence on the output recommendation. An example is shown in Figure 5, where Body Mass Index (BMI) is shown as the most influential factor for hypertension diagnosis.

5.1 Producing and analysing the DCP responsibility model

5.1.1 Producing the initial responsibility model. This section discusses development of the initial model (Figure 6). As discussed in section 2, we include the AI development lifecycle in our model to understand the impact of development issues on operational risk. This means we have included actors influencing the ML development process, regulators, and operating roles. Then we have considered each of their outputs as resources and/or occurrences.

The initial model is in Figure 6. A number of the human actors (*Clinical Staff*) have the qualifier 1..N. This indicates where there are many actors performing this role, which we do not have detailed information for, and they are not attached to a specific institution. This is another example of the many hands issue discussed in section 2.2. For example, the *Training patient database* is a resource that has been contributed to over many years, by different clinics with different unknown staff, each performing a *Maintaining patient database* occurrence. The *AI software tools* have also been developed by many different actors; in this case we assume they have come from specific software development institutions, for example, commercial tools. We could extend the model to consider each of these individually if valuable.

For regulatory oversight we have listed the health regulator (e.g., Medicines and Healthcare products Regulatory Agency (MHRA) in the UK) and National Institute for Health and Care Excellence (NICE). MHRA provide approval for medical device safety assessments, and NICE

provide clinical guidelines which we used to assist in determining the high/low risk of different co-morbidities. We have a node for *AI development good practice*, which does not currently have an associated actor with the task of development, due to the lack of established AI specific medical device guidance. There are, for example, guidelines being developed by the U.S. Food and Drug Administration (FDA), but these are in draft and not proven [13].

The operational actors (right hand side of the diagram), include *Clinical Staff* responsible for maintaining non-electronic and electronic records for the patient. Additionally, we include the clinician performing a consultation with a patient, and the DCP. Unlike the previous example in section 4, the impacted party is included, i.e., the type II diabetes patient, as they have a specific role during consultation. Here, the clinician will interact with the patient, asking them questions about their condition and general health. We also include that the clinician has a (*moral obligation*) role for *Duty of care*.

**5.1.2 Analysing and developing the model.** In this section we demonstrate how a forwards-looking analysis of Figure 6 helps develop the revised model (Figure 7). As described in section 3, we need to consider hypothetical issues with both occurrences and resources, and think about how they may impact on safety of the AI-SCS. This should consider safety impact on the performance of the AI-SCS and on interacting actors. Issues may include, missing occurrences, actors with conflicting or overloaded roles, insufficient or late resource provision and so on. An extract of the HAZOP style analysis is shown in Table 2.

When applying the guidewords we have proposed some additional tasks (i.e., occurrences) to mitigate or prevent potential safety issues. For example, in the highlighted path in Figure 7, issues with the training data have *causal* impact on performance of the DCP. This is shown as eventually (potentially) contributing to (*Incorrect*) *Clinical Decision and Treatment*. Mitigation measures can be assigned to the AI Developer, by improving how data is processed and extracted prior to training. In Figure 7, we hypothesise that by not performing task they are (*(attributability) moral*) responsible for (*Insufficient*) *Training and Assurance of AI*. However, the clinician would likely be held accountable (*(accountability) moral*) for the incorrect diagnosis. The Type II diabetes patient would be affected. As previously discussed, stricter criteria to determine issues like moral accountability and attributability are an avenue for further research.

We need to assign the task of assuring the AI software tools. The chain of impact of bugs in tools is similar to that for training data quality issues. Bugs could remain undetected throughout development and into operation, affecting the predictions and explanations from the DCP, and thus impacting the decision making of the clinician. There are several reasons for this. First, the software tools may be complex and not designed for safety-critical use (for example, generic library tools used to manage and alter data, perform training of multiple neural network models etc.), making it difficult to assure effectively. Second, there may be no access to source code or documentation. Third, we cannot perform complete testing of the DCP to find bugs, due to the enormous variation of potential patient data and uncertain requirements for accurate predictions. A detailed causal impact analysis, considering impact, size of contribution and likelihood, for these chains forms an avenue for further research. As previously noted, several models may be needed for different outcomes.

Another potential issue identified by applying the guidewords, is whether explainability information provides useful insight into the AI prediction for the clinician [27] [26]. As discussed, a visualisation of the relative influence of FOIs on the prediction can be provided (see Figure 5). However, some FOIs could influence the DCP prediction in similar proportions whether it indicates high or low risk of a heart attack, or be close to NICE guideline boundaries. The clinician may choose not to use explainability data if it is ambiguous.

Occ.	Resource(s)	(task) role Actor	Uses Actor	G/word	Issue	Mitigation
Develops tools	SW tools	Software supps (1..N)	DCP developer	Insuff.	Tool quality is unknown. Bugs impact on DCP performance.	Perform assurance assessment task to assess potential issues
Maintain database	Training dataset	Clinical staff (1..N)	AI developer	Insuff.	Data poorly distributed, missing values. DCP outputs are insufficient, e.g., perform poorly on patients matching missing elements.	Perform training data quality assessment and compensate where possible
Training and as- surance of AI	Prediction from DCP	AI Dev.	Clinician	Insufficient	Inconsistent and misleading performance of DCP	Use explainability, follow up patient progress to assess DCP performance over time
Training and as- surance of AI	Prediction from DCP	AI Dev.	Clinician	Conflict	DCP provides FP/FN, impacts clinical decisions	Use explainability, follow up patient progress to assess DCP performance over time, prioritise clinical expertise
Training and as- surance of AI	Prediction from DCP	AI Dev.	Clinician	Conflict	DCP provides TP/TN but conflicts with clinician	Use explainability, follow up patient progress to assess DCP performance over time, prioritise clinical expertise
Clinical decision and treat- ment	N/A	Clinician	Type II di- abetes pa- tient	Incorrect	Wrong treatment recommendation due to influence of DCP prediction (FN/FP)	Patient discussion, use explainability, prioritise clinical expertise
Generating explana- tion	Explanation data	DCP	Clinician	Insuff.	The FOI data adds limited insight into the prediction	Patient discussion, prioritise clinical expertise

Table 2. Extract of the model HAZOP analysis for the DCP

Table 2 proposes that incorrect (i.e., false negative/false positive) and conflicting predictions from the DCP are mitigated by the Clinician, by discussion with the Patient, other information, and use of their clinical expertise. This has added additional tasks, potentially overloading them and extending consultation times. Additionally, as the clinician is mitigating many potential problems with the AI partially introduced by other actors, we could argue their burden of risk has increased, rather than being reduced, by the use of the DCP. Again, we are concerned about them becoming a "liability sink" in the future. The model and analysis provides detail and clarity on where additional risks have come from.




Fig. 6. DCP (task)role relationships

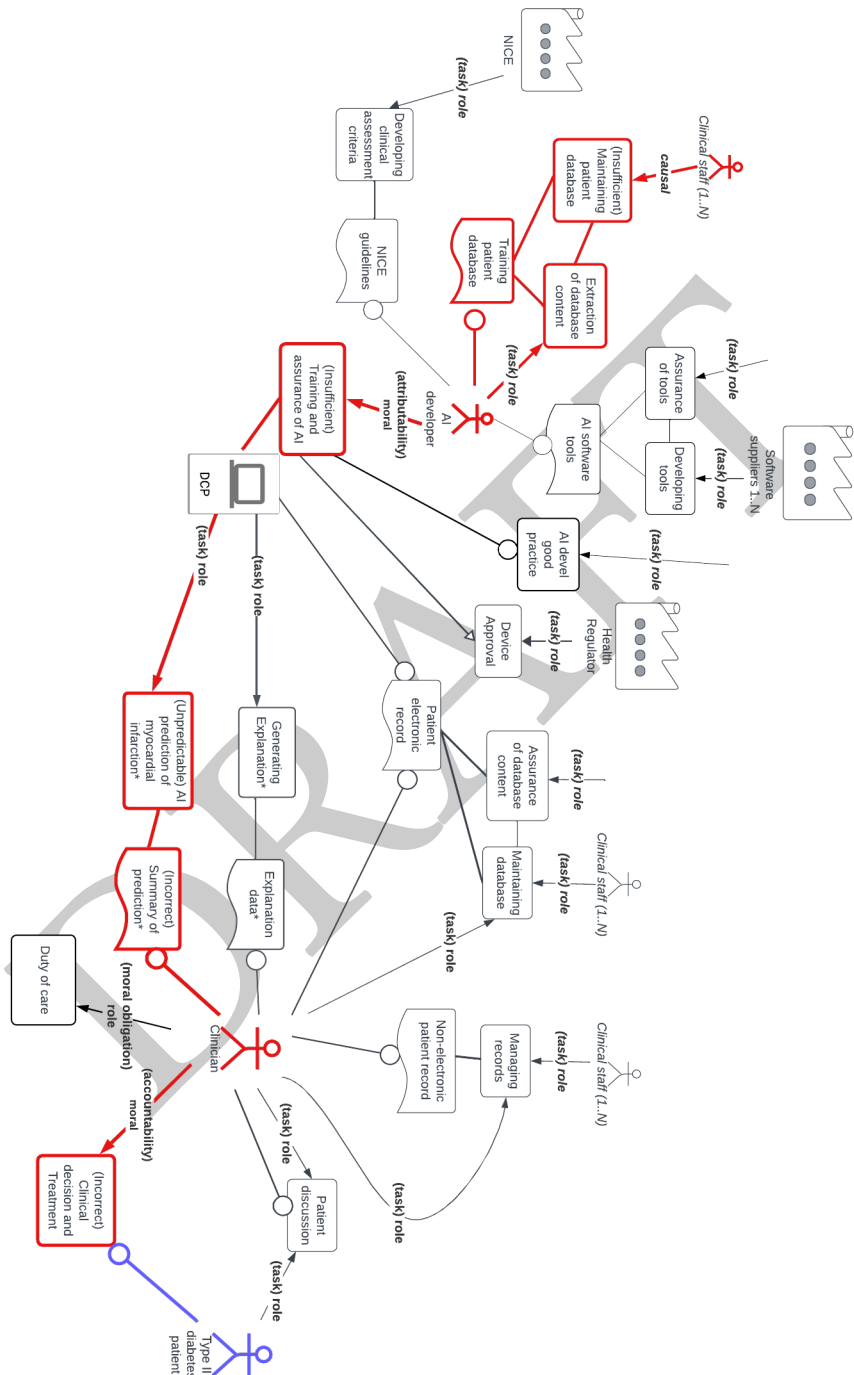


Fig. 7. Revised DCP model (red highlighted path showing operational safety impact of training data errors)

## 6 DISCUSSION AND FUTURE WORK

There are many concerns around responsibility relating to the development and operation of AI-SCS, including "responsibility gaps", where it is difficult to ascribe blame or determine liability for unsafe behaviour of an AI agent, "the problem of many hands" where it is hard to identify responsible individuals, and "liability sinks" where operators have increased burden of risk when using an AI-SCS. We argued that modelling and analysing responsibilities during design of AI-SCS provides the potential to reduce risk, e.g., by ensuring there are no conflicting or missing responsibilities. We also hope to avoid unfair or disproportionate blame being placed on operators or developers.

We presented a practical method for responsibility modelling which can support a SAC, required for most AI-SCS. We presented two different examples, demonstrating how modelling responsibility provides clarity on different roles and responsibilities in development of a AI-SCS, as well as the impact on operational safety, and the potential risk burden placed upon operators.

We first demonstrated how to develop a responsibility model, using prior knowledge of the Uber Tempe fatal collision in 2018. We used our notation and analysis method to produce a nuanced model showing different types of responsibility, such as moral attributability, accountability and liability. Additionally, we developed a picture of how responsibility issues (from multiple actors) during development led to operational shortfalls, ultimately contributing towards the accident.

The second example of an AI-based medical decision support tool, showed how the method can be used as a predictive analysis tool, proposing a number of additional risk reduction tasks, and illustrating the risks being mitigated by a clinician. To validate our work further, we need to demonstrate that our model and analyses have helped identify, and found means to alleviate, real safety responsibility issues. To a certain degree this has been accomplished via expert and peer review of our work. Additionally, we uncovered some missing roles and tasks, and also demonstrated how the burden of risk on the clinician may be increased.

For additional validation we could revisit the DCP responsibility model following an incident or accident. We could test whether the model and analysis had uncovered relevant issues and/or provided valuable evidence relating to responsibility. This is not practical, or ethical, to pro-actively pursue, but it remains a possibility if such an event occurs.

As future work will develop the analysis methodology, including criteria for determining responsibility types such as moral accountability and liability. We will explore causal contribution in more depth, including severity and likelihood of impact, which also relates to accountability and liability. Inclusion of third-party actors with no specific role is also a future consideration. We additionally plan to develop a SAC framework incorporating the models.

This work is potentially generalisable to other SCS, but our focus was on AI-SCS due to the known issues around responsibility.

## ACKNOWLEDGEMENTS

This work was supported by the Engineering and Physical Sciences Research Council (EP/W011239/1) and the Assuring Autonomy International Programme, a partnership between Lloyd's Register Foundation and the University of York. We would like to thank Paul Noordhof (Philosophy Department, University of York), Philip Morgan (York Law School, University of York), Tom Lawton (Bradford Teaching Hospitals) and Berk Ozturk (AAIP, University of York) for their invaluable insight and comments.

## REFERENCES

- [1] April 2021. Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain Union Legislative Acts. <https://artificialintelligenceact.eu/wp-content/uploads/2022/05/AIA-COM-Proposal-21-April-21.pdf>.



- [2] February 2023. Artificial intelligence liability directive briefing. [https://www.europarl.europa.eu/thinktank/en/document/EPRS\\_BRI\(2023\)739342](https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2023)739342).
- [3] UK General Public Acts. 1974. Health and Safety at Work etc. Act 1974. <https://www.legislation.gov.uk/ukpga/1974/37/contents>.
- [4] Rob Ashmore, Radu Calinescu, and Colin Paterson. 2021. Assuring the machine learning lifecycle: Desiderata, methods, and challenges. *ACM Computing Surveys (CSUR)* 54, 5 (2021), 1–39.
- [5] Uber ATG. 2018. Safety Report Supplement Internal and External Safety Reviews. [https://aurora-dev.cdn.prismic.io/aurora-dev/4f3e03fe-7d41-4bed-a998-7b2d918b9579\\_UberATGSupplementSafetyReview2018.pdf](https://aurora-dev.cdn.prismic.io/aurora-dev/4f3e03fe-7d41-4bed-a998-7b2d918b9579_UberATGSupplementSafetyReview2018.pdf)
- [6] Lisanne Bainbridge. 1983. Ironies of automation. *Automatica* 19, 6 (1983), 775–779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
- [7] G. Baxter and I. Sommerville. 2011. Responsibility Modelling For Resilience. In *PROCEEDINGS OF THE FOURTH RESILIENCE ENGINEERING SYMPOSIUM*. 22–28. <https://doi.org/10.4000/books.pressmines.960>
- [8] A. J. Blyth, Chudge J., J. E. Dobson, and M. R. Strens. 1993. ORDIT: A new methodology to assist in the process of eliciting and modelling organisational requirements. In *Proceedings on the Conference on Organisational Computing Systems*. 216–227.
- [9] Simon Burton, Ibrahim Habli, Tom Lawton, John McDermid, Phillip Morgan, and Zoe Porter. 2020. Mind the gaps: Assuring the safety of autonomous systems from an engineering, ethical, and legal perspective. *Artificial Intelligence* 279 (2020), 103201.
- [10] Jennifer Cobbe, Michael Veale, and Jatinder Singh. 2023. Understanding Accountability in Algorithmic Supply Chains. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 1186–1197. <https://doi.org/10.1145/3593013.3594073>
- [11] A. Feder Cooper, Emanuel Moss, Benjamin Laufer, and Helen Nissenbaum. 2022. Accountability in an Algorithmic Society: Relationality, Responsibility, and Robustness in Machine Learning. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 864–876. <https://doi.org/10.1145/3531146.3533150>
- [12] Sidney Dekker. 2018. *Just Culture: Restoring Trust and Accountability in Your Organization, Third Ed.* CRC Press.
- [13] U.S. Food and Drug Administration. Jan 2021. Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan. <https://www.fda.gov/media/145022/download>.
- [14] Office for Nuclear Regulation. 2023. Allocation of Function Between Human and Engineered Systems. NS-TAST-GD-064.
- [15] The PyTorch Foundation. [n. d.]. Pytorch. <https://pytorch.org/>. Accessed August 2023.
- [16] J Fulbrook. 2021. Reverberations from Uber v Aslam in personal injury claims? *Journal of Personal Injury Law* 2 (2021), 59–67.
- [17] David J Gunkel. 2012. *The machine question: Critical perspectives on AI, robots, and ethics.* mit Press.
- [18] H. L. A. Hart. 2008. 210POSTSCRIPT: RESPONSIBILITY AND RETRIBUTION. In *Punishment and Responsibility: Essays in the Philosophy of Law*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199534777.003.0009> arXiv:<https://academic.oup.com/book/0/chapter/160311870/chapter-pdf/38995148/acprof-9780199534777-chapter-9.pdf>
- [19] Richard Hawkins. [n. d.]. The Principles of Software Safety Assurance. <https://api.semanticscholar.org/CorpusID:14955881>
- [20] Richard Hawkins, Colin Paterson, Chiara Picardi, Yan Jia, Radu Calinescu, and Ibrahim Habli. 2021. Guidance on the Assurance of Machine Learning in Autonomous Systems (AMLAS). arXiv:2102.01564 [cs.LG]
- [21] IEC. 2010. IEC-61508 - Functional Safety of Electrical/Electronic/Programmable Electronic Safety-related Systems.
- [22] Independent. 2023. The backup driver in the 1st death by a fully autonomous car pleads guilty to endangerment. <https://www.independent.co.uk/news/uber-ap-phoenix-maricopa-county-tesla-b2384012.html>.
- [23] Alan Turing Institute. 2023. AI Safety Standards Hub. <https://aistandardshub.org/ai-standards-search/>. Accessed November 2023.
- [24] ISO. 2018. ISO-26262 Road Vehicles Functional Safety.
- [25] ISO. 2019. ISO-14971 Medical devices. Application of risk management to medical devices.
- [26] Yan Jia, John McDermid, Tom Lawton, and Ibrahim Habli. 2022. The role of explainability in assuring safety of machine learning in healthcare. *IEEE Transactions on Emerging Topics in Computing* 10, 4 (2022), 1746–1760.
- [27] Yan Jia, John Alexander McDermid, Nathan Hughes, Mark-Alexander Suján, Tom Lawton, and Ibrahim Habli. 2023. The need for the human-centred explanation for ML-based clinical decision support systems. In *2023 IEEE 11th International Conference on Healthcare Informatics (ICHI)*. IEEE.
- [28] T Kelly. 1998. Arguing Safety, a Systematic Approach to Managing Safety Cases. *PhD Thesis, Department of Computer Science, University of York* (1998).

- [29] Edwards L. 2022. Regulating AI in Europe: four problems and four solutions. <https://www.adalovelaceinstitute.org/wp-content/uploads/2022/03/Expert-opinion-Lilian-Edwards-Regulating-AI-in-Europe.pdf>.
- [30] V7 Labs. [n. d.]. Open Datasets. <https://www.v7labs.com/open-datasets>. Accessed August 2023.
- [31] Tom Lawton, Phillip Morgan, Zoe Porter, Alice Cunningham, Nathan Hughes, Ioanna Iacovides, Yan Jia, Vishal Sharma, and Ibrahim Habli. 2023. Clinicians Risk Becoming "Liability Sinks" for Artificial Intelligence. (2023).
- [32] Russell Lock, Tim Storer, Ian Sommerville, and Gordon Baxter. 2009. Responsibility Modelling for Risk Analysis. *ESREL, ESRA: Eur Safety Reliab Assoc.* <https://doi.org/10.1201/9780203859759.ch152>
- [33] Carl Macrae. 2022. Learning from the Failure of Autonomous and Intelligent Systems: Accidents, Safety, and Sociotechnical Sources of Risk. *Risk Analysis* 42, 9 (2022), 1999–2025. <https://doi.org/10.1111/risa.13850> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/risa.13850>
- [34] A Matthias. 2004. The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and Information Technology* 6 (2004), 175–183. <https://doi.org/10.1007/s10676-004-3422-1>
- [35] J A McDermid, M Nicholson, D J Pumfrey, and P Fenelon. 1997. Experience with the application of HAZOP to computer-based systems. In *IEEE proceedings of the 10th Conference on Computer Assurance Systems Integrity, Software Safety and Process Security*. 37–48.
- [36] Helen E Monkhouse, Ibrahim Habli, and John McDermid. 2020. An enhanced vehicle control model for assessing highly automated driving safety. *Reliability Engineering & System Safety* 202 (2020), 107061.
- [37] Lauritz Munch, Jakob Mainz, and Jens Christian Bjerring. 2023. The Value of Responsibility Gaps in Algorithmic Decision-Making. *Ethics and Information Technology* 25, 1 (2023), 1–11. <https://doi.org/10.1007/s10676-023-09699-6>
- [38] National Transportation Safety Board. 2019. Collision Between Vehicle Controlled by Developmental Automated Driving System and Pedestrian, Tempe, Arizona, March 18, 2018, NTSB/HAR-19/03. <https://www.nts.gov/investigations/accidentreports/reports/har1903.pdf>
- [39] Helen Nissenbaum. 1996. Accountability in a computerized society. *Science and engineering ethics* 2, 1 (1996), 25–42. <https://doi.org/10.1007/BF02639315>
- [40] Tor-Olav Nævestad, Ingeborg Storesund Hesjevoll, and Ross Owen Phillips. 2018. How can we improve safety culture in transport organizations? A review of interventions, effects and influencing factors. *Transportation Research Part F: Traffic Psychology and Behaviour* 54 (2018), 28–46. <https://doi.org/10.1016/j.trf.2018.01.002>
- [41] Berk Ozturk, Tom Lawton, Stephen Smith, and Ibrahim Habli. 2023. Predicting Progression of Type 2 Diabetes using Primary Care Data with the Help of Machine Learning. In *Medical Informatics Europe 2023*.
- [42] P Morgan et. al. 2023. *Tort Liability and Autonomous Systems Accidents*. Edward Elgar Publishing, Ltd.
- [43] Zoe Porter, Joanna Al-Qaddoumi, Philippa Ryan Conmy, Phillip Morgan, John McDermid, and Ibrahim Habli. 2023. Unravelling Responsibility for AI. arXiv:[2308.02608](https://arxiv.org/abs/2308.02608) [cs.AI]
- [44] Z. Porter, A. Zimmermann, P. Morgan, J. McDermid, T. Lawton, and I. Habli. 2022. Distinguishing two features of accountability for AI technologies. *Nature of Machine Intelligence* 4 (Sept 2022), 734–736. Issue 9. <https://doi.org/10.1038/s42256-022-00533-0>
- [45] Open Robotics. 2023. Gazebo. <https://gazebo.org/>. Accessed November 2023.
- [46] RTCA/EUROCAE. 2011. DO-178C Software Considerations in Airborne Systems and Equipment Certification.
- [47] Philippa Ryan Conmy, John McDermid, Ibrahim Habli, and Zoe Porter. 2023. Safety Engineering, Role Responsibility and Lessons from the Uber ATG Tempe Accident. In *Proceedings of the First International Symposium on Trustworthy Autonomous Systems* (Edinburgh, United Kingdom) (TAS '23). Association for Computing Machinery, New York, NY, USA, Article 25, 10 pages. <https://doi.org/10.1145/3597512.3599718>
- [48] Philippa Ryan Conmy, Berk Ozturk, and Ibrahim Habli. 2023. The Impact of Training Data Shortfalls on Safety of AI-based Clinical Decision Support Systems. *SAFECOMP 2023* (2023).
- [49] D Shepardson. 2023. Backup driver in 2018 Uber self-driving crash pleads guilty. <https://www.reuters.com/business/autos-transportation/backup-driver-2018-uber-self-driving-crash-pleads-guilty-2023-07-28/>.
- [50] Kuldeep Sohal, Dan Mason, John Birkinshaw, Jane West, Rosemary R C McEachan, Mai Elshehaly, Duncan Cooper, Rob Shore, Michael McCooe, Tom Lawton, Mark Mon-Williams, Trevor Sheldon, Chris Bates, Megan Wood, and John Wright. 2022. Connected Bradford: a Whole System Data Linkage Accelerator. *Wellcome open research* 7 (2022), 26. <https://doi.org/10.12688/wellcomeopenres.17526.2>
- [51] Mark Suján, Dominic Furniss, Kath Grundy, Howard Grundy, David Nelson, Matthew Elliott, Sean White, Ibrahim Habli, and Nick Reynolds. 2019. Human factors challenges for the safe use of artificial intelligence in patient care. *BMJ health & care informatics* 26, 1 (2019).
- [52] Mark A Suján. 2023. Looking at the Safety of AI from a Systems Perspective: Two Healthcare Examples. In *Safety in the Digital Age: Sociotechnical Perspectives on Algorithms and Machine Learning*. Springer, 79–90.
- [53] Mark A Suján, Ibrahim Habli, Tim P Kelly, Simone Pozzi, and Christopher W Johnson. 2016. Should healthcare providers do safety cases? Lessons from a cross-industry review of safety case practices. *Safety science* 84 (2016), 181–189.

- [54] Dennis F. Thompson. 1980. Moral Responsibility of Public Officials: The Problem of Many Hands. *American Political Science Review* 74, 4 (1980), 905–916. <https://doi.org/10.2307/1954312>
- [55] Dennis F. Thompson. 2017. *Designing Responsibility: The Problem of Many Hands in Complex Organizations*. Cambridge University Press, 32–56. <https://doi.org/10.1017/9780511844317.003>
- [56] Daniel Tigard. 2021. There Is No Techno-Responsibility Gap. *Philosophy & Technology* 34 (09 2021). <https://doi.org/10.1007/s13347-020-00414-7>
- [57] UK Maritime and Coastguard Agency. 2022. Improving Safety and Organisational Performance Through A Just Culture . [https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment\\_data/file/286139/just\\_culture.pdf](https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/286139/just_culture.pdf).
- [58] UK Ministry of Defence. 2017. Safety Management Requirements for Defence Systems 00-056 Part 1, Issue 7. Online: accessed 2022.
- [59] Human Rights Watch. 2023. How the EU’s Flawed Artificial Intelligence Regulation Endangers the Social Safety Net: Questions and Answers. <https://www.hrw.org/news/2021/11/10/how-eus-flawed-artificial-intelligence-regulation-endangers-social-safety-net>.
- [60] Bradley W. Weaver and Patricia R. DeLucia. 2022. A Systematic Review and Meta-Analysis of Takeover Performance During Conditionally Automated Driving. *Human Factors* 64, 7 (2022), 1227–1260. <https://doi.org/10.1177/0018720820976476> PMID: 33307821.
- [61] Wired. 2023. The Legal Saga of Uber’s Fatal Self-Driving Car Crash Is Over. <https://www.wired.com/story/ubers-fatal-self-driving-car-crash-saga-over-operator-avoids-prison/>.
- [62] Vahid Yazdanpanah, Enrico Gerding, Sebastian Stein, Mehdi Dastani, Catholijn Jonker, Timothy Norman, and Sarvapali Ramchurn. 2022. Reasoning About Responsibility in Autonomous Systems: Challenges and Opportunities. *AI and Society* 38 (12 2022). <https://doi.org/10.1007/s00146-022-01607-8>

Received December 2023