Net-HARMS, AcciNet and SafetyNet: A new safety management toolkit for complex systems

Paul M Salmon[1], Neville A Stanton[2], Guy H Walker[3], Patrick Waterson[4] & Adam Hulme[1]

[1]University of the Sunshine Coast, Australia, [2]University of Southampton, United Kingdom, [3]Heriot Watt University, United Kingdom, [4]Loughborough University, United Kingdom

ABSTRACT

Risk assessment and accident analysis methods based on systems thinking are currently popular, but few can be used together in an integrated manner. This article describes and demonstrates the Systems Thinking Accident and Risk Toolkit (START) which comprises the Networked Hazard Analysis and Risk Management System (Net-HARMS) risk assessment method, the Accident Network (AcciNet) accident analysis method, and the Safety Network (SafetyNet) intervention evaluation method. The three methods were designed to be used in an integrated manner as part of organisational safety management activities. START is described and demonstrated via a case study focussed on autonomous vehicles. The findings highlight the benefits of integrating risk assessment and accident analysis activities, including how accident data can be used to strengthen risk assessment outputs, and how the efficacy of specific risk controls can be considered in accident analysis efforts. Practical guidance on using the methods is offered, as well as recommendations for future research and applications in practice.

KEYWORDS

Systems thinking, risk assessment, accident analysis, Net-HARMS, AcciNet, SafetyNet

Introduction

Risk assessment and accident analysis are critical and mandated components of safety management (Dallat et al., 2019; Hulme et al., 2019). Formal risk assessment involves the use of structured methods to proactively identify potential hazards that may create adverse outcomes during specific work tasks (Chemweno et al., 2018). Accident analysis is a component of accident investigation and involves the post hoc description and modelling of adverse events to identify contributory factors. It is acknowledged that risk assessment and accident analysis methods are inherently related and should be used in an integrated manner (Hollnagel, 2008). Most formal safety management systems incorporate both risk assessment and accident analysis processes and recommend their use as part of an integrated framework of methods (e.g. Li & Guldenmund, 2018).

Although a number of Human Reliability Assessment (HRA) methods developed during the 1990s and 2000s were designed to be used both proactively for error identification and retrospectively for error analysis (e.g. Shorrock and Kirwan, 2002), many of the risk assessment and accident analysis methods used in practice are fundamentally different and thus difficult to use in an integrated manner. Differences include the underpinning theory, modelling approach, taxonomy or classification schemes, and representation of outputs. Indeed, aside from the Systems-Theoretic Accident Model and Processes (STAMP; Leveson, 2004) methods (i.e., STPA and CAST), few fully integrated risk assessment and accident analysis methods exist. This is particularly the case for

'systems thinking' based risk assessment and accident analysis methods which are currently popular (Dallat et al., 2019; Hulme et al., 2019; 2021a).

Whilst the STAMP methods are popular with researchers, there are various barriers which prevent their use in practice (Stanton et al., 2019; Waterson et al., 2015). To fully realise the benefits of integrating systems thinking-based risk assessment and accident analysis methods, more exploration on the development and testing of integrated methods that are usable in practice is required. In this article we outline new systems thinking-based risk assessment, accident analysis, and safety intervention evaluation methods which were developed specifically to be used together in an integrated manner: the Networked Hazard Analysis and Risk Management System (Net-HARMS; Dallat et al., 2018); the Accident Network (AcciNet; Salmon et al., 2020); and, the Safety Network intervention evaluation method (SafetyNet). The methods are described and demonstrated in a case study application focussed on autonomous vehicle safety.

The Systems Thinking Accident and Risk Toolkit (START)

Net-HARMS, AcciNet, and SafetyNet were designed specifically to be used together in an integrated manner. Together they form the Systems Thinking Accident and Risk Toolkit (START) which was designed to support the implementation of systems thinking approaches during organisational safety management. The integrated toolkit is driven by two components that are shared across the three methods: a task network describing the work system of interest and a risk mode taxonomy which is used to prompt analysts to identify risks (Net-HARMS), contributory factors (AcciNet), and emergent properties associated with safety interventions (SafetyNet). The task network provides the description of the work system that is the basis on which to identify risks, identify contributory factors, and to test and refine new controls or safety interventions.

Net-HARMS

Net-HARMS (Dallat et al., 2018) is a systems and network theory-based risk assessment method that supports the proactive identification of risks. This is achieved through the use of task network of the work system under analysis and a risk mode taxonomy which analysts use to identify task and emergent risks. Task risks are defined as task specific risks that could occur when undertaking key work tasks. Emergent risks are defined as unexpected risks that could foreseeably emerge when tasks risks interact with one another. Net-HARMS was designed specifically to provide two key advances over existing methods: first, to enable analysts to identify risks across the overall sociotechnical work system, as opposed to sharp-end risks only, and second, to enable analysts to identify emergent risks that arise when different risks interact with one another (Dallat et al., 2018).

Applying Net-HARMS involves first developing a Hierarchical Task Analysis (HTA; Annett et al., 1971) that describes the overall work system. The HTA is then converted into a task network which shows the tasks required for safe delivery of the work in question along with the relationships between tasks. Task risks are identified by applying a risk mode taxonomy (Table 1) to each node within the task network. For risk modes that are deemed credible (i.e. could conceivably occur), the analyst provides a description of the risks and their consequences, ratings of their probability and criticality (low, medium or high), and suggested risk controls. Following task risk identification, the risk mode taxonomy is applied once more to identify emergent risks that are likely to arise during instances when the identified task risks occur and influence the conduct of other tasks. This is a critical feature that attempts to support the identification and management of new and unexpected risks that are created when performance elsewhere in the work system is sub-optimal.

AcciNet

The AcciNet method is based on three fundamental tenets of accident causation:

1. That all accidents are created by an interacting network of behaviours associated with multiple actors, both human and non-human, who reside across the sociotechnical system (Rasmussen, 1997);
2. That the interacting network of contributory factors involved in accidents includes 'work as imagined' (i.e. undertaken in line with procedures), 'normal performance' (i.e. 'work as done' where performance was appropriate and no discernible failure occurred), and decisions and actions whereby performance can reasonably be classified as sub-optimal (Dekker, 2011; Hollnagel, 2012); and
3. That emergent risks play a critical role in accident causation. Emergent risks occur when multiple behaviours interact with one another to create unexpected and difficult to foresee behaviours. These emergent behaviours occur across the sociotechnical system and may be proximal or distal to the accident event.

AcciNet was developed to be used in conjunction with Net-HARMS and specifically to enable practitioners to identify and depict the three tenets of accident causation described above. The method uses a task network for the system under analysis as its primary input, with analysts subsequently using the task network to identify contributory factors, their interrelations, and associated actors. The Net-HARMS risk mode taxonomy (Table 1) is then used to classify identified contributory factors. The use of the task network to identify contributory factors removes of the need for analysts to identify the relationships between contributory factors (as the relationships between tasks are already specified within the task network).

Table 1. Net-HARMS risk mode taxonomy.

Classification	**Risk Mode**	**Description**
Task	T1: Task mistimed	Task is undertaken too early or too late within process
	T2: Task omitted	Task is not undertaken
	T3: Task completed inadequately	Task is undertaken but is completed in an inadequate manner
	T4: Inadequate task object	The object(s) used to complete the task are inadequate
	T5: Inappropriate task	An inappropriate task is performed instead of the required task
Communication	C1: Information not communicated	Information required to complete the task is not communicated
	C2: Wrong information communicated	The wrong information is communicated
	C3: Inadequate information communicated	Information is communicated but is inadequate e.g. incomplete communication with missing information
	C4: Communication mistimed	Communication is undertaken too early or too late within process
Environment	E1: Adverse environmental conditions	Adverse environmental conditions influence task performance

SafetyNet

SafetyNet is a safety intervention analysis method designed to support organisations when developing and implementing new risk controls and safety interventions. SafetyNet is applied following Net-HARMS and AcciNet and involves an analysis of the likely impact of proposed

safety interventions via the use of the task network and classification scheme. This involves using the task network and a modified version of the Net-HARMS taxonomy to identify any positive and negative consequences of implementing risk controls and safety interventions within current organisational practice.

SafetyNet works by inserting new risk control and safety intervention nodes into the task network and then linking the nodes to existing system tasks. The SafetyNet taxonomy is systematically applied to the new nodes within the updated task network to identify potential positive and negative effects for each safety intervention once implemented in practice. Controls and interventions are subsequently refined based on the findings and tested further until analysts are satisfied that they will have the desired effect.

Uber-Volvo case study

On the morning of Sunday 18th March 2018, a Volvo XC90 Sport Utility Vehicle (SUV) fitted with Uber's self-driving system struck and killed a pedestrian on Mill Avenue in Tempe, Maricopa County, Arizona (NTSB, 2018). The vehicle was being tested as part of Uber's Arizona testing program and was occupied by a vehicle safety operator at the time of the collision. The test vehicle was travelling northbound along a two-lane road with a posted speed limit of 45mph and collided with a female pedestrian who was attempting to cross the road from a centre median strip.

At the time of the collision, the vehicle had been in self-driving mode for approximately 19 minutes and was negotiating its second loop of the test-route. According to the NTSB, the Uber system first registered radar and LIDAR observations of the pedestrian around 6 seconds prior to impact. The self-driving system initially classified the pedestrian as an unknown object, then as a bicycle, and was unable to identify its intended path. Around 1.3 seconds before impact, the self-driving system determined that an emergency braking manoeuvre was required. Due to the vehicle's City Safety system being disabled, it was not possible to initiate the required emergency braking manoeuvre (NTSB, 2018; 2019; Stanton et al., 2019). The vehicle safety operator noticed the pedestrian and intervened less than a second before impact. Whilst she initially engaged the steering wheel, she did not brake until after the impact with the pedestrian, with recent accounts suggesting that she was distracted at the time by a streaming service on a mobile phone (Stanton et al., 2019).

Method

Four authors (PS, AH, GW, NS) developed a task network for the task 'Conduct on-road autonomous vehicle testing program' and used it as the basis on which to conduct Net-HARMS, AcciNet, and SafetyNet analyses. The primary sources of information for the task network were two NTSB reports; however additional documentation was also used including academic articles describing the incident (e.g. Stanton et al., 2019).

One author (PS) used the task network and Net-HARMS taxonomy to identify potential task and emergent risks for the 'Conduct on-road autonomous vehicle testing program' system. The Net-HARMS analysis was subsequently reviewed by a second co-author (AH) and any disagreements were discussed until consensus was achieved.

An AcciNet workshop was held to identify relevant contributory factors deemed to have played a role in the collision. All authors reviewed the task network along with the NTSB reports and identified contributory factors that played a role in the collision. Each contributory factor was then coded into a contributory factor type using the Net-HARMS risk mode taxonomy. Due to space constraints the SafetyNet analysis is not presented in the current article.

Results

Task network

The 'Conduct on-road autonomous vehicle testing program' task network is presented in Figure 1. Tasks are represented by nodes and relationships between tasks are represented via arrows linking the nodes. Tasks are deemed to be related with one another if the conduct of one task influences or is dependent on, another task or if tasks are undertaken together. For example, the tasks 'Train vehicle safety operators' and 'Drive vehicle' are linked as the Uber test vehicle safety operators could only drive the test vehicle once they have completed the vehicle safety operator training program.

Net-HARMS analysis

An extract of the Net-HARMS task risk and emergent risk analysis is presented in Tables 2 and 3. The main benefit of applying Net-HARMS is that it enables the identification of risks across entire sociotechnical systems. For example, in addition to risks associated with the vehicle operator and vehicle automation, risks associated with the design and communication of standards, regulation of autonomous vehicle testing, design of autonomous vehicles and autonomous vehicle testing programs, and the training of autonomous vehicle safety operators were identified.

AcciNet analysis

The AcciNet analysis is overlaid on the task network in Figure 1. The red shading denotes tasks that were completed sub-optimally or not at all, and the green shading denotes tasks that were successfully completed as required. Nodes that are shaded both red and green represent tasks where selected agents performed the task successfully as required (green) but other agents performed them sub-optimally or not at all (red). Each of the red nodes includes a classification of the contributory factor based on the risk mode taxonomy presented in Table 1, along with a description of the agent associated with the contributory factor. Agents who completed tasks successfully are required are also included within the green shading.

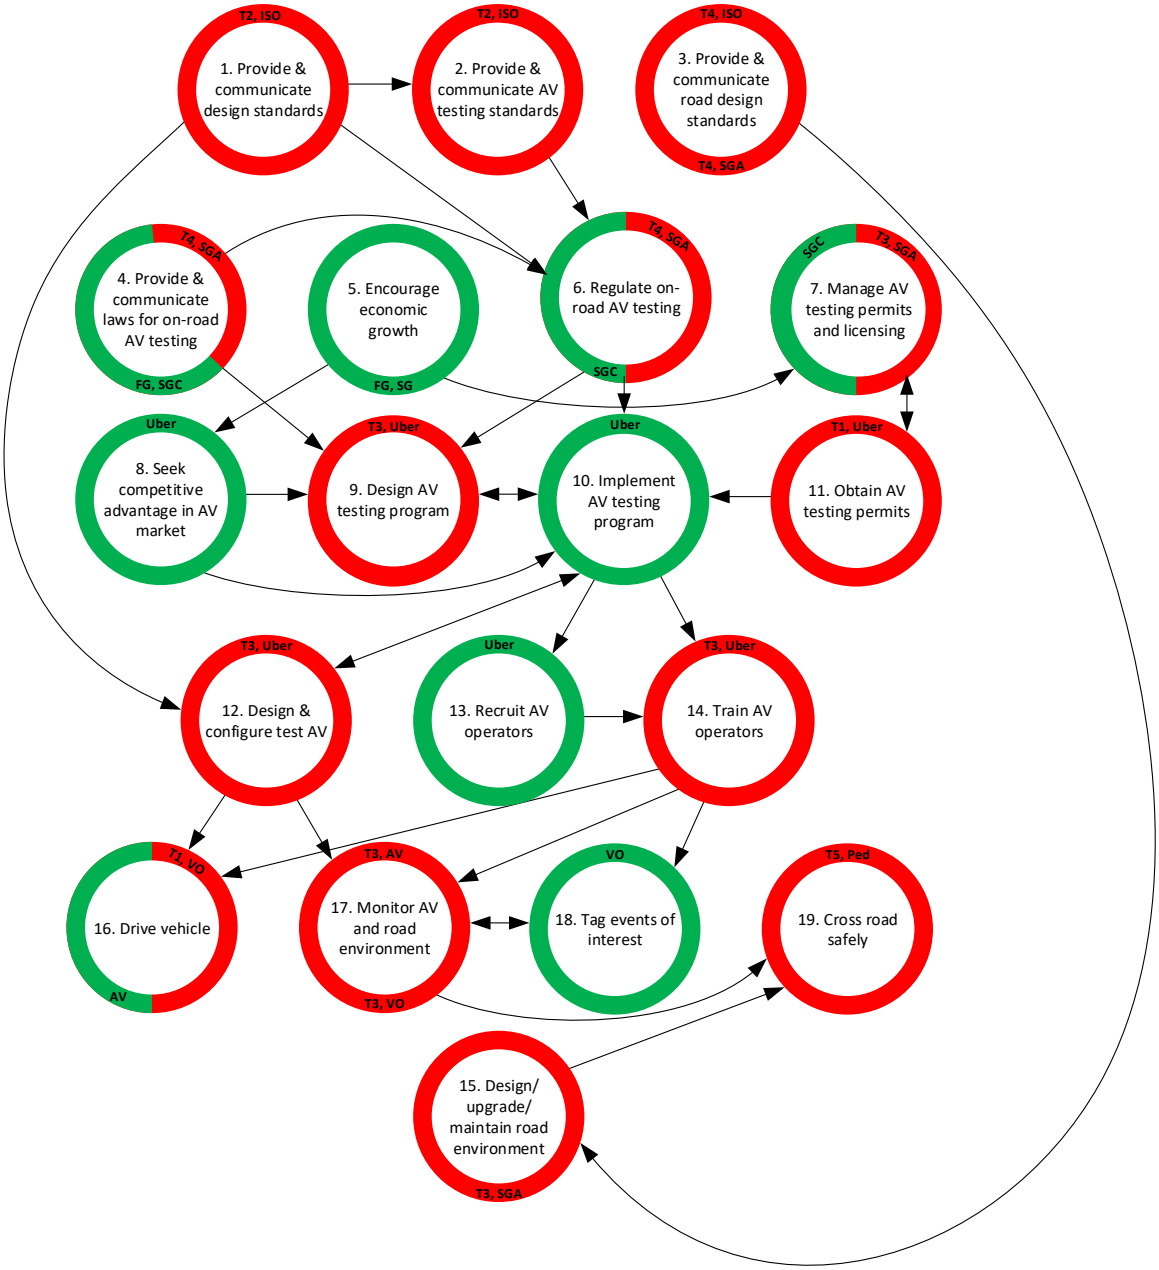

Figure 1. Uber-Volvo task network overlaid with AcciNet analysis. Red shading represents failure; Green shading represents normal performance. ISO = International Standards Organisation; SGA = State Government Road Arizona; SGC = State Government California; FG = Federal Government; AV = Autonomous Vehicle; VSO = Vehicle safety operator; Ped = Pedestrian

Table 2. Net-HARMS task risk analysis extract for 'Provide and communicate design standards' task

Task	**Actor**	**Risk mode**	**Task risk description**	**Task risk consequences**	**P**	**C**	**Risk controls**
1. Provide and communicate design standards	ISO	T1 – Task mistimed	Design standards are developed and/or communicated too late (after on-road testing of autonomous vehicles has begun)	Autonomous vehicle may not be designed to appropriate standard which may introduce risks around features of driving performance e.g. detection and response to other road users	H	M	Introduce regulation restricting testing until appropriate design standards are available
	ISO	T2 – Task omitted	Design standards are not developed at the time of testing	Autonomous vehicle may not be designed to appropriate standard which may introduce risks around features of driving performance e.g. detection and response to other road users	H	M	Introduce regulation restricting testing until appropriate design standards are available
	ISO	T3 – Task completed inadequately	Design standards are inadequate (e.g. do not provide clear standards or guidance on all aspects of autonomous vehicle design)	Autonomous vehicle may not be designed to appropriate standard which may introduce risks around features of driving performance e.g. detection and response to other road users	H	M	Use formal process of consultation and review to ensure design standards are fit-for-purpose

Table 3. Net-HARMS emergent risks extract

Task	Actor	Risk mode	Task risk description	Related task	Emergent risk mode	Emergent risk description	Emergent risk consequences	P	C	Risk controls
1. Provide and communicate design standards	ISO	T1 – Task mistimed	Design standards are developed and/or communicated too late (after on-road testing of autonomous vehicles has begun)	12. Design and configure autonomous vehicle test	T1 – Task mistimed	The design of the autonomous vehicle is delayed as there are insufficient standards available	Testing program is delayed	L	L	- Ensure design and testing program timelines are flexible and have the capacity to accommodate delays
1. Provide and communicate design standards	ISO	T1 – Task mistimed	Design standards are developed and/or communicated too late (after on-road testing of autonomous vehicles has begun)	12. Design and configure autonomous vehicle test	T3 – Task completed inadequately	The autonomous vehicle is not designed in line with appropriate standards	The autonomous vehicle is unsafe and the risk of collisions during testing program is heightened	H	H	- Restrict sign off on autonomous vehicle design until appropriate design standards have been met - Prevent initiation of on-road testing program until appropriate design standards have been met

As shown in Figure 1, the Uber-Volvo incident was created by a network of contributory factors that included both failures, work as imagined, and instances of normal performance. For example, the task of monitoring the autonomous vehicle and road environment was performed sub-optimally by both autonomous vehicle and vehicle safety operator. The design of the testing program was problematic, with vehicle safety operators working eight-hour shifts and being required to monitor the vehicle as it drove around a pre-set route and take over control only when necessary (Stanton et al., 2019). Finally, Stanton et al. (2019) report that there was a lack of international and national standards for automation design and testing, meaning that Uber had little technical guidance for appropriate interfaces, safety standards, or testing regimes. The AcciNet also shows the role of tasks that were completed appropriately. For example, 'Regulate on-road autonomous vehicle testing' and 'Manage autonomous vehicles testing permits and licensing' were undertaken as required by the Californian state government. Initially, Uber was planning to undertake its testing in California but a dispute over the need for permits led to Uber's vehicle registrations being revoked. In response, Uber moved its testing program to the more lenient state of Arizona (encouraged by the Arizona Governor; Stanton et al., 2019a). As such, regulation of testing represents a task that was undertaken as required, but played a key contributory role in the incident, as it ultimately led to Uber's testing program being moved to Arizona.

Discussion

A limitation of most state-of-the-art risk assessment and accident analysis methods is that they cannot be used in an integrated manner. The START toolkit includes systems thinking-based risk assessment (Net-HARMS), accident analysis (AcciNet), and safety intervention evaluation methods (SafetyNet) designed specifically to be used together for organisational safety management. Whilst each method has important strengths when used in isolation (see Dallat et al., 2018; Salmon et al., 2020), it is worth reiterating that the approaches will be most useful when used in an integrated manner. This will enable organizations to use outputs from risk assessments to direct accident analysis activities and then feed accident analysis findings back into their risk assessment processes. This includes using data from accident analyses to validate risk assessment efforts by showing the types of risks that have emerged during adverse events and to inform probability and criticality assessments. Proposed risk controls and safety interventions can also be evaluated and refined via SafetyNet, using previous risk assessment and accident analysis outputs to guide judgements on likely positive and negative emergent properties. As a result, the quality of both risk assessments and accident analyses will be enhanced as both activities become increasingly data driven, and the risk controls and safety interventions proposed will likely be safer and more effective. Further testing and applications of START and its component methods is encouraged. In particular, future research should continue to establish the reliability and validity of the Net-HARMS, AcciNet and SafetyNet approaches (Hulme et al., 2021b, In Press) and also examine the potential use of software support to enable use in practice.

References

Annett, J., Duncan, K. D., Stammers, R. B. and Gray, M. J. 1971, Task Analysis (London: HMSO).

Chemweno, P., Pintelon, L., Muchiri, P. N., Van Horenbeek, A. (2018). Risk assessment methodologies in maintenance decision making: A review of dependability modelling approaches. *Reliability Engineering & System Safety*, 173, 64-77

Dallat, C., Salmon, P. M., & Goode, N. (2018). Identifying risks and emergent risks across sociotechnical systems: The NET-HARMS. *TIES*, 19:4, 456-482.

Dallat, C., Salmon, P. M., Goode, N. (2019). Risky systems versus Risky people: To what extent do risk assessment methods consider the systems approach to accident causation? A review of the literature. *Safety Science*, 119, 266-279

Dekker, S. (2011). Drift into Failure: From Hunting Broken Components to Understanding Complex Systems. Aldershot: Ashgate

Hollnagel, E. (2012). The functional resonance analysis method: modelling complex sociotechnical systems. Ashgate, Aldershot, UK.

Hulme, A., Stanton, N. A., Waterson, P., et al. (2019). What do applications of systems thinking accident analysis methods tell us about accident causation? A systematic review of applications between 1990 and 2018. *Safety Science*, 117, 164-183

Hulme, A., McLean, S., Dallat, C. et al. (2021a). Systems thinking-based risk assessment methods applied to sports performance: A comparison of STPA, EAST-BL, and Net-HARMS in the context of elite women's road cycling. *Applied Ergonomics*, 91, 103297

Hulme A, Stanton NA, Walker GH, et al. (2021b). Are accident analysis methods fit for purpose? Testing the criterion-referenced concurrent validity of AcciMap, STAMP-CAST and AcciNet. *Safety Science*, 144

Hulme A, Stanton N. A, Walker G. H, et al. (In Press). Testing the reliability and validity of risk assessment methods in HFE. *Ergonomics*, DOI: 10.1080/00140139.2021.1962969

Leveson, N. G. (2004). A new accident model for engineering safer systems. *Safety Science*, 42:4, 237—270.

Li, Y., Guldenmund, F. W. (2018). Safety management systems: a broad overview of the literature. *Safety Science*, 103, 94-123.

NTSB. (2018). National Transportation Safety Board Preliminary Report Highway: HWY18MH010. Retrieved from: https://www.ntsb.gov/investigations/AccidentReports/Pages/HWY18MH010-prelim.aspx (12 December 2018).

Rasmussen, J. (1997). Risk management in a dynamic society: A modelling problem. *Safety Science*, 27:2/3, 183-213.

Salmon, P. M., Hulme, A., Walker, G.H., et al. (2020). The Accident Network (AcciNet): A new accident analysis method for describing the interaction between normal performance and failure. *HFES Society Annual Meeting*, 64:1, 1676-1680

Shorrock, S.T. and Kirwan, B. (2002). Development and application of a human error identification tool for air traffic control. Applied Ergonomics, 33(4), pp.319-336.

Stanton, N. A., Salmon, P. M., Walker, G., Stanton, M. (2019). Models and Methods for Collision Analysis: A Comparison Study based on the Uber collision with a pedestrian. *Safety Science*. 120, 117- 128.

Waterson. P.E., Robertson, M.M., Cooke, N.J., Militello, L., Roth, E. and Stanton, N.A. (2015). Defining the methodological challenges and opportunities for an effective science of sociotechnical systems and safety. *Ergonomics*, 58, 650-8.