The Safety-Critical Systems Club Newsletter

Safety Systems

Vol 34 No. 1 – Feb 2026



SEA CHANGE
A way ahead for maritime autonomy?

THE TWISTS OF FATE
The Lisbon Funicular Disaster

JOINING THE DOTS
Our ontological odyssey continues!

WORK AS IMAGINED?
Reflections from an early careers engineer

SCSC

thescsc.org



SCSC Publication Number: SCSC-209

While the authors and the publishers have used reasonable endeavours to ensure that the information and guidance given in this work is correct, all parties must rely on their own skill and judgement when making use of this work and obtain professional or specialist advice before taking, or refraining from, any action on the basis of the content of this work. Neither the authors nor the publishers make any representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability or availability with respect to such information and guidance for any purpose, and they will not be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever (including as a result of negligence) arising out of, or in connection with, the use of this work. The views and opinions expressed in this publication are those of the authors and do not necessarily reflect those of their employers, the SCSC or other organisations.

© SCSC CIC 2026 except where explicitly stated for specific articles, reports or images.

Front cover: AI generated image
Editorial: Costa Concordia ID 44074775 © Alessandro Canova | Dreamstime.com
Editorial: Lloyd's Register, © Paul Hampton, licensed as Public Domain
In Brief images may be subject to copyright – contact the SCSC Newsletter Editor for details.

Contents

WELCOME

Editorial 3

Opening words from the SCSC Newsletter Editor.

In Brief 4

Recent system safety news items from around the world.

FEATURES

Maritime Autonomy – A Way Ahead? 5

Nick Tudor discusses a way forward to enable Maritime Autonomy.

Cracking Tales of Safety-Critical Cables 11

Wendy Owen and Jon Wiggins discuss the Lisbon Funicular Disaster

Joining The Dots 21

The SCSC Ontology Working Group provides the next episode in a series explaining ontological concepts.

Reflections from an Early-Career Engineer 31

Beatriz Coutinho provides insights as an early career professional.

REPORTS

Seminar: Complying with the EU AI Act 35

Highlights from this SCSC seminar in Brussels discussing the EU AI Act.

Seminar: New Developments in Rail Safety Systems 43

Highlights from this SCSC seminar discussing developments in rail safety systems.

SCSC Timeline 51

Paul Hampton introduces the new SCSC Timeline.

60 Secs with ... Chris Hobbs 55

Chris answers some quick-fire questions on system safety and the future.

GROUPS

Working Groups 61

Details of the SCSC Working Groups.

SCSC Steering Group 69

Who's who in the Steering Group.

EVENTS

Calendar 71

Events Diary 72

Upcoming Seminar: Soft Skills for Safety Practitioners



scsc.uk/e5055

Editorial

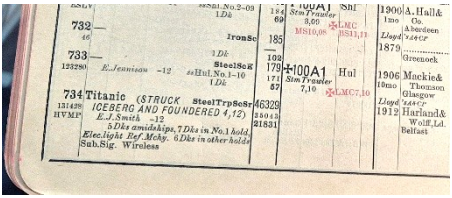
Welcome to the February 2026 edition of the SCSC's *Safety System* Newsletter, the first of the New Year. I hope you all had a relaxing seasonal break and are ready for the challenges that this year will bring!

In my last editorial I compared two significant 'accidents' that spanned almost a century and noted some commonalities that led me to ask how much we were actually learning from accidents. Through this line of enquiry, I happened upon the work of Nippin Anand, who had asked the very same question in his 2024 book *"Are We Learning from Accidents? A Quandary, A question and a way forward"*.



The book focuses on the events of the *Costa Concordia* tragedy in 2012. I provide a review of Nippin's book in this edition and I'd encourage you to read the book as it may well challenge your thought processes when thinking about the Concordia event and in investigating accidents in general.

The Concordia incident marked almost the centennial anniversary of the Titanic disaster in 1912. Coincidently, at the end of October I attended an exhibition on the Titanic that was being held at Westpoint in Exeter. The events of the Titanic sinking are of course well-known but I was struck by a quote from J. Bruce Ismay, Director of the ship owners, White Star Line, when intervening in the design to reduce the number of lifeboats from the proposed 64 to 32. *"I'll not have so many little boats, as you call them, cluttering up my decks and putting fear into my passengers".*



There is strong resonance with *Costa Concordia* in this regard, with the pursuit of improved customer experience being a disproportionately powerful cultural factor. I am reminded in my own career as a safety engineer that commercial pressures often come into conflict with safety activities. Standing up for our principles can be challenging as it may lead to uncomfortable situations, especially for those carving early careers, where there is often a steep seniority gradient. The SCSC is therefore delighted to be able to run a workshop in this area (scsc.uk/e5055) aimed at developing the softer skills required by Safety Practitioners. We are very excited about this event as we plan to make it interactive with some roleplay activities!

We lead with an article from Nick Tudor covering ideas around accelerating the adoption of autonomous vessels in the maritime domain. We then have a fascinating insight from Wendy Owen and John Wiggins into the Lisbon Funicular tragedy as well as other issues with the reliance on cabling. The Ontology Working Group continues their serialisation of their introduction to ontology based on familiar aviation-related concepts. I am also delighted to introduce Beatriz Coutinho, who provides fascinating reflections on her experiences as an early career's professional. In contrast, we have a 60-seconds interview with Chris Hobbs, who has many decades experience in the industry. Finally, we have two event reports from our two trips to mainland Europe in Brussels and Munich.



In Brief








Nvidia unveils 'reasoning' AI technology for self-driving cars



In January 2026, Nvidia boss Jensen Huang announced Alpamayo, a tech platform the company says will help self-driving cars think like humans.

"Alpamayo brings reasoning to autonomous vehicles, allowing them to think through rare scenarios, drive safely in complex environments, and explain their driving decisions," Huang said on stage at the annual CES technology conference in Las Vegas. The vehicle will be released in the US in the coming months before being rolled out in Europe and Asia. *bbc.co.uk*



Space Debris Could Soon Disrupt Flights Like Bad Weather

A 2025 study conducted by the University of British Columbia estimates a 26% probability that, within a year, uncontrolled space debris will re-enter Earth's atmosphere over heavily trafficked airspace. And by 2030, the probability of an individual commercial aircraft encountering a piece of debris could reach 1 in 1,000, based on earlier modelling.. *dailygalaxy.com*

Drones used to carry blood in trial aimed at saving lives



Specially commissioned drones will be used to fly blood donations as part of a new trial. Currently, blood donations are processed in south Wales then transported by road, a journey that can take hours.

The ultimate ambition of the Dragon's Heart project is to fly life-saving blood samples to the scenes of accidents using drones." *bbc.co.uk*



Is texting behind the wheel of a self-driving Tesla crazy?

As self-driving cars get closer to reality, Tesla is striving to remain a big player. Tesla boss Elon Musk told investors: "We're going to look closely at the safety statistics, but we will allow you to text and drive essentially". However, some consumers conducting a free trial of Tesla's latest self-driving software are still reluctant to fully depend on the system. *bbc.co.uk*

From Iran to China to Venezuela – how tanker seized by US hid true location



The oil tanker seized by US forces on 10th December 2025 had a track record of faking or concealing its location information, apparently to hide its activities, ship tracking data shows.

Experts told BBC Verify that the *Skipper* was likely a part of the so-called "dark fleet" - a global network of oil tankers that seek to evade oil sanctions by obscuring their ownership, identities and travel histories. *bbc.co.uk*

Maritime Autonomy – A Way Ahead?



Maritime autonomy allegedly holds out the potential for significant economic gain, should anyone be able to crack all of the problems that would make a scalable business. Nick Tudor explains why it is not currently possible to make such a business and proposes a way ahead that could be seen as radical by some, necessary by others.

Full Autonomy

Currently, we treat autonomous vessels as if they were 2-year old children; we don't trust them to do things on their own. We give the child a task and we watch them like a hawk and we don't let them anywhere near sharps knives, fires, etc. We need to change this so that we treat them, more like, and I accept that the analogy might break down here, teenagers. In this scenario, the chances are that they are capable of doing the job asked of them unsupervised, to a reasonable/acceptable standard and when they have finished or if they are unable to do so, then they will report back. The focus here will be on what I am going to refer to as 'full autonomy' which I am going to define as follows: The vessel is given a task, it goes and does it unsupervised and it only reports to a human if something has stopped it from doing so, or it has completed the task; a teenager at least!

> **"We treat autonomous vessels as if they were 2-year old children; we don't trust them to do things on their own"**

The Law

Without going into specifics, I am not a lawyer, let alone a maritime lawyer, international law requires that every vessel shall have a Master. In conversations with maritime lawyers, Class Societies, insurers and others, this is so that there is someone to blame when it has all gone wrong. This might be a bit cynical, but it appears to be not far from the truth. The role of a Master is to be the final arbiter, the ultimate decision maker on all matters at sea. For autonomous vessels of any size, this role is still required by law and therefore the regulation that enables compliance to that law means that some aspects of autonomous vessel design operation and deployments have to adopt some human based practices. However, there is no-one on board the vessel.

Master of all they Survey?



In a number of circumstances, what exactly can a Master do? For example, the 'Officer-of-the-Watch' also known as the 'Operator' in a facility called a Remote Operation Centre (ROC), while supervising the autonomous vessel (they're only allowed to supervise one at a time), notices that the communications have failed.

Such problems have to be raised to the Master who maybe anywhere else in the world. They are woken up, arrive on the bridge/arrive in the ROC/phone in to be briefed that the Operator cannot communicate with the vessel.

Questions is: what exactly are they supposed to do about it; there are no communications? (There are similar other such issues for which it is probably not right to assume functions will be available for an autonomous vessel that are analogous to a crewed vessel; a list can be provided!) To repeat, there is no-one on board.

> ***"it does not make sense for there to be a role of Master for an autonomous vessel in the same way that one is required for a crewed vessel"***

This means that the vessel has to behave as though there were humans on board even though there aren't and this has to be apparent, as far as is reasonably practicable, to other maritime traffic. I argue that it does not make sense for there to be a role of Master for an autonomous vessel in the same way that one is required for a crewed vessel.

Removing the requirement for a Master would require international agreement to change the law, but we need changes to international and national laws to enable scalable use of autonomous vessels anyway, so this is a difference that makes no difference. So, who is going to be the one in the dock when it all goes wrong? I'll come back to that, shortly.

Class

Another big issue is Class. The problem as I see it is that the mantra currently adopted is that autonomous vessels have to be 'as safe as crewed vessels'. The Class determines what functions a vessel has from many perspectives, especially safety. However, it appears that there is no objective definition of what constitutes 'safe'. Arthur Dent in *The Hitchhikers Guide to the Galaxy* [1] had the same problem when Ford Prefect enlightened him to his predicament and told him he was 'safe' on the Vogon constructor ship[1].

Currently, it appears that good practice has evolved since at least Nelson was a boy and an assessment is undertaken by Class Societies. Opinion varies as to which of these Societies is best; another level of subjectivity, and one needs to define 'best'! These vessels have no-one on board, I may have mentioned this before, so why do we need to classify them in the same way as crewed vessels?

There will be many arguments why this should be so, but historical precedent is not necessarily relevant in this case. Let's turn the question round and rephrase to 'why not have a different class'? We can still adopt a swathe of good practice but stop where it does not make sense. This would allow a number which defines what we mean by 'safe' (enough) as is the case in most other sectors ranging from aerospace, automotive, rail and industrial plant. As a result, lots of problems can be solved.

Safety Argumentation & Evidence

If we have an objective safety target to meet, then argumentation can be put in place to show that any aspect of the vessel has been assessed from a design perspective. We therefore show that we understand the contribution to the safety of the vessel through the mitigation of hazards. By doing so, evidence can be produced to show that the claims have been met, or otherwise. Without an objective safety target, subjectivity leads to non-standardised vessels and consequent questionable safety; depending on your criteria for 'best' with regard to the Class Society used.

Insurance

Insurers are all about the numbers; they love them! Give them such an understanding of the risk and they can then properly cost that risk; what's not to like in this regard?

[1] Arthur: If I asked you where the hell we were, would I regret it?
Ford: We're safe.
Arthur: Oh good.
Ford: We're in a small galley cabin in one of the spaceships of the Vogon Constructor Fleet.
Arthur: Ah, this is obviously some strange use of the word *safe* that I wasn't previously aware of.
— **Douglas Adams, The Hitchhiker's Guide to the Galaxy**

Software

The maritime world has no sector specific software standard and that's great; the world doesn't need any more of them. Software systems are not necessarily a one-for-one replacement of electrical or mechanical systems. Development of software in aircraft has been carefully developed to meet an internationally agreed standard which is 'goal' based and aligns well with safety argumentation.

Other sectors have developed software standards but these tend to be technique oriented which can lead to a tick-box approach to assurance of the software system. Given that sensor systems can be built to detect on-board and off board issues, that there is a sufficiency of them (spectrum coverage and redundancy) set in a systems architecture, then there merely needs to be some software that makes decisions, high integrity decisions. This software cannot be artificial intelligence (AI), the safety of the vessel depends on it and AI is a bet.

The software has to be shown to play its part in the mitigation of hazards, so deploying robust techniques to build this are required. Sounds expensive? Well, so is losing the vessel! Nevertheless, while other approaches are available, check out the techniques deployed by D-RisQ which can save a considerable amount of time/cost while providing evidence to support the safety case.

Other Sectors

Interestingly, a few years ago, Volvo declared that if (when?) they produced an autonomous car, they would accept liability for any accidents; it's apparently still 'when'. In the air industry, the aircraft manufacturer is liable, same for most other products. However, like all such products, they may decide that the issue was not caused by them but their supplier and hence the next lawsuit is raised and lawyers make their money.

Maritime Autonomy & Liability; A Suggestion

By removing the role of Master, liability can shift to the vessel producer. So long as the vessel is operated within the defined context of use (defined in the safety case), then any issues can be pinned on the vessel designer; if it's out of the context of use, the operator /service provider becomes liable. Proving this should not be difficult. Someone is still likely to be liable, which should satisfy the lawyers!

The Market



Should there be a market for autonomous surface vessels, it is not going to be sufficiently addressed by having one person supervising one vessel; it's not a scalable business model. Given that the high integrity software can ensure that the vessel behaves at least as well as a 'teenager' then there is no real reason for a one-to-one relationship. The service has to be reliable and should be able to continue in the vast majority of cases without supervision. One might decide that higher risk areas (eg around harbours) there has to be closer supervision. This would make the market viable through being scalable. Note that the operator still has to operate the vessel in accordance with the design specification and the design specification has to be maintained in accordance with the manufacturer's requirements, in the same way that aircraft and cars are required to do every year or so.

Conclusion

"The significant problems we face cannot be solved by the same level of thinking that created them": Albert Einstein. If we remove the role of Master, have a completely different class of vessels, build them such that it can be argued to be adequately safe against a design safety target, then there is a chance that a market could be made and addressed. Demonstrating this approach might cost a few quid, and cause some changes to international agreements and the odd law here and there, but then away anchors!

"The significant problems we face cannot be solved by the same level of thinking that created them": Albert Einstein

References

[1] The_Hitchhiker's_Guide_to_the_Galaxy, https://en.wikipedia.org/wiki/The_Hitchhiker's_Guide_to_the_Galaxy_(novel)

Image attribution:
autonomous: ID 346620250 | Autonomous © Galih Wisnu | Dreamstime.com
ROC: AI generated: Midjourney
oceanus12: ID 382361426 | Autonomous Marine © IanBushwacker | Dreamstime.com
class silhouette: Gdr at English Wikipedia., CC BY-SA 3.0 http://creativecommons.org/licenses/by-sa/3.0 , via Wikimedia Commons

Nick Tudor, D-RisQ



Following a full career as an RAF Officer Engineer, Nick has been working in software and high integrity systems for the past two decades. As co-Founder of D-RisQ, he has worked in multiple sectors including aerospace, defence, automotive, rail, autonomous systems in air, land, sea, nuclear decommissioning and cyber-security. Perhaps the only claim to fame he might have is as one of the key authors of DO-333, the Formal Methods Supplement to DO-178C (the aerospace software standard).


ISSN 2754-1118
Vol 4, Issue 2
The Safety-Critical Systems Club
SAFETY-CRITICAL
SYSTEMS eJOURNAL
SCSC
For Everyone Working in System Safety

scsc.uk/scsj

Cracking Tales of Safety-Critical Cables



Wendy Owen provides a fascinating insight into the Lisbon Tram/Funicular accident. She describes the circumstances surrounding the accident, the reported causes and the impact and aftermath of the incident. Wendy and Jon Wiggins then provide a few similar case studies to show this is not an isolated case.

The first time I heard about the Lisbon Tram/Funicular accident (3rd Sept 2025) was on the BBC news. Witnesses described it as "out of control" and the whole incident happened in less than a minute.

A Review of the Lisbon Accident

The Glória funicular derailed and crashed into a building in Portugal's capital. 62 rescue personnel and 22 emergency vehicles attended the scene.

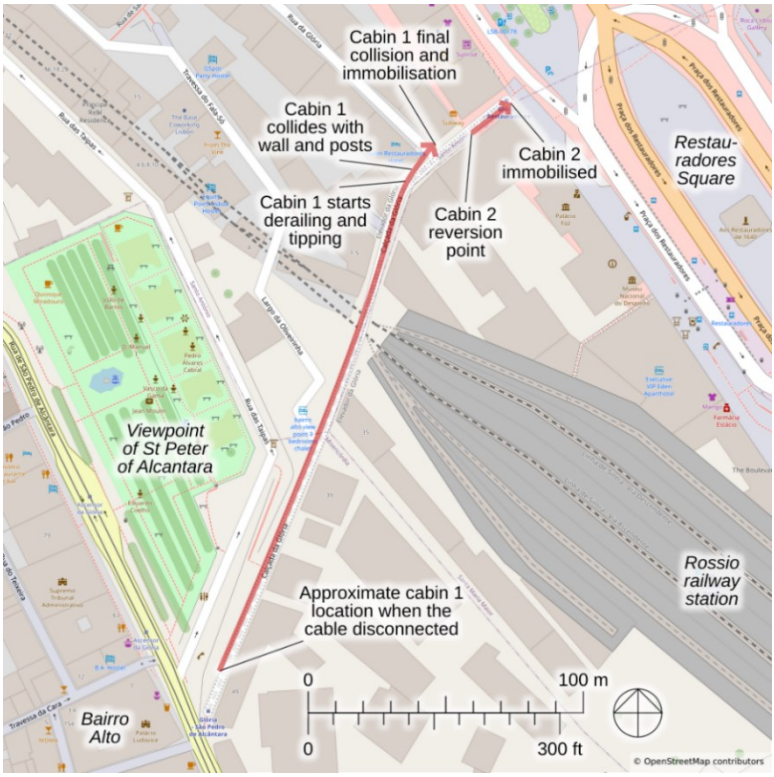
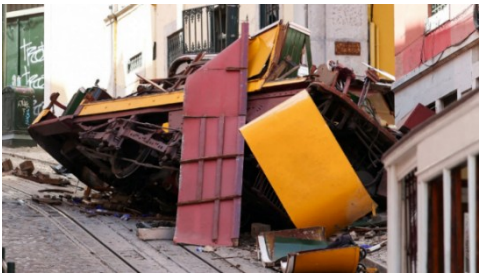


Figure: Map of the incident

Here's a map of the incident. The car at the top of the hill (lower left of the map on the figure) sped uncontrollably downhill before derailing and hitting a building near the bottom.

Not long after, news came in that three Brits were involved, including news in my local papers – David Young, 82, was a former customs officer and lifelong transport enthusiast [1], a Scot living on Anglesey. He was among 16 people who died (of twelve nationalities involved, all adults, including one hit by the railcar). Additionally, 23 people were injured. David's family said [2] it was a comfort that *"his final moments were in pursuit of the hobby which gave him so much happiness… In retirement he enjoyed visiting heritage railways and tramways around the world*". Cheshire Police confirmed the other two were Kayleigh Smith, 36, and her partner Will Nelson, 44. Llinos Medi, MP for Ynys Môn, sent her condolences to Mr Young's family and "*to those who have been affected*". Prime Minister Sir Keir Starmer said he was "*deeply saddened*" by the incident, "*His thoughts are with their families and those affected by this terrible incident*," said a spokesperson. There are many people like David who very much enjoy travelling on heritage transport.

About twenty years ago I was an independent safety reviewer for a new tram project. I am also more generally familiar with a few heritage funiculars in the UK. I'm particularly fond of old trams and funiculars, and I'd travelled on that now 140-year-old iconic funicular in Lisbon some years ago. I had travelled on a fantastic funicular in Bridgnorth not long before the Lisbon accident.



It was notable that those affected were often referred to as victims rather than fatalities or deaths. A crime had taken place, apparently. I really, really, wanted to know more about what had happened – was this a crime, an unpredictable mishap, or something else?

We didn't have to wait long to get some initial answers, only a few days [3].

Figure: Immediate aftermath of the derailment

Investigators determined the carriage was travelling at ~60km/h (37mph) when it hit the building and that a cable connecting two carriages had snapped, although the rest of the cable mechanism was functioning properly. They added that the driver (AKA brake guard) tried to apply the emergency brakes but failed to prevent the derailment. Although he activated both the pneumatic brakes and a manual brake when the cable came loose, at this early stage of the investigation it was not clear whether another – automatic - brake initiated as it was supposed to. Sadly, the driver was one of the deceased.

The cable was reportedly only 337 days into its expected 600-day operational life. Let's pause at this point – this is saying that the cable snapped at just over half-way through its expected lifetime, that its expected lifetime is not even two years, and this is also given that the trams perform a very frequent daily service.

A fuller report was due within 45 days, but again we didn't have to wait long, the larger report (by Portugal's Air and Rail Accident Investigations Bureau) was released on 10th October 2025 [4].



Figure: Immediate aftermath of the derailment, as seen from Restauradores Square, with dark yellow derailed car toppled into building on right.

Litany of Failings

This reported a "*litany of failings*"; for readability here, I have grouped these failings into five topics. A note that the words carriages, cabins and cars are loosely interchangeable.

Cable Failure Mode

- An underground cable – which acted as a counterweight between two carriages – was defective and broke, causing the crash
- After travelling only about 6 metres, the connecting cable between the two cabins failed at its attachment point to the upper car
- While a visual inspection of the scheduled cable was allegedly performed on the morning of the accident and did not show any issues, the section of the cable where it eventually failed was not accessible for observation.



Figure: Cable broken during the derailment. Bottom left: cross section. Right: unbroken cable inserted into the *trambolho*.

Cable Specification Issues

- The cable had never been certified for passenger transportation. The cable was "*not technically suitable*". It did not comply with the specifications in force at the CCFL [Lisbon's transport operator] to be used for the Gloria tram.
- Yet, the specific type of cable had been used on this funicular for about six years.

Braking System Issues

- The vehicle's speed could not be reduced, despite the pneumatic and hand brakes being applied immediately, once the driver lost the balancing force of the cable
- Activation of the emergency braking system failed to stop the derailment. This system, which the driver correctly tried to apply when the cable snapped, either did not function at all or did not function properly
- The emergency braking system was never tested in advance

Procurement & Installation Issues

- The cable was acquired in 2022 by Carris (who run Lisbon's transport system). There had been no oversight by Carris engineers
- The cable was not tested before it was installed

Maintenance Background & Inspection Issues

- Maintenance regimes of the Ascensor da Glória
  - It was last placed under general maintenance (held every four years) in 2022
  - Its most recent immediate maintenance (held every two years) was conducted Aug-Sept 2024
  - A routine daily visual inspection was allegedly carried out nine hours before the derailment

- There were several instances when the maintenance programme referred to "*non-existent, inapplicable or outdated standards*"
- The supervision and maintenance of the funicular (by a company outsourced by Carris) did not work properly. Apparently, the Glória funicular was given the all-clear on the morning of the disaster.

The preliminary report emphasised that the information gathered to date was "*incomplete*", with further tests and analysis needed. "*The guilt or responsibility of any organisation or person involved in the incident should not be presumed*".

Meantime, Lisbon's mayor reaffirmed that "*the unfortunate tragedy... was due to technical and not political causes*".

Orders to Operators

The investigation report concluded: "*At this point in time, it cannot be said whether the use of this type of non-compliant cable intervened, or what intervention it had, in the rupture... and it is certain for the investigation that there were other factors that had to intervene.*"



Figure: An operator's view from the cabin of car No. 1 in July 2023, showing the location of its fatal derailment.

It ordered/advised:

a) All Lisbon cable cars to be taken out of service until the necessary safety checks were in place
b) All braking systems to be reviewed. These need to be capable of immobilising the cabins in the event of a cable break
c) Carris to implement a new safety management system in line with European best practice
d) Closing a loophole whereby Lisbon's historic funiculars are exempt from legal and supervisory oversight that applies to other cable forms of transport

Carris issued a few statements in response:

- They had dismissed their head of maintenance of funiculars and trams
- Said it was *"not possible at this time to state whether or not the non-conformities in the use of the cable are relevant to the accident*"
- Noted that other cables had been in use on the Glória funicular for 601 days without incident [cue tut-tuts from the RAM engineers]
- The cables had been brought into use under the current board of directors, who took office in May 2022, but the acquisition process took place under a previous board.

What Else?

Other pertinent information that arose during the writing of this article not already mentioned above:

- A few hours after the disaster, Carris' CEO stated to the press that "*the maintenance protocol was scrupulously followed*" by MAIN and that it included monthly, weekly and daily inspections
- An incident had occurred in May 2018, when one of the cars derailed due to a lack of maintenance of its wheels. No one was injured. The outsourced contractor supplying maintenance at that time is not named.

The Aftermath

A full investigation report will take about 11 months to complete. A National Day of Mourning was declared following the accident. The historic tramway is out of action.

Lisbon municipal council formed a technical team of Carris engineers, university experts, Ordem dos Engenheiros and the National Civil Engineering Laboratory (LNEC) to develop new safety mechanisms and determine if, and when, the funicular is safe to reopen.

On the face of it, it does appear as a *"litany of failures*":

- Ultimately, no-one spotted that it was the wrong type of cable being installed and put into service. There are some big questions here about how and why it was not spotted.
  - It definitely was not the specified cable (assuming that the spec itself was correct), reflecting a lack of due diligence by both procurement and installation teams.
  - Maintenance also missed it, on many occasions, although there is some evidence that it wouldn't have been possible to spot this type of fault during a visual inspection.
  - It also wasn't tested post-manufacture or upon installation (as far as I can tell). Nor does it have a long lifespan.
- There's something seriously amiss with the emergency braking.

We shall wait and see what the technical team come up with.

Note: In the UK, funiculars require formal authorisation by an Approval Body and HSE/ORR (depending). In general, they are not classed as "light rail" but fall within the remit of the Cableway Installations 2018 (a post-Brexit version of an EU regulation).

To conclude this already fairly long article, here is a small collection of other cracking cable tales.

Case Study 1 (Jon): Post-Tensioned Concrete Tanks

The storage of large volumes of water or media is a common occurrence in the water industry. A cost-effective method of constructing and installing large volume tanks is therefore a requirement.

A common solution is to utilise post-tensioned Concrete tanks. In essence, these are pre-cast concrete sections arranged on a pre-cast base in a circle with a series of steel cables (or tendons) threaded through them and tensioned to allow the seals between the section to become watertight and remain so under the highest load conditions as illustrated below.



A common use is in storm diversion tanks. These tanks take the excess water which cannot be treated by a sewerage works during a storm and then later discharge through the works once the high flow has abated. As such, these tanks go through regular stress cycles with water loading and temperature variations causing expansion and contraction of the tendons.

There have been some notable cases of the tendons failing catastrophically causing localised damage and weakening to the tank and in extreme cases collapses, opening out like petals of a flower, a failure mode nicknamed "petalling" as illustrated below.

The implications of these "petal" failures are clear

- Potential for death or serious injury due to tendon ejection
- Potential for death due to tank collapse
- Potential for environmental and property damage from sewage leaking



The case-by-case failure mechanism can be difficult to pinpoint and studies into post tensioned structure failures have drawn several hypotheses.

Two common factors arise:

- Corrosion. Steel and water, especially sewage, is a mixture made for corrosion [9]. In addition, effects such as embrittlement and chloride ion corrosion have been postulated [10]
- Creep and stress cycling. Where structures are subjected to repeated loading/unloading cycles, the movements and loading patterns have been noted to cause cracking in the concrete structure, and creep and damage to the tendons.

The regimes to maintain and inspect these structures have varied, although a lack of maintenance and inspection may be seen as a contributory factor in failures. These structures require that tendons are maintained and inspected at regular intervals. Inspections are predominantly visual checks. Tendons may be removed for inspection and re-greasing alongside a tension check.

Other measures, such as the addition of secondary supporting cables (to prevent petal failures) and physical barriers to mitigate the risk of tendon ejection have additionally been installed. Awareness of the hazard and the need to carry out regular inspections is key to risk reduction and improvements.

Case Study 2 (Wendy): The Maverick Procurement Team

I have a tale about another cable where no-one spotted it was the wrong type of cable before it went into service. It was a very different type of cable to the above, no-one was hurt, it was during a trial, but there are some similarities. I have previously used this as a case study for training new RAMS engineers.

It happened during a step-change technology upgrade at a small regional control centre. I was called in on a fact-and-fault-finding mission on behalf of the manufacturer of the new technology. Something was awry with some readings, and after a few weeks no-one had been able track down the problem.

I sat there (remotely) for a day observing the operational data and talking to a few of the operators. The first sign of the problem was visual, on monitors and printouts. Sometimes it was fine, occasionally it was "wobbly", and even less occasionally it was blank. Whatever was happening was certainly erratic. They had checked and re-checked the monitors, the printer, the systems under control. Nothing, apparently. Eventually, after noticing something about the patterns on the printouts and the odd way the monitors looked at times, I asked them to double-check a few of the cables. Then I asked them to replace two of the cables, one at a time, then both together. The replacement cables looked slightly different if you looked very closely. I asked them to compare, contrast and triple-check the cables.

To cut a long story short, there was one "faulty" cable – it was supposed to be a high quality, high integrity EMC/EMI protectively sheathed cable, but it was not. This affected data transfer to/from the monitors and printers. The interference was being "set off" by various electrical equipment in the room, both permanent and transitionary/temporary equipment, and possibly mobile phones too. Someone mentioned the cables were procured in a rush prior to the trial and someone else mentioned getting them for a very good price. The receipt for the dodgy cable was eventually tracked down, revealing purchase at a local hardware shop.

A few months later, there had been no erratic behaviour, and the trial concluded. The system as installed is still in use today and there has, to my knowledge, been no more wobbles.

End Note

The day before this article was completed, there was widespread disruption on the UK rail network. It centred around a power outage at Leeds Station, reportedly from a *"defective cable*" [11].

References

[1] https://www.bbc.co.uk/news/articles/cq8ev844eqpo Lisbon funicular crash victim was 'transport enthusiast' 8/9/2025
[2] Families pay tribute to British couple killed in Lisbon funicular crash 5/9/2025
[3] https://www.bbc.co.uk/news/articles/c2dnp9pre2lo Cable snapped before Lisbon funicular crash, investigators say 6/9/2025
[4] https://www.bbc.co.uk/news/articles/c20pg8pzp2no Lisbon funicular crash initial report reveals litany of failings 20/10/2025
[5] https://www.bbc.co.uk/news/articles/cnvryg271ymo Lisbon funicular worker among those killed in crash 4/9/2025
[6] https://www.reuters.com/world/europe/lisbon-design-new-mechanism-ensuring-funicular-safety-after-crash-2025-09-08/ Lisbon to design new mechanism ensuring funicular safety after crash". Reuters. 8/9/2025.
[7] https://en.wikipedia.org/wiki/2025_Ascensor_da_Gl%C3%B3ria_derailment 2025 Ascensor da Glória derailment [checked 30/11/2025]
[8] https://assets.publishing.service.gov.uk/media/5ef0c82ee90e0741fa96c841/cableway-installations-2018.pdf The Cableway Installations Regulations 2018 Guidance July 2018
[9] https://www.cross-safety.org/uk/safety-information/cross-safety-report/post-tensioned-precast-concrete-tank-failure-147
[10] https://www.sciencedirect.com/science/article/abs/pii/S0141029616313633 Damage and collapse mode of existing post tensioned precast concrete bridge: The case of Petrulla viaduct. Feb 2018.
[11] https://www.bbc.co.uk/news/articles/c1m8nx0dv94o Power outage at Leeds Station leads to widespread rail disruption. Nov 2025.

Wendy Owen (Independent Consultant)



Wendy has over 35 years' experience in the field of safety and systems assurance. Approximately a third of her career to date has been in the transport sector, predominantly in rail. In the mid-00s, she was a co-author/editor of the UK rail sector's Engineering Safety Management guidance (AKA Yellow Book https://www.rssb.co.uk/en/standards-catalogue/CatalogueItem/YB-Iss-4), particularly the ISA sections. Yellow Book subsequently formed a basis for CSM-RA and the book itself transformed into an international version, iESM (https://www.intesm.org/), since used in many other countries. She spent a portion of her childhood in view of the funicular at Aberystwyth and is always delighted to see a tram.

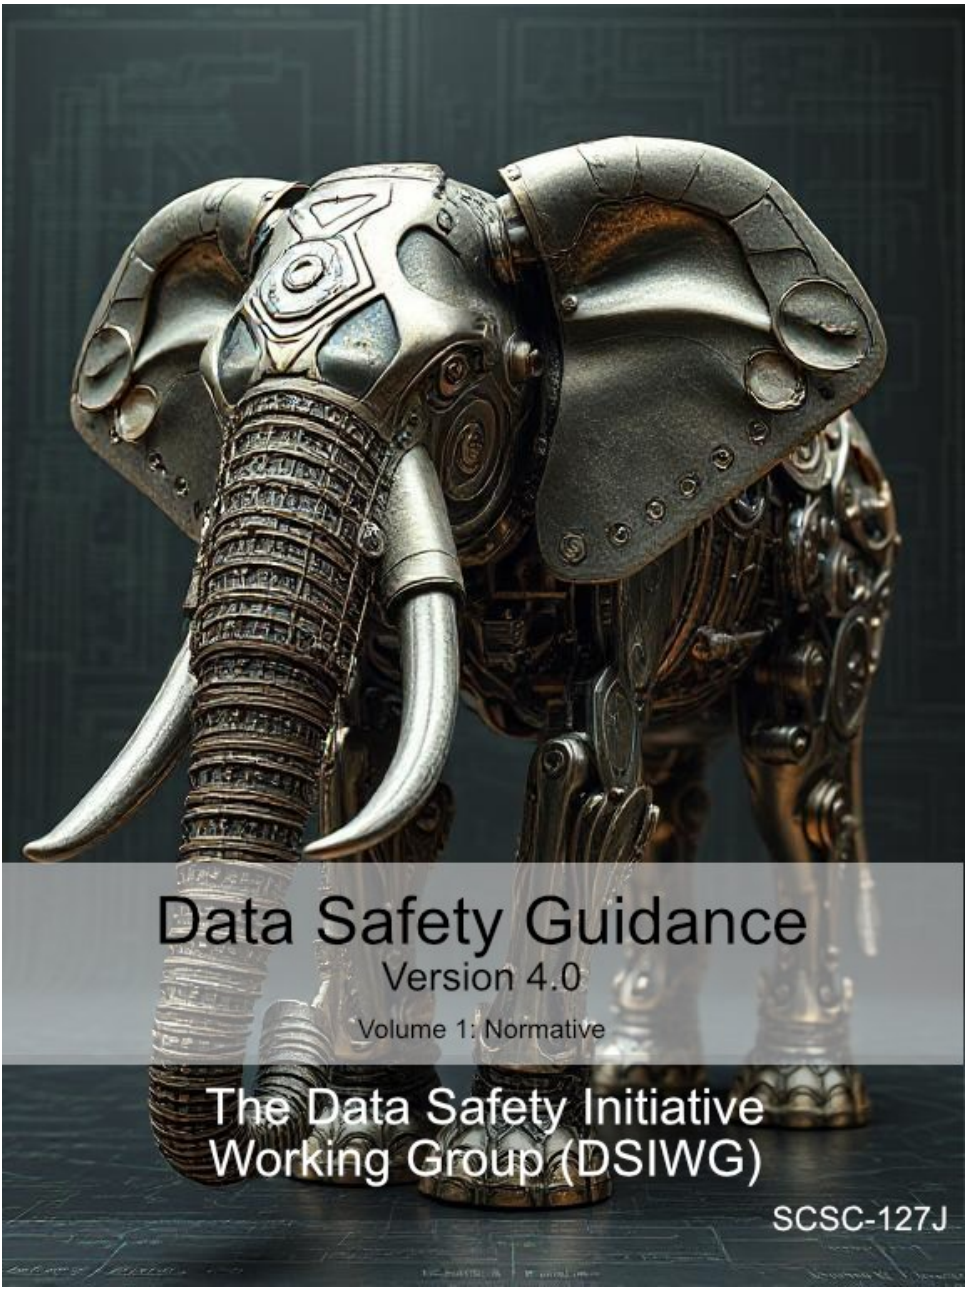
Jon Wiggins (Functional Safety & Cybersecurity Technical Managed, DEKRA)



Working in a wide variety of sectors from Utilities to OEM automation system design, Jon has built a wide range of experience in the construction, design, maintenance and assurance of systems. As a chartered engineer and UK expert to the IEC and ISO Jon advises on the safety of systems and machinery as well as matters of cybersecurity. Jon is a serial author and advocate for improved safety and security throughout UK industry.

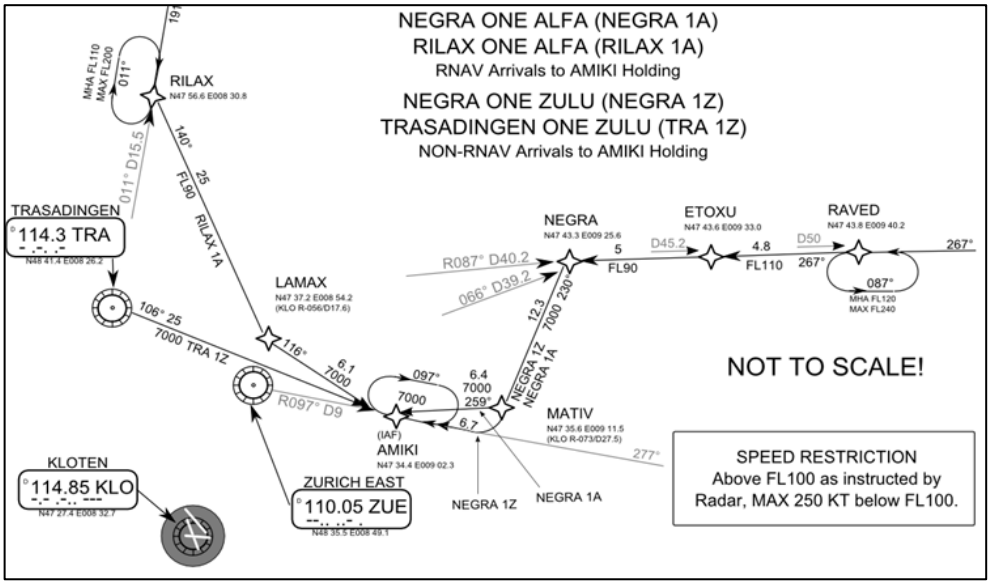
The authors retain copyright of this article.

Image Attribution:
cover image: ID 38285338 © Dennis Van De Water | Dreamstime.com
map: OpenStreetMap contributors Labels: Gabinete de Prevenção e Investigação de Acidentes com Aeronaves e de Acidentes Ferroviários (GFIAAF) Edit: cmglee (talk · contribs) - Base map: http://openstreetmap.org/?#map=18/38.715443/-9.143841 Labels: http://www.gpiaaf.gov.pt/upload/processos/d054239.pdf, ODbL, https://commons.wikimedia.org/w/index.php?curid=174827233
aftermath: By infoLibre - YouTube: Así fue el accidente de un funicular en Lisboa (Time: 26s) – View/save archived versions on archive.org and archive.today, CC BY 3.0, https://commons.wikimedia.org/w/index.php?curid=174400430
broken cable: By Gabinete de Prevenção e Investigação de Acidentes com Aeronaves e de Acidentes Ferroviários (GPIAAF) - https://www.gpiaaf.gov.pt/upload/processos/d054239.pdf, Public Domain, https://commons.wikimedia.org/w/index.php?curid=174735778
operator's view: By Fabaklava - Own work, CC BY 4.0, https://commons.wikimedia.org/w/index.php?curid=174383672
concrete structures: AI-generated (Gemini)



scsc.uk/scsc-127J

Joining The Dots



In this, the 3rd episode of a series of articles exploring the principles of ontological modelling, the Ontology Working Group (OWG) builds on previous episodes, introducing some new terms from the Unified Foundation Ontology (UFO) and expanding further on the worked example relating to flights and flight planning.

With these articles, the OWG aims to make the topic more accessible through a series of connected newsletter articles. Each episode expands on ontological aspects and by the end, will hopefully provide the reader with enough knowledge to read and contribute to models in the future.

What Is an Ontology?

As we mentioned in earlier articles, an ontology is not some esoteric formalisation, which is nice to have, but is essential to people's day-to-day work on real systems.

When we build a system, we are modelling our world and this involves developing concepts, terms and their relationships – this is called an ontology, or more formally: "**an explicit specification of a conceptualization**" [1].

If not explicitly specified, establishing a shared understanding of interrelated concepts will be much harder and miscommunication may be more common. For the system to operate correctly, any model used to support a part of the lifecycle of a system needs to be correct and consistent. Hopefully the concepts in the model are defined in writing, but at worst, meaning can just be held in the minds of the designers. An ontology also serves as a contract of meaning between stakeholders, who make what are known as 'ontological commitments'. That is, they commit to use vocabularies in a way that are consistent with the ontology's theory. [1]

Where Did We Get To? – A Recap

We are now on our 3rd episode (1st in May 2025, and 2nd in Oct 2025). This article may still be of interest if you have not read the previous two, but you will get a lot more value out of it if you read them first.

To recap on the earlier articles, the OWG uses the **Unified Foundation Ontology (UFO)** [2]. This is tailored for ontological modelling and implemented in **OntoUML** [3] (a customisation of the Unified Modelling Language, UML [4], with associated software support via application plugins). The following notation was previously introduced.

- We use boxes and connecting lines (more formally **Classes** and **Relationships** using the modelling terminology, which are highlighted in this section in **bold**) to represent concepts in our worldview. For readability, classes mentioned in the text of this article are shown in **bold purple**.
- Ontological aspects are expressed as **Stereotypes** (the part of the class definition with Guillemets like this: **«role»**). **Stereotypes** mentioned in the text in this article are in **bold blue**.
- Arrows with a closed white triangle are **Generalizations**, or in other words, one 'thing' subtypes the more general 'thing' that is being pointed to (for example, a **Cessna** is a subtype of **Aircraft** and **Aircraft** is a **Generalization** of **Cessna**)
- An open arrow ⇒ shows an **Association** between things. For example, in the phrase *"Pilot 'flies' Aircraft"*, the concept 'flies' would be modelled as an **Association**. Note that **Associations** also have **Stereotypes**.
- **Association end**s are the numerical expressions at the end of each association, such as **0..*** and **1..*** (read as *zero or more* and *one or more* respectively). These show **Multiplicity** – the range of the number of each type of thing that can be related. This is used to help express constraints, such as a **Pilot** can be trained to fly one or more **Aircraft**, but a given **Aircraft** can usually be flown by, at most, two **Pilot**s at any one time.

In the last episode, we covered the important concept of identity principles: these form sets of individuating characteristics that, in combination, give us the ability to separate one individual from another (individuate, more formally).

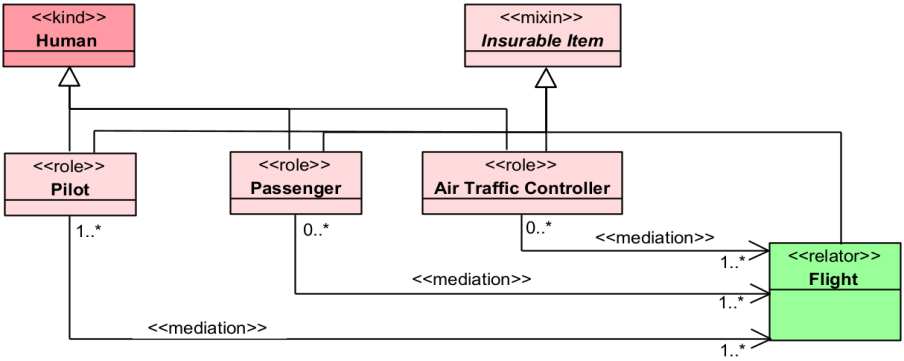
We have also covered some core UFO concepts, expressed as **Stereotypes**:

- **Kind**s reflect the objects we find around us like people, aircraft and dogs that share a common identity principle.
- **Role**s and **Phase**s reflect things that **Kind**s might be during their lifetime.
- **Mixin**s are used to describe shared properties of objects that do not all belong to the same **Kind** and so can span multiple **Kind**s that do not share the same identity principles.

We used an example model of a **Flight** to help explain the concepts and we'll further expand this model in this episode.

Our Story So Far

We've developed our story over the last two episodes to arrive at the following figure:



There are a lot of lines at the top of the diagram, but they are simply saying that:

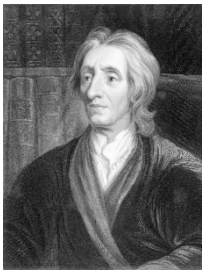
Pilot, **Passenger** and **Air Traffic Controller** are roles of **Human**

Pilot, **Passenger**, **Air Traffic Controller** and **Flight** are also all specialisations of ***Insurable Item***, even though they do not share a common identity principle. This is represented using the **Mixin** stereotype to indicate that ***Insurable Item*** can apply to a 'mix' of different subtypes. Note also that ***Insurable Item*** is written in italics, which is the UML convention to indicate a type that cannot be directly instantiated, known as an abstract **Class**. Abstract classes can only be instantiated through one of their subtypes.

The **Multiplicity** tells us some interesting things too. A **Flight** must have a least one **Pilot** (1..*), it does not always need to have an **Air Traffic Controller**, and it does not necessarily need to have a **Passenger** (0..*), although it might have many. It is also the case that a **Passenger** can only exist if it is associated with at least one flight (no flight means no **Passenger**s). Note also that this model will be over constrained for some scenarios because, for example, the minimum of 1 **Pilot** means that unpiloted flights by drones cannot be modelled. It could also be under constrained for a model of commercial flights, which usually must have a copilot.

Making Things Count

Kinds, **Role**s and **Phase**s are all examples of concepts that allows us to individuate and count their instances. These concepts have the higher-order categorisation of **Sortal** (higher-order categorisation will be coloured **bold green** here). We can thank John Locke for coining the term in 1690! There are a couple of other important higher-level categorisations of UFO concepts that we will introduce later.



Examples of **Sortals** are: **People**, **Runways**, **Computers**, **Locks**, **Aircraft** and **Companies**. **Role**s and **Phase**s are **Sortals** too – we can certainly individuate and count, for example, the number of **Broken Seat**s (a phase of Seat) and **Students** (a role of a Human).

If you can't individuate and count instances of a class, it is a **NonSortal**. Examples of **NonSortals** are generally more abstract, such as adjectives like **Loud**, substances like **Metal** or **Fuel**, and characterising qualities like **Redness**.

Are we saying we can't count **Fuel** and **Metal**? That would be a problem if we want to ensure an aircraft made of **Metal** has sufficient **Fuel** to reach its destination! We can count **Metal Fuselages** and the **Tanks** that hold the **Fuel**, but if we consume some **Fuel** then what remains is still referred to in the singular as the **Fuel**. In such cases, instead of counting we have to measure amounts of the substance. To solve this problem, OntoUML provides a **Quantity** stereotype that we can use. This provides identity and so is also a **Sortal**.

Close Encounter of the Subkind

We've said that a **Kind** represents the things in our world concepts that provide an identity principle for their instances and do not require a relational dependency. Because **Kind**s are **Sortal**s, we can individuate the instances of that particular **Kind**. There are, however, often many things that we consider to be different, but are all readily identifiable as belonging to a common **Kind**. For example, a **Boeing 737 Transport Aircraft**, **Airbus 320 Transport Aircraft** and a **Bombardier Dash 8 Transport Aircraft** are all **Aircraft** and share the same identity principle.

In these circumstances, it is useful, not least for convenience (avoiding repetition for example), to indicate that a class has an overarching identity provider. We therefore introduce **Subkind** as a new stereotype that allows us to build classic taxonomies of things. **Boeing 737** can then be a **Subkind** of **Boeing Transport Aircraft**, which in turn can be a **Subkind** of the **Kind Aircraft**.

Rigidity

We have touched on the idea of **Roles** and **Phases** being temporal properties of a particular **Kind**, an example might be an **Airborne Aircraft** being a **Phase** of the **Kind Aircraft**. The **Aircraft**, being a **Kind,** will remain an **Aircraft** throughout its existence, so like all other **Kind**s and **Subkind**s, it is a **Rigid Sortal** – whereby, 'rigid' should be understood to refer to properties or classifications that are considered essential and unchanging for an entity across all possible situations and times.

By contrast, if a concept includes temporal properties, that is, properties that exist only for part of the thing's lifetime, then that is called **AntiRigid**. So, **Airborne Aircraft**, being a **Phase**, is **AntiRigid**. In fact, because individual **Airborne Aircraft** can be distinguished and share identity principles, they are also a **Sortal**. As such, a **Phase** is referred to as an **AntiRigid Sortal**. Other examples of **AntiRigid Sortal**s include: **Student Pilot** (**Role**), **Airline Staff** (**Role**) and **Healthy Human** (**Phase**).

Thinking about identity and concepts such as **Sortal**s and **Rigid**ity may seem to add additional complexity but, as before, this provides structure to how we organise and reason with things in our world, so we don't end up modelling things that don't make sense. It can also help us to improve the amenability of our models to automated semantic checking, which is becoming essential for complex system safety. There are also additional constraints that apply, such as:

- A **Rigid** type can't be a subtype of an **AntiRigid** type. For example, a **Computer** cannot be a subtype of '**Operational Computer**' as it is still a **Computer** even when it is switched off.

There are many other rules that apply, and these can thankfully be encoded in modelling tools, so we are guided towards building consistent models. Those with a deeper interest may wish to explore the full OntoUML specification [3].

Maintaining Good Relations!

Any model, such as an architectural design, needs more than just a collection of objects to give it meaning. It is only through the conceptual relationships between those objects that we really harness the expressive power of our models.

In the example above we saw that a **Flight** was a relationship between **Pilot**, the **Passenger** and the **Air Traffic Controller.** A **Relator** is the thing that must exist for two or more individuals to be connected in a material way. Other examples of **Relator**s are: **Marriage**, **Investigation**, **Enrolment**, **Employment** and **Subscription**. As the **Relator** exists for the entire time the relationship exists, then it is also **Rigid** as we saw with our **Flight Relator**. In fact, all **Relator**s are actually **Rigid Sortal**s as you can individuate and count this type of relationship.

We can now expand our model to include a further **Relator**: **Flight Authorisation**. The **Flight Authorisation** is the authorisation given by the **Air Traffic Controller** for the **Pilot** to follow the associated **Flight Plan,** which includes the specific **Flight Route** to be taken for the **Flight**.

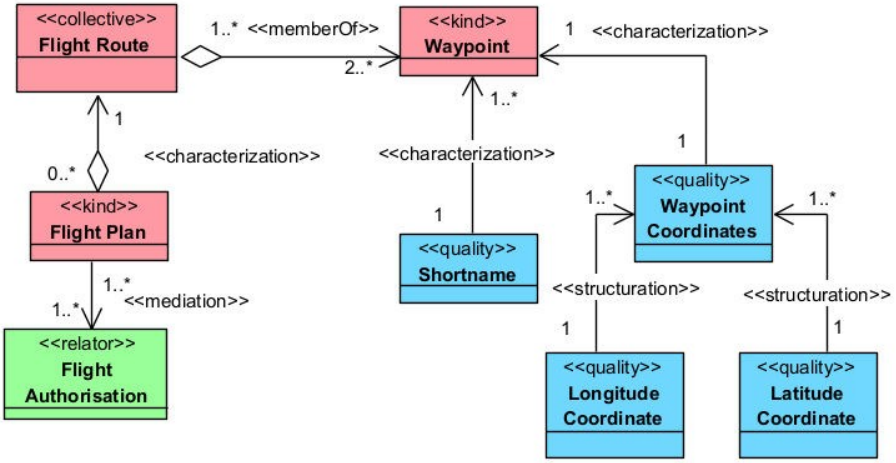
Note that **Flight Plan** will comprise many things but as we are interested specifically in exploring **Flight Route** we've modelled that explicitly. This is shown as a **Shared Aggregation** relationship with a hollow ◇ diamond at the containing element end, to indicate the **Flight Plan** contains a part known as **Flight Route**. This form of **Aggregation** relationship allows the two things it connects to have different lifecycles and multiple membership – the **Flight Route** for this particular **Flight Plan** can exist in isolation from this **Flight Plan** and may well be adopted by other **Flight Plan**s, possibly at the same time.

While multiple **Flight Plan**s may contain the same route, additional constraints may be required to ensure that two **Aircraft** do not fly the same **Flight Route** at the same altitude at the same time.

But what is a **Flight Route** in reality? In many cases it is a collection of **Waypoint**s which are representations of physical locations on the Earth the Pilot intends to navigate between.

We now introduce another useful stereotype to help us model a collection, using the stereotype **Collective**. This is used to bundle similar objects that are all perceived as contributing equal parts to the whole. Our **Flight Route** is therefore a type of **Collective** that has a number of **Waypoint**s as its members[2].

Waypoints have exactly one set of **Waypoint Coordinates**, which is modelled as a combination of a **Latitude Coordinate** and **Longitude Coordinate,** being alphanumeric values, representing physical locations. For example, in the chart at the beginning of the article, for the **Waypoint** with **Shortname** **ETOXU**, the **Latitude Coordinate** is **N47 43.6** and the **Longitude Coordinate** is **E009 33.0**. Note there is no altitude component to our **Waypoint Coordinates**, which exist at all altitudes at once (so more than one aircraft can pass **Waypoint Coordinates** at the same time, assuming they honour the rules of separation).



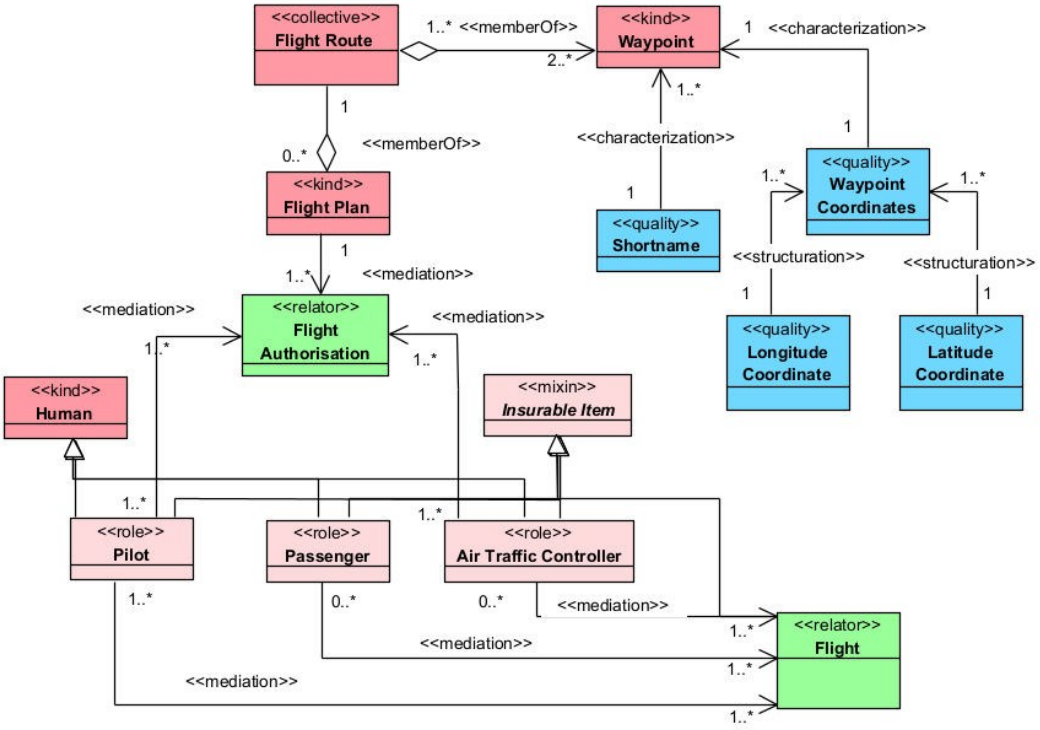
To model the **Latitude Coordinate** and **Longitude Coordinate** component of our **Waypoint Coordinates** we use another concept called **Quality**, which has three subtypes. A **Perceivable Quality** of an object is one that you can measure like **Weight**, **Height**, **Colour**, **Position** or **Speed**. A **NonPerceivable Quality** is one that cannot be directly measured by an instrument, like **Passenger Satisfaction**. Finally, a **Nominal Quality** is used to make reference to an individual or quality, like a **Human Name**, a **Book Title** or a **Boarding Card Number**.

[2] In the diagram, classes below Flight Authorisation have been omitted for clarity but still form part of our model.

Note that, importantly, we are, at the moment, modelling only the representation that **Pilot**s use in their **Flight Plan**s. In our example, **Shortname**, **Waypoint Coordinates**, **Latitude Coordinate** and **Longitude Coordinate** are all modelled as **Nominal Quality** because they are character strings on a chart. If we were to extend the model to include the physical points on the Earth that the strings represent, these points would be modelled as measurable **Perceivable Quality**s.

Up until now, we've only focused on class stereotypes, but there are also many stereotypes for relationships, of which we've already seen **Mediation** (how a **Relator** mediates between things) and now we introduce **Characterization** (how a **Quality** relates to an object). For example, **Shortname** has an association relationship with the **Waypoint** shown using the **Characterization** relationship stereotype. We'll introduce just one more useful relationship stereotype: **Structuration**. This allows us to provide structure to a **Quality** so **Waypoint Coordinates** can be split into the textual representations of **Latitude Coordinate** and **Longitude Coordinate**.

In our model, **Waypoint Coordinates** has one and only one **Latitude Coordinate** and one and only one **Longitude Coordinate** value and can only be associated with a single **Waypoint**. As mentioned previously, altitude, is not relevant to our **Waypoint** as so is not a component of our **Waypoint Coordinates** representation. Also, a particular **Waypoint** has exactly one short abbreviated name (**Shortname**).



We can now congratulate ourselves on expanding our model to include **Flight** planning aspects. The model is still fairly simple, but we can already use it to communicate and reason about our worldview.

The model forces us to think about the multiplicity of the relationships we've established, such as the number of **Waypoint**s in a **Flight Plan**. We probably need a minimum of two; the place we are starting from and the place we will land (even if this is the same location), and that might actually be OK for a really short flight. Longer flights will of course have more and in fact, in real air traffic control domains, navigational organisations may add additional **Waypoint**s to help safely control flight movement across their airspaces.

We should also permit the same **Waypoint** to appear more than once in a given **Flight Route**. This might be required for "there and back" type flights or maybe a circular, or figure of eight type route.

What We've Learned

Here's a recap of what we've covered in this edition:

- We recalled why ontological modelling is important and applies to everyone's day-to-day endeavours and refreshed our memories on what we learned in the previous article covering **Classes**, **Stereotypes** and identity principles.
- We learned the term **Sortal** to refer to categorisations that allow us individuate and count their instances and explored the temporal phases of things through their **Rigid**ity
- We learned about **SubKinds** and explored relationship stereotypes by providing an example of a **Flight Authorisation** covering the **Relator** class stereotype and the **mediation** relationship stereotypes
- We finally expanded on the example to introduce the idea of a **collective** of **Waypoints** forming the **Flight Route** and how **Quality** can be used with the stereotyped **Association** **characterization** to describe what individuates a **Waypoint**. We also introduced **structuration** to decomposed different aspects of a **Quality**

This episode hopefully shows the power and value of ontological models: they make us be very explicit and clear about the relationships we are trying to articulate and if we find difficulty in modelling our ideas through an ontological model, then it may be a sign that they are not fully developed yet.

As well as articulating what we care about, ontologies also help show deficiencies in our thinking, and this is perhaps even more important from a safety risk management perspective.



We will leave you with a couple of questions to consider, which explore how robust our model is at the moment:

- In reality, an individual **Human** cannot be both a **Pilot** and a **Passenger** on the same **Flight** – does the model prohibit this situation?
- Does the model prohibit **Waypoints** located at different **Waypoint Coordinates** from having the same **Shortname**? If not, do you see any problems that might arise?

In the next episode we will answer these questions and explore the idea of an *anti-pattern* – if we think about a design pattern as being a common arrangement that has worked well in the past, then an anti-pattern can be considered an arrangement that frequently seems at first like a good approach but has flaws and so should be avoided.

References

[1] T. R. Gruber, A Translation Approach to Portable Ontology Specifications, in Knowledge Acquisition, 5(2):199-220, Academic Press, 1993, http://www-ksl.stanford.edu/kst/what-is-an-ontology.html

[2] G. Guizzardi, A. Botti Benevides, C. M. Fonseca, D. Porello, J. P. A. Almeida, T. P. Sales, Ufo: Unified foundational ontology, Applied ontology 17 (2022) 1–44, https://doi.org/10.3233/AO-210256

[3] OntoUML Specification, Marek Suchánek, https://ontouml.readthedocs.io/, accessed Jan 2026

[4] Unified Modelling Language, https://www.omg.org/spec/UML

[5] European Geostationary Navigation Overlay Service, https://en.wikipedia.org/wiki/European_Geostationary_Navigation_Overlay_Service, accessed Sept 2025

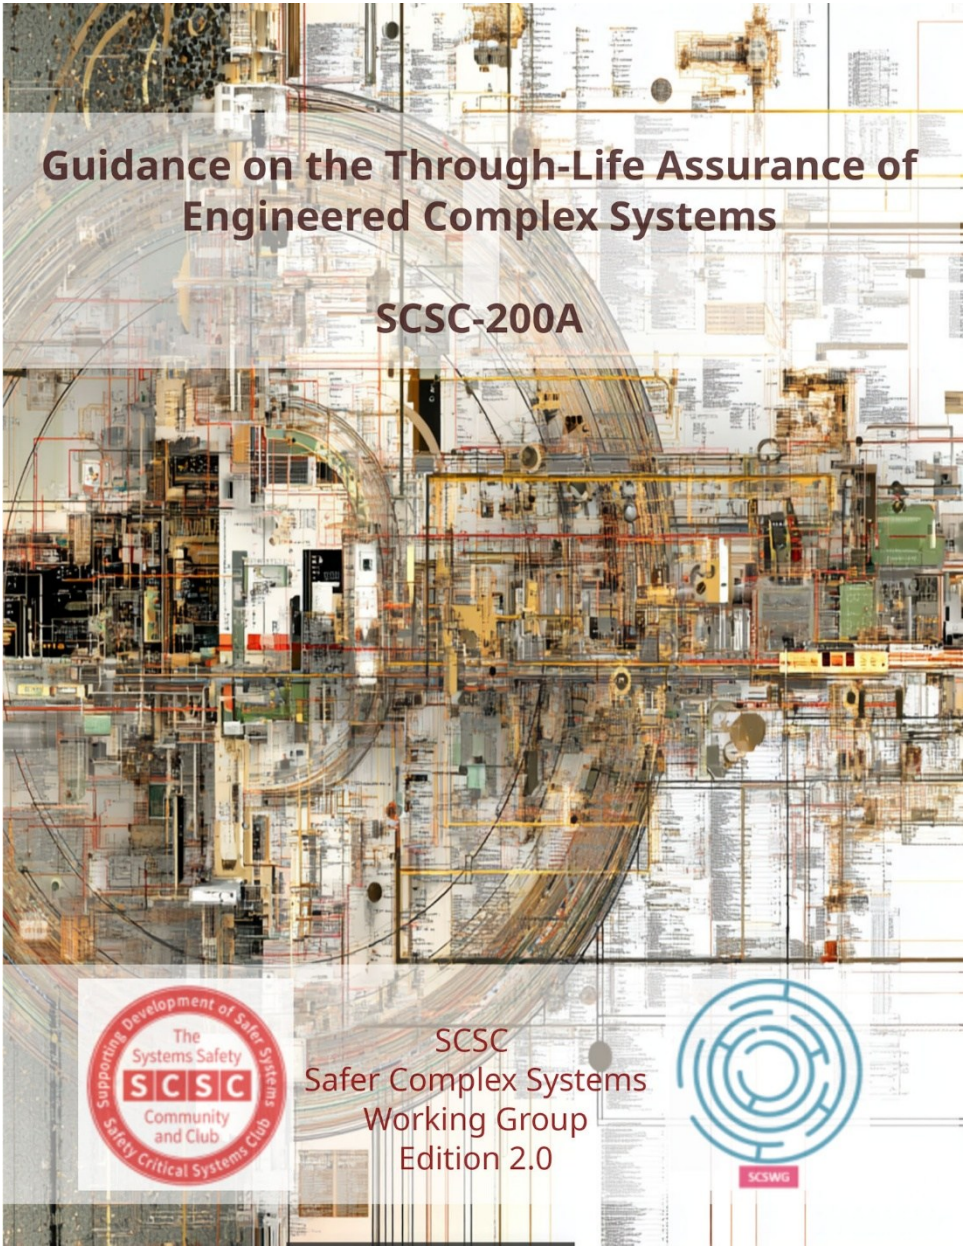
Image attributions:
navigation mao: El-mejor, CC BY-SA 3.0 https://creativecommons.org/licenses/by-sa/3.0, via Wikimedia Commons
John Locke: ID 19445045 © Georgios Kollidas | Dreamstime.com

The Ontology Working Group (OWG)



The OWG was set up in 2020 to help address issues with the lack of formalisation in terms and conceptual relationships used in standards and guidance documents. In particular, the similar but at times different, interpretations of terms and concepts used between safety and security domains is an ongoing cause of confusion and misunderstandings and presents barriers to more formal integration of these disciplines.

The OWG has its origins in the Data Safety Initiative Working Group (DSIWG), which, between 2016 and 2019, explored the development of an ontology model for the concepts described in the Data Safety Guidance (DSG). However, it was realised that the work couldn't be completed without a wider ontological model for risk, and so the OWG was established and is currently working on developing a clear, systematic terminology, and unambiguous model of risk and value concepts and their relationships.



scsc.uk/scsc-200A

Safety By Design: Reflections from an Early-Career Engineer



When Beatriz Coutinho graduated with a Master's degree in Aerospace Engineering in 2022, she expected a clear and straightforward transition into an engineering role. Like many other graduates, she imagined herself in a technical role, where she would be designing complex systems, applying all the theories she had spent years learning, and being at the forefront of innovation. The reality proved itself to be more nuanced and far more instructive. Beatriz offers her perspective on entering this profession, some of the lessons she has learned so far, and how safety principles have come to shape not only how she works, but how she thinks.

In the two years since my graduation, my career has taken me through various detours. From roles that felt outside my field to a position that has introduced me to the world of systems and functional safety. Along the way, I have learned that a career is an evolving system and that the best engineers aren't defined by the roles they start in, but by how they learn, adapt, and grow.

"The best engineers aren't defined by the roles they start in, but by how they learn, adapt, and grow"

Starting Out: From study to practice

The transition from education to employment often feels abrupt – like you're walking through fog and are not quite sure what's the next right step. Academic work provides a strong technical foundation but can leave uncertainties about how that knowledge translates into professional practice.

In my own case, my first post-university positions were not traditional engineering roles. They sat on the periphery, involving skills like project coordination, data management, and operational support. At the time, I felt like I was drifting away from my goal, yet they proved to be invaluable and formative experiences.

"Technical understanding alone is not sufficient, engineering outcomes are a product of systems, processes, people, and culture working in concert"

Those roles introduced me to how organisations function, how to communicate across disciplines, and how to adapt quickly. Those so-called "non-engineering" roles taught me that technical understanding alone is not sufficient, engineering outcomes are a product of systems, processes, people, and culture working in concert. Understanding that interplay has shaped my perspective as I have transitioned into a career in systems and functional safety.

Discovering Safety Engineering

My formal introduction to safety came when I joined **3SK**, a consultancy specialising in systems engineering and functional safety, as a Graduate Project Engineer. Although safety had featured in my studies only in passing, it quickly became central to my daily work, I was introduced to new languages, concepts and frameworks.

Safety engineering is humbling, it asks us not simply to understand how systems operate but to anticipate how they might fail, and to consider the consequences of that failure. It demands technical rigour, but also imagination and empathy. It's about considering not just the systems, but the people that interact with them (e.g., operators, maintainers, and end users).

What struck me most was that safety work is not about achieving perfection, but about designing systems that fail safely, recover predictably, and keep people protected. It is about learning from the evidence, questioning assumptions, and maintaining transparency.

This mindset extends beyond my day-to-day role, it influences how I approach challenges even in my personal life, more as opportunities rather than barriers.

Working as a Woman in Engineering

One thing that became clear early on, is that engineering remains a profession in which women are still underrepresented. During my studies, this disparity was visible in lecture theatres and, when transitioning to the industry, this became evident in meeting rooms and office spaces.

Initially, I viewed this imbalance as something to quietly navigate; I was often cautious in technical discussions, hesitant to voice uncertainty or to question opinions. However, with time and experience I came to realise that confidence is not a prerequisite for contribution but a result of it. Engaging fully, sharing ideas, and asking questions are acts of participation that strengthen both the individual and the team and help normalise the presence of women's voices in technical dialogues.

Representation isn't about being the "only woman in the room" but to ensure you're not the last.

I have been fortunate to work within an organisation that values inclusion and actively promotes it. This support has reinforced my belief that diversity is not simply an ethical aspiration, but a safety asset. In safety-critical engineering, diversity of thought leads to stronger analysis, fewer blind spots, more resilient systems and innovative problem-solving approaches. Fostering psychological safety is therefore as vital to system safety as any technical measure.

Safety in Practice and in People

Over the past two years, I have come to appreciate that safety is as much a way of thinking as it is a discipline. It requires modesty, precision, and continuous reflection. Working on safety-critical projects has reminded me that systems just like people are dynamic; they adapt, they interact, and sometimes they fail. I often associated mistakes with incompetence, but in safety, errors are treated not as personal failings but as data points, simply as new opportunities to understand, improve, and strengthen the system.

Safety work constantly reminds us that no system, procedure, or individual is infallible. My path hasn't been perfectly linear, and I doubt it will ever be. I have had moments of overconfidence, times when I hesitated, and occasions where I underestimated the complexity of a problem. Each instance became part of my personal growth, reinforcing the importance of communication, traceability, and learning.

The longer I work in this field, the more I see that safety is as much about people as it is about processes. Behind every requirement, every standard, and every risk assessment lies the intent to protect lives and that purpose gives meaning to the detail.

Equally, safety depends on culture. The technical frameworks provide structure and discipline, but culture determines how effectively they're applied. I have seen how small acts, like taking time to explain a concept to a colleague, asking a question in a review, acknowledging uncertainty, or simply listening to others allows both safety systems and teams to function far more effectively.

"safety by design extends beyond the system itself it encompasses the people and relationships that sustain it."

In this sense, safety by design extends beyond the system itself it encompasses the people and relationships that sustain it.

Key Lessons Learned

- Growth is iterative: Engineering careers, like systems, evolve through feedback and refinement. Each project adds insight that informs the next.
- Safety starts with awareness: The earlier a potential issue is identified, the more manageable it becomes
- Confidence is built, not granted: You rarely feel "ready" before contributing. Taking part, even when uncertain, is how confidence is built
- Diversity strengthens safety: Inclusive teams are not only fairer but demonstrably safer
- Culture matters: Psychological safety and technical safety reinforce each other
- Keep learning, always: Safety is an evolving discipline and continuous learning is its defining principle

Continuing the Journey

Although my time in industry has been relatively short, I see a pattern emerging, one defined less by titles or tasks and more by mindset. I have learned that safety is not a static goal but a continual process more like a mindset, and that the safest systems are those designed with humanity in mind.

Engineering, at its best, is about designing environments where people and systems can operate safely, effectively, and with confidence. As I continue to build my career, I hope to keep contributing to both kinds of safety: the technical integrity of the systems we create, and the cultural integrity of the teams that make them possible. The challenge, and the opportunity, lie in designing systems that are not only safe in operation but created within environments where safety is lived and valued.

Beatriz Coutinho, Project Engineer, 3SK, MEng, MIET, MWES



Beatriz Coutinho is a Project Engineer at 3SK specialising in systems engineering and functional safety. With a master's degree in Aerospace Engineering, they focus on safety analysis, V&V strategies, and the application of modern engineering methods to emerging technologies. Their work includes research, technical writing, and the development of engineering solutions aimed at improving real-world safety and system performance.

The author retains copyright of this article.

Seminar: Complying with the EU AI ACT



The "Complying with the EU AI ACT: What is needed for Products and Systems?" seminar was held at the DoubleTree by Hilton in Brussels, Belgium, on 9th October 2025 and hosted by Mike Parsons and Davy Pissoort. This seminar covered the EU AI ACT itself and discussed the implications for organisations and engineers working with AI technology.

This event was the SCSC's second seminar to be held on mainland Europe after the highly successful seminar (Safe Autonomous Transport - the Good, the Bad and the Ugly scsc.uk/e1134) that was held the previous year in Munich. This time the seminar was held in Brussels, Belgium with local support provided by Prof. Davy Pissoort who helped arrange speakers for the event.



As with the previous trip, all the audio/video equipment needed to be brought by train from the UK so the UK organising team met at St Pancreas for another Eurostar trip. The team were also joined by Karin Rudolph who was one of the speakers based in the UK.



The hotel was only a short train journey from the main Brussels Midi terminal and proved to be an excellent venue for the seminar.

The room was spacious and the coffee break snacks and lunch were excellent and catered for all dietary requirements.

The hotel had an Art Deco feel while retaining the Belgian heritage with waffles and other local savoury treats available at breakfast.

Mike Parsons opened the seminar describing the challenges of AI and Large Language Models (LLMs) and how these technologies are presenting challenges to system safety. Davy Pissoort, a professor at KU Leuven university, then took over to introduce the speakers. He said they have several projects working on AI safety, and that standardisation in AI is a significant concern for many companies. There were six talks in total:

- **Jelle Hoedemaekers**, Agoria - The AI Act: A stress test for standardisation
- **Jan De Bruyne,** Professor IT law at KU Leuven and Head of the Centre for IT & IP Law - The Regulation of Artificial Intelligence under the AI Act and Liability – Challenges and Ways Forward
- **Karin Rudolph**, AI Ethics and Governance consultant and Founder Collective Intelligence - Who's Accountable? Ethics and Liability in the New AI Regulatory Landscape
- **Thor Myklebust and Dorthea Mathilde Kristin Vatn**, SINTEF Digital - The AI Act and The Agile Safety Plan
- **Isabella Ferrari**, Professor at Università degli Studi di Modena e Reggio Emilia - The Protection of Intellectual Property in the Training of Artificial Intelligence: Balance and possible Implications
- **Mathias Verbeke,** Faculty of Engineering Technology, Flanders Make@KU Leuven - From Standards to Practice: Bridging the Technological Gaps in EU AI Act Compliance Towards Safer AI Systems

Jelle Hoedemaekers



Jelle Hudemaekers presented the first talk of the day on AI standardization related to the EU AI Act. He explained the importance of standards as "unsung heroes" of our technological lives, giving examples such as the standard for container dimensions which dramatically improved global trade and safety.

He differentiated between de facto industry standards and formal standards developed by Standards Development Organizations (SDOs) like ISO, IEC, CEN, and CENELEC.

He explained the New Legislative Framework (NLF) used in Europe, where high-level requirements are in legislation while technical specifications are in standards. The AI Act follows this approach.

For high-risk AI systems, implementing harmonised standards will provide "presumption of conformity" with the AI Act, making market access easier.

He described the standardization process and outlined the key challenges in AI standardisation:

- Applying a physical product framework to digital products
- Issues with standards accessibility (Malamut case)
- More mandatory requirements in European standards than international ones
- Tight timelines with August 2026 compliance deadline
- Lack of clarity on which standards companies need to implement
- Criticism about standardization process transparency

He emphasised the need for more experts to actively participate in standards development rather than being critical of the work "from the sidelines".

Jan De Bruyne



Jan de Bruyne was next up and gave a presentation on the EU AI Act and related liability issues. He first introduced the AI Act which was adopted in 2024 and builds upon previous ethical guidelines for trustworthy AI.

He said the AI Act is part of a broader regulatory landscape that includes ethics, technology, and market mechanisms. Some provisions became applicable in February 2025, but most will apply from August 2026 with Type 1 high-risk AI systems have until August 2027 to comply.

The Act has a risk-based approach covering:

- Prohibited AI systems (e.g., social scoring systems)
- High-risk AI systems (allowed but with strict requirements)
- Systems requiring transparency (e.g., chatbots, deepfakes)
- Minimal risk systems (not heavily regulated)
- General-purpose AI models (with separate requirements)

For high-risk AI systems he said there were two categories of AI system making a distinction between those that are:

- Safety components of products requiring third-party conformity assessment
- Used in specific sectors for specific purposes

Requirements in the act include risk management, data governance, technical documentation, transparency and human oversight. Importantly, it was noted that there is a reversal of the burden of proof from the victim having to prove defective operation to the manufacturer having to prove correct operation.

Jan concluded with some liability issues:

- Distinction between contractual and extra-contractual liability
- Revised Product Liability Directive now includes software as a "product"
- Focus on "legitimate safety expectations" as a standard
- Challenges with burden of proof for victims of AI harm

Jan emphasized that while the AI Act sets high-level requirements, standards (as discussed by Jelle) will be crucial to make these requirements concrete and implementable for companies.

Karin Rudolph



After a coffee break we heard from Karin Rudoplh from Collective Intelligence.

Karin's presentation focused on accountability, ethics, and liability in AI systems. She began by discussing the complexity of determining responsibility when AI causes harm, using examples like autonomous vehicle accidents and misleading AI chatbots.

She said that the new Product Liability Directive, now includes software and digital products, making it challenging to prove causation between AI defects and harm. She covered ethical frameworks that could be applied to AI:

- Utilitarianism (consequence-based ethics)
- Deontology (rule-based ethics using principles like human rights)
- Virtue ethics (focusing on developing good character)

She then gave examples of ethical approaches in AI:

- The "Utilitarian Self-Driving Car" dilemma (People accept it is better for the vehicle to sacrifice the vehicle occupant rather than kill 20 pedestrians unless they are the occupant themselves!)
- Claude's constitution based on human rights principles to build ethical and transparent AI systems

Karin went on to discuss the aspects such as:

- The concept of moral agency in AI systems and whether machines should make moral decisions
- The EU Code of Practice for general-purpose AI models, focusing on risk mitigation throughout the AI lifecycle
- The challenge of systemic risk assessment with AI that has emerging capabilities that cannot be fully predicted
- The shift from narrow AI governance to dealing with more complex general-purpose AI and AI agents
- Future considerations including the debate around AI personhood and the changing relationship between humans and AI

Karin concluded by highlighting implementation challenges, the need for agility in regulation, the changing role of engineers, and the importance of being curious and open-minded about these rapidly evolving technologies.

Thor Myklebust and Dorthea Mathilde Kristin Vatn



After lunch, the next talk had two speakers with Thor first covering AI Safety Case considerations and then Dorthea covering Human Factor aspects of AI systems.

Thor started the presentation by introducing the book he had published called "AI Act and The Agile Safety Plan", which provides guidance on how to produce Safety Plan for an agile project. It covers all the relevant topics and issues, discussing compliance in high-risk systems development and establishes a connection to the safety cases to help produce these efficiently.

Thor went on to discuss some of the projects his company has been involved in especially in helping clients gain a foothold in the marketplace. He said the development of safety plans and safety cases for AI components is in its infancy and changes are expected to the AI Act as it is open to different interpretations by stakeholders. The remainder of his talk focussed on the challenges of producing a Safety Plan such as definitions of the system, managing software including the AI lifecycle, data, biases and issues with AI training material and outliers, compliance, proliferation of unqualified tools, documentation, procurement and sub-contractors.

Dorthea presented on the human factor's perspective of safe AI development and use.

She started by saying that the human factors perspective emphasises that knowledge of how the human brain processes information should guide technology design, especially in safety-critical domains. She highlighted examples of technology failures, including car accidents with automated features where drivers blamed their vehicles, demonstrating the importance of considering human limitations in design.

She then introduced the concept of situational awareness, developed by Mika Ensley in the 1990s, which is crucial for operators who serve as the ultimate fallback mechanism when automated systems fail. She discussed Article 14 of the AI Act regarding human oversight, which requires humans to maintain ultimate control and decision-making power over AI systems in safety-critical contexts.

She said that explainable AI (XAI) is not a fixed characteristic but a capability that needs continuous development, as machine learning systems and user needs are constantly changing.

From her PhD research with a Norwegian renewable energy company, she found that developing XAI capabilities requires not only technical resources but also organisational structures, expertise, and processes for collaboration and mutual learning.

She concluded with key takeaways:

- the importance of human factors expertise
- early user focus in development
- consideration of human factors in safety-critical contexts
- the need for organisations to invest in tools, people and processes to develop XAI capabilities.

Mathias Verbeke



Matthias's presentation focused on the technological gaps in AI safety and the role of standards in addressing them. AI integration in mechatronic systems offers opportunities but creates new safety challenges. He noted the AI Safety Paradox: as AI systems become more capable, they become both harder to assure for safety but it is more important that this is done.

He said that traditional safety methods struggle with AI's non-deterministic behaviour. For example, it is difficult to incorporate all possible scenarios into model training and models degrade over time due to data drift.

While standards like those from the CENELEC Joint Technical Committee 21 (JTC 21) provide good coverage of the AI Act requirements, they lack harmony and technical depth. They explain the "what" and "why" but are vague on the "how."

He said work is in progress to develop a more coherent structure of standards to address interrelationship gaps and provide clearer mapping to standardization requirements.

There still remain technical implementation gaps with three key areas that need to be addressed:

- Safety specification (defining system purpose and associated risks)
- Uncertainty quantification (measuring model output uncertainty)
- Model monitoring (tracking performance over time and addressing degradation)

He said there are however, promising developments: Standards like ISO/IEC 5469 provide useful guidance on incorporating AI in safety-critical systems, but more detailed technical specifications are needed.

Matthias concluded by saying that he and his colleagues are actively researching formal safety specification methods, robust uncertainty quantification, and runtime safety monitoring to help bridge the gap between current standards and practical implementation.

Isabella Ferrari



Our final talk of the day was given by Isabella Ferrari who presented remotely. Isabella focused on legal issues in AI training, particularly copyright, intellectual property protection, and privacy. She explained how AI training relies on datasets that can be protected by copyright, with data falling into four categories:

- open domain
- requiring authorisation
- illegally scraped
- corporate confidential data

She discussed different types of licenses like MIT and GPL, which have varying requirements for attribution and redistribution. She then highlighted several legal cases; the first was Dr. Steven Taylor's experiments with trying to get AI systems recognised as inventors/creators (which were rejected due to AI lacking personhood).

Secondly, she covered in detail the ongoing New York Times lawsuit against OpenAI and Microsoft for using copyrighted articles for training ChatGPT without permission. This case centres on whether AI training constitutes "transformative use" under the "fair use" doctrine.

For privacy concerns, Isabella emphasised the importance of data anonymisation techniques in different sectors. In healthcare, pseudonymisation preserves useful information while removing personal identifiers. For automotive applications, she suggested using synthetic data from video games to avoid privacy issues with real-world imagery.

Mike Parsons – Discussion

Mike Parsons then hosted a general discussion. The first observation was that, in the aviation industry, there's currently no recognised way to certify AI systems for safety-critical applications since existing standards don't account for AI's non-deterministic nature.

There was consensus that clarity on AI regulation and standards might take several years, with estimates ranging from 3-4 years.

Karin highlighted the unprecedented complexity of AI systems, especially as they begin to interact with each other, creating new dimensions of challenges.

Jelle raised concerns about the pace of standards development compared to technological advancement, questioning whether standards can keep up with both AI development and the need for legislation.

There was acknowledgment that standards typically take about 10 years to develop while legal systems take even longer, creating a significant gap between technology and governance.

The overall sentiment was one of uncertainty about the future of AI regulation, with recognition that many questions remain unanswered despite the day's discussions.

Brussels Bustle!

The evening presented an opportunity to visit the centre of Brussels and experience its vibrant night life and many colourful and intriguing shops – not least its rubber duck shop! There were also a few shops with no prices of the handbags and watches being offered for sale in the window – as the saying goes, "if you need to ask the price, you probably can't afford it!"



There was also ample opportunity to shop for Belgian's other great export – chocolate! The central Grand-Place plaza we visited was particularly bustling as a local university was holding its graduation ceremony in the square. The square itself was delightful especially in early evening when all the lights of the buildings, restaurants and food stalls gave a warm and enticing ambiance to the scene.



After taking in the marvels of the gothic façade of the Town Hall and other building dating back to the 14th century, we found a restaurant where we enjoyed a 11 samples of Belgian beer, mostly of which were over 8% alcohol!





With a few free hours before our train back to the UK the next day, we returned to the centre briefly the following morning to see the 1000-year old St. Michael & St. Gudula Cathedral with its amazingly intricate pulpit.







Report by Paul Hampton, SCSC Newsletter Editor

Seminar: New Developments in Rail Safety Systems



The "New Developments in Rail Safety Systems" SCSC Seminar was held at the Eurostars Book Hotel in Munich, Germany, on 4th December 2025 and was hosted by Carmen Carlan and Mike Parsons. Although the event had 'rail' in the title, many of the talks addressed the role of Artificial Intelligence (AI) and Machine Learning (ML). These are topics that have applicability in many other sectors. Roger Rivett provides highlights of the day.

This was the second seminar held in Munich and a follow up to the successful seminar held in November 2024. Carmen Carlan was again instrumental in organising the event, arranging the speakers and helping select a suitable venue. She also co-hosted with Mike Parsons, who, along with Alex King, Brian Jepson and myself, made the journey via train from the UK.

It was again decided to travel by train and carry the high-value AV equipment. Fortunately, the trains to Germany were on time and all connections were made without the difficulties encountered during the 2024 trip. After leaving St Pancras at 11:04 we arrived at the hotel at the relatively civilised time of 21:15. However, the plans for the journey back suffered an immediate set back when the train from Munich to Frankfurt was two hours late. We did all eventually get home in the early hours of Saturday, except for Brian who had to stay over-night in London and take a train in the morning; he arrived home 24 hours after he started!

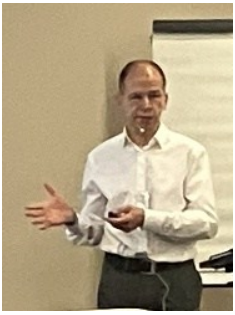
The book-themed hotel features the title of books, and the names of authors, as decoration and if you had trouble sleeping you could always read the quote thoughtfully written on the ceiling above the head of the bed. The event was once again well-served with a large room for the event and good catering, including a 3-course meal for lunch.

Mike Parsons opened the event with an introduction to the SCSC and the reason for organizing a seminar focused on rail safety systems, and the role of AI and ML in improving train operations. Mike then handed over to Carmen Carlan who introduced each of the speakers in turn.

The seven speakers presented talks discussing AI-based obstacle detection systems, machine learning graph compilers, autonomously driven trains, the operational design domain and the use of mobile mapping technology for railway infrastructure surveying. All topics were focussed on contributing to the safety and reliability of railway systems.

- **Lennart Kilian**, Siemens Mobility GmbH - safe.trAIn: AI based obstacle detection for trains - risk analysis, functional scalability, safety on system level and lessons learned
- **Jelena Frtunikj**, NVIDIA - Hardening Machine Learning (ML) Inference for Safety-Critical Applications: From Static Libraries to ML Graph Compilers
- **Carsten Thomas and Philipp Jaß**, HTW Berlin - N-Version Architectures for Railway Track Segmentation with Deep Neural Networks
- **Arnold Bitner**, Modelwise - In automation we trust - also for FMEDAs?
- **Gereon Weiß**, Automation Systems at Fraunhofer IKS - Safe Autonomy in Rail – The Role of Operational Design Domains
- **Stephen Bull**, Ebeni - Assuring Changes to a Complex System – Lessons From Experience
- **Volker Wegener**, DATADEV - Railway geodesy - 5 mm at 80 km/h: A modern method for efficient and precise surveying and digitalisation of railway infrastructure without any safety risk

Lennart Kilian



Lennart reported on work performed as part of the EU Safe.trAIn project https://safetrain-projekt.de/en/, which is contributing to the goal of 'Fully Automated Train Operation'. The benefits of this are that it addresses the shortage of drivers and provides opportunities to improve efficiency and delivery of a better service. The latter will also contribute to reducing emissions if it leads to a shift from personal to public transport.

This project focused on the detection of objects, based on the assumption that this can only be done using ML. A number of challenges were noted including that the application falls under the European AI ACT as a high-risk application and the fact that there is no safety standard for AI-based perception systems in the rail domain. There are also no clear requirements for interpretability and robustness that have to be met and no established tools and processes.

The safety objectives are to protect both passengers in the train and persons on the track. This involves the detection of objects and persons on the track. The safety target is 'overall as good as a driver' and quantified as a Probability of Failure on Demand (PFD) of 1%, although the use of PFD is controversial in European Rail Groups.

The PFD is considered to be equivalent to 1–recall, where recall =TP/(TP+FN), TP is a true positive and FN is a false negative.

The 1% PFD target is derived by comparing event trees for the driver and the object detection system. The analysis is independent of the assumed initial event probability, which means that it's the ratio that is relevant, not the absolute scale. It was observed that expanding the sensor capabilities was more effective than decreasing the PFD and that increasing the PFD to 10% is significant while decreasing it to 0.1% has little effect.

Overall, it was considered that safe perception for driverless regional trains could be achieved and that, although posing many challenges, these were already being overcome in the automotive domain.

Jelena Frtunikj



Jelena described how, since 2012, market demand has driven an evolution of neural network architectures. This has seen an increase in core tasks and representative models, combined with model complexity.

To date, these architectures have been based on the use of pre-validated static kernels. In these, the safety subset of the kernels is compiled ahead of time and each kernel is validated according to ISO 26262. The compilers are qualified and the kernel executed on an ISO 26262 safety-certified runtime environment. This has the benefit of clear, deterministic execution. However, it cannot support new operations, it cannot fuse operations for efficiency and new hardware requires months of validation. Performance is also suboptimal because it is limited to what is in the library, and generic libraries cannot keep pace with modern deep learning.

The proposed way forward is the use of an ML compiler that translates a computation graph, representing the ML neural network, into a form that can be executed on the target platform. It consists of two main components, a graph compiler and a kernel compiler that performs complex across-layer optimizations that holistically evaluate and optimize the entire network. The graph compiler takes advantage of model-aware optimization, hardware-aware specialization, operation fusion, memory optimization and scheduling and concurrency.

While this new approach improves performance, it also has benefits for safety as it moves from being tool-centric ('Is the compiler qualified?') to being output-centric ('Is the generated compiler code correct?'). The old paradigm was based on compliance with ISO 26262 for the compiler, the safety runtime and the static libraries. The result was safe but limited. With the new approach, the output of the graph compiler is validated at build time and each generated kernel is subjected to a multi-layer verification (graph level, kernel level, application level) before being executed on a deterministic safety runtime developed according to ISO 26262. The result is safe and flexible.

A question was asked about the need for a correct semantic definition of the graph compiler, as for a code compiler. This is not possible as the semantics are application-dependent; instead, validation is performed at different levels. Kernal is developed as a Safety Element out of Context (SEooC) and the user is provided with a safety manual, a safety case and tools for specific contextual validation to identify any non-deterministic behaviour.

Another question concerned having a validated graph compiler. Again, this is considered not to be possible due to the rapid development of the technology and the lack of semantics, due to optimisations across multiple measures.

Carsten Thomas and Philipp Jaß



Carsten explained that while driverless metro systems had been in operation since the late 1960s, autonomous train operation on mainline railways is a lot more challenging. It is a less controllable environment, longer braking distances are required and the network is used by different types of trains.

In the conventional approach, the train driver must observe the track, the signals, the level crossings and the overhead power lines in order to avoid hazards such as obstacles on the track, unjustified emergency braking, and persons to persons on, or close to the track. In this way, harm is prevented from occurring to persons in the train and those outside it.

To avoid obstacles, it is necessary to determine which track, the ego-track, the train will be on after crossing points. An autonomous train cannot rely solely on GPS, because it is not accurate enough, or train schedules, because that are not reliable enough. The approach taken here is to use an AI-based perception system to identify all possible future tracks and combine this with switch and signal detection and localization.

The most critical failure of the AI-based perception system is not detecting an obstacle that is present (false negative). The acceptance criteria used was that these are not more frequent than for a human train driver; this equates to a PFD < 1% (which cannot be achieved by AI alone) and falls between SIL1 and SIL2. However, this criterion is still subject to debate.

Safety is a system property, so although there is no such thing as safe deep learning, what is required are safe cognitive systems. Possible safety measures include the use of dissimilar redundancy in sensing and implementation, plausibility checks using tradition algorithms and checking against external data. The work described here considered redundancy and dissimilarity in the implementation.

Philipp noted that a single neural network is not suitable for safety-critical systems before discussing n-version architectural dissimilarity, which is a well-known method from software safety standards for non-ML software. This uses redundancy to detect and potentially correct errors. An architecture was presented in which one input is processed by three ML models with a non-ML component producing the output.

The ML models were not pre-trained as the rail domain lacks a database of pre-existing datasets. Instead, the models were trained on a newly collated data set of 15,519 images collected from multiple sources. The predictions of the ML models were combined using two different algorithms, Maximum Confidence Voting (MCV) and Pixel Majority Voting (PMV). The combination of individual confidence scores was based on the similarity of individual predictions.

Overall, the project demonstrated that the use of N-version architectures may be a promising approach to ML safety. The diversity in the Deep Neural Network architectures improves

segmentation reliability while the combined confidence score correlates better with the actual prediction quality. Pixel-level combination enables error correction.

Arnold Bitner



Arnold presented on the integration of safety engineering processes with engineering processes. He identified 5 problems with the safety analysis process:

- Missing integration between requirements, safety and product development
- Months of manual work for Failure Mode, Effects, and Diagnostic Analysis (FMEDA) per design iteration
- Accuracy depending on scarce and expensive experts capable of work
- Budget and time constraints making full analysis impossible
- Suboptimal fit of design to requirements due to over, or under, engineering

These highlight the need for automation and data integration across the supply chain and the challenges of managing complex systems. There is a need to abstract and simplify data for human review while maintaining traceability. It is also important to connect stakeholders across the supply chain and share data securely.

His proposal was to take the engineering design, the safety requirements and the failure modes and rates and combine them in a single mode. Then, to use AI to perform automated reasoning on the failure behaviour and produce an FMEDA report that can be assessed; this reduces the human workload. Such a system should be scalable, efficient, easy to use and certifiable.

Gereon Weiß



Gereon outlined the problem he was addressing as being that autonomous systems in the real world must safely operate in an unpredictable and chaotic place and that understanding and defining the context in which they operate is key.

He then introduced the concept of the Operational Design Domain (ODD) as a means to help achieve this by confining the real world to that in which the autonomous system is intended to operate.

He mentioned two automotive standards (ISO 34503:2023-08 and ASAM OPENODD BASE STANDARD 1.0.0 SPECIFICATION V1.0.0:2025-04-03.) and two examples of automotive implementation (Waymo and Mercedes).

He introduced two definitions, one commonly used for Automated Driving Systems (ADS):

> *A set of operating conditions, including, but not limited to, environmental, geographical, and time-of-day restrictions, and/or the requisite presence or absence of certain infrastructure characteristics in which a given automated driving system, or feature thereof, is specifically designed to function*

and one for AI Systems:

> *A set of operating conditions, including all entities relevant for the AI system, in which it, or feature thereof, is specifically designed to function*

The ODD should be seen as a central artifact that provides traceability and consistency for many different aspects of the development process such as architecture, safety case, training, verification and validation and performance monitoring.

ODDs can vary in subject matter and detail depending on the context in which they are used. The ODD covering development aspects needs to be detailed whereas the ODD covering safety aspects can be more coarse-grained and abstract so as not to be over-specified.

There can be application-specific refinement of the ODD, for example starting with a more permissive one for describing the system environment, a refined one for system development and assurance and one specifically refined for ML-methods.

The level of complexity of the ODD decreases in the move from manual to fully automatous systems as manual systems can cope with a more complex ODD than fully automatous systems, which tend to be designed for narrow ODDs.

Gereon than introduced a new document for the Specification of Operational Design Domain in Rail, DIN DKE SPEC 99004, which defines a taxonomy for the ODD for fully automated rail vehicles covering important operating conditions for railway-specific AI applications such as Automated Train Operation (ATO), e.g. environmental or track conditions, https://www.dinmedia.de/de/technische-regel/din-dke-spec-99004/390496147. This has a wide range of potential usage, including:

- Risk assessment for safety and security
- Reference System comparison
- System development
- ODD limit detection
- Scenario-based testing
- Conformity assessment and approval
- Intended purpose adequacy
- Reuse
- Increasing confidence
- Project planning

Gereon concluded his talk by giving three example ODDs for rail applications, each of which was described by a taxonomy, depicted graphically and also specified using YAML (a human-readable data serialisation language).

Stephen Bull



Stephen highlighted the fact that changes to a railway system had to take place while the railway was still operating and compared this to changing the engine of a plane while in flight.

A railway is a system of systems supporting the needs of many people including the train driver, the signaller, the maintainer and the customer. While being difficult to define what makes the railway complex, Stephen highlighted 15 characteristics that contributed to the notion of complexity.

He illustrated some of these with three examples from experience, the first being the high number of interfaces and legacy components with which the UK East Coast Digital Programme has to contend with. The second was the interaction between the need for very curved platforms at Tottenham Court Road and the clearance of the platform screen doors from the train on the London Elizabeth Line. The last one concerned radio communication capacity where the unexpected loss of communication resulted in the continuous attempt to regain a connection. This led to the system being swamped in what looked like a denial-of-service attack.

Stephen then considered a number of emerging challenges.

AI is both a risk and an opportunity. The key issues to be addressed for assurance being transparency, bias, traceability, cyber security and ethics.

Given the large quantity of data involved it is necessary to understand the different types and how it is used; reference was made to SCSC Data Safety Guidance document (see scsc.uk/dsiwg). An example was given where the data representing the gradient of the track has to be trustworthy because the train relies on this to determine when and how hard to break.

The link between security and safety was highlighted; the increasing attack surface is leading to an increasing number of attacks resulting in a higher risk of harm. Reference was made to the SCSC Security Informed Safety guidance (see scsc.uk/siswg).

The problem of standards becoming outdated was mentioned, along with the fact that they never reflect the latest technology by the time of publishing and that different ones often contradict each other.

Stephen summed up the key challenges as being:

- Understand the interactions and interfaces in your system
- Understand your assurance goals (e.g. for standards, AI, data)
- Know how you will integrate your system

His advice for addressing these was start early, agree a strategy, identify regulations and standards, decide how to demonstrate compliance, review regularly and plan to manage change.

Volker Wegener



Volker explained that the management of a railway requires knowledge of the infrastructure, for example the precise location of the rail, bridges and tunnels. It also needs up-to-date information concerning the state of the track and signals. There is also a need to create situation maps after floods, earthquakes and storms. Such knowledge has traditionally been collected using methods which were not efficient and disrupted the rail service. Therefore, there is a need for a more cost-effective method that does not disturb regular service and is also safe for customers and railway staff. Ideally, the information would be stored such that it can interface with existing IT systems.

To address this, the work had evaluated the use of a multi-sensor approach consisting of a Global Navigation Satellite System (GNSS), an inertial measurement unit (IMU), Light Detection and Ranging (LiDAR) and cameras mounted on a locomotive. These were used in conjunction with ground control points and a vehicle positioning systems with targets along the tracks. This was able to collect data at speeds ranging from 0 to 100 km/h and achieved a ±5 mm accuracy at 80 km/h.

The benefits of this approach include the fact that there are no persons in the track danger zone. The fast coverage of long track, with automatic detection of damage, cracks, markings and traffic signs leads to reduced inspection costs, better planning of renovations and it also minimizes downtime.



The data is collected in a format which can be used in combination with existing data sets providing a uniform database for all departments helping with improved decision-making through spatial analysis and forms the basis for smart city applications and digital twins.

The approach also helps with security and disaster management by providing mobile mapping in the event of floods, earthquakes and storms and leads to faster response times and improves deployment planning.

The overall conclusions were that multi-sensor systems are the most efficient railway surveying method providing high accuracy without operational impact and being fast, scalable and cost-efficient.

The presentations and videos for this event are available to SCSC Members, see scsc.uk/e5004 for details.

SCSC Timeline



Paul Hampton, SCSC Web Master, introduces the SCSC Safety Timeline, describing the history of the timeline, what has been done with the captured data and how you can help develop it in the future!



For those that attended the SCSC's Symposium in Feb 2024, you may recall we asked attendees to think about important historical events in the world of system safety. These might have been specific accidents or other significant event such as the publication of a standard or landmark court case etc. These were then captured on "post it" notes on a large timeline display that spanned multiple presentation boards.



We have now collated this initial set of data and enriched it to include fuller descriptions and images, and are using a timeline presentation tool to allow users to interactively browse the data. The timeline can be viewed in 3D as in the top image, and you can scroll back and forward in time. There is also a 2D view as shown above.

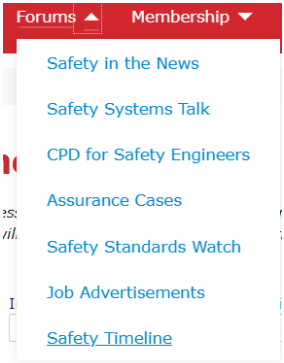
If you are interested in finding out more about a particular event then you can click on an individual item and a brief description will be presented. You can also then click 'Find out more' to be taken to an external resource, such as Wikipedia to get a full description of the event



You can see the current timeline by visiting the following link:

tiki-toki.com/timeline/entry/2211124/System-Safety-Timeline

It is also accessible via the SCSC website from the **Forum->Safety Timeline** menu option as shown.



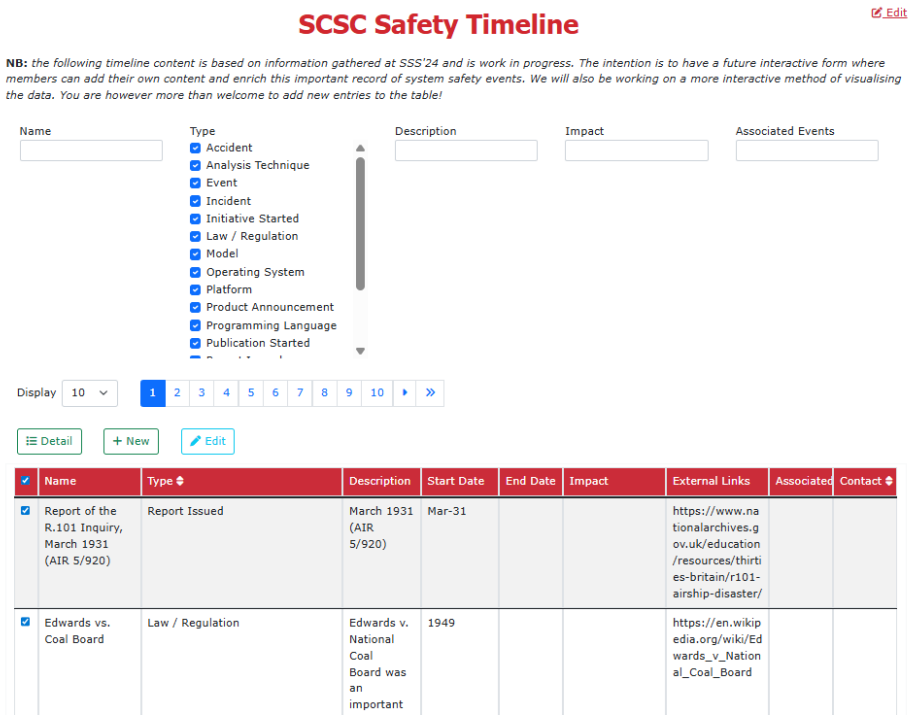
The timeline is currently viewable on an external website but in the future, we plan to embed the viewer in the SCSC's own website so we have a more consistent look-and-feel.

We have over 150 entries at the moment, but we think there are still many more events that need to be captured!

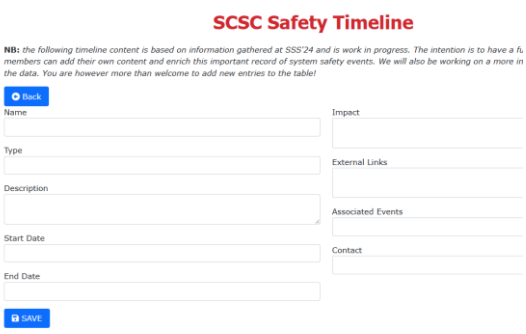
We are therefore looking for your continued support in this endeavour, so we would very much like you to continue to add items to the timeline.

To make this easier, we now have an online method for adding events. This is available to registered users of the website and is accessible through the same **Safety Timeline** menu option as shown.

Once you are logged on to the website you will also see the following table. This allows you to search the entries we have already and to add new ones. It is of course worthwhile checking that an entry has not already been captured, so you can use the **Name** and **Description** fields to search for events. You can also filter by event types by checking/unchecking the types of events you are interested in.



Click on + New the button to add a new entry. You will be presented with the following screen:



Add as much detail as you can and then press SAVE

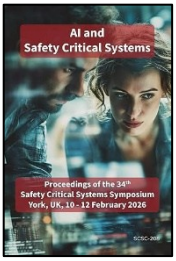
Try and be as specific with the dates as you can but a month and year or just a year will work fine too.

In External Links it would be useful if it led to a full description of the event (eg. a Wikipedia entry) and also if there is a particular image you'd suggest for the main timeline display, then a web link to it would be helpful.

You can change your entry at any time from the main table listing screen. Click the checkbox on the side of your entry and then press Edit Any of the fields you previously provided can then be amended.

The timeline is active now so please feel free to add new content! We will curate your submissions to check for duplicates etc. and then these will be added to the main 2D/3D timeline views. We very much look forward to seeing your contributions!

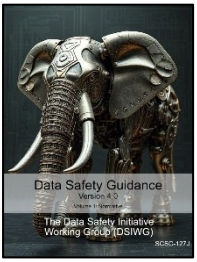
Recent Safety Publications



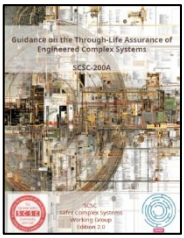
Proceedings of the Thirty forth Safety-Critical Systems Symposium
scsc.uk/scsc-208



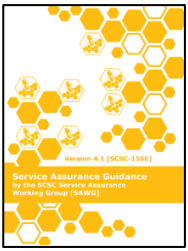
Safety-Critical Systems eJournal vol.4 no.2 Winter Issue 2025
scsc.uk/scsc-210



Safety-Critical Systems Data Safety Guidance v4.0 Parts 1, 2 and 3
scsc.uk/scsc-127K



Guidance on the Through-Life Assurance of Engineered Complex Systems v2.0
scsc.uk/scsc-200A



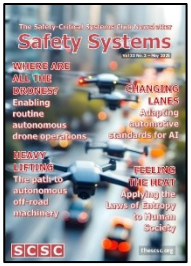
Service Assurance Guidance v4.1
scsc.uk/scsc-156E



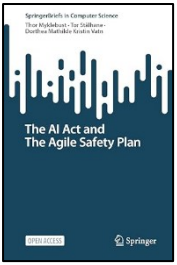
The Safety-Critical Systems Club Newsletter, Volume 33, Number 3
scsc.uk/scsc-206



Safety-Critical Systems eJournal vol.4 no.1 Spring Issue 2025
scsc.uk/scsc-204

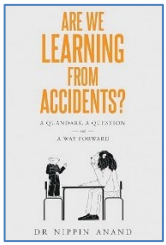


The Safety-Critical Systems Club Newsletter, Volume 33, Number 2
scsc.uk/scsc-205



The AI Act and The Agile Safety Plan
amazon.co.uk/dp/3031805038

Book Review



Paul Hampton, SCSC Newsletter Editor, reviews the book: "Are We Learning from Accidents? A Quandary, A question and a way forward" [1], by Dr Nippin Anand.

In the SCSC's last newsletter editorial, I asked whether we were still learning from accidents and I happened across this apt 2024 work from Dr Nippin Anand that asks the very same question, so I was intrigued and interested not least, to hear a fuller description of the events of the *Costa Concordia* sinking in 2012 [2].

The *Costa Concordia* was a €450 million cruise ship that grounded and partially sank after striking a rock in 2012. At the time, the ship was carrying over 4,000 passengers and crew, from which 32 died.

The ship struck rocks after the Captain, Francesco Schettino, ordered an unscheduled route change to pass within half a nautical mile of Isola del Giglio, Tuscany in order to perform a 'sail-by salute'. Schettino was vilified for causing the accident and leaving the boat prematurely during the evacuation. He was found guilty of manslaughter and sentenced to 16 years in prison.

I, and I suspect like others, thought the incident at the time was pretty much an 'open and shut case': a rogue overconfident captain going 'off-piste' to please the crowd, and I actually wondered how much could actually be learned from it.

My narrow expectation was soon dispelled – Nippin's exclusive interviews with Schettino reveal a much more nuanced and complex situation and, in fact, I realised that the *Costa Concordia* incident was really being used as a springboard to convey much more far-reaching ideas. In particular, Nippin shows how perceptions are shaped by our Western culture, our *mythology* (in the sense of the stories we believe, sometimes unconsciously, to be true) and our deeply held assumptions and cherished belief.

Examples are our 'Hollywood' culture that celebrates the heroes of our stories and vilify the anti-heroes, to always seek a 'root cause' to an accident – to find something or someone to blame, a scapegoat, and to assume human fallibility can in some way be 'fixed' through more procedures and regulations. Much to my surprise, the book helped highlight my own prejudices and how I have participated in helping promote these same entrenched cultural narratives.

Nippin's argues that there is never one 'truth' about an event and we should embrace and hold in balance multiple and nuanced narratives to truly understand why accidents happen and to get meaningful learning from them. He uses the "One Brain Three Minds (1B3M)" model based on the Social Psychology of Risk to help explore the sense-making of an individual in the lead up to and during the aftermath of an accident. This model brings in not only the rational and logical analysis of a human but also incorporates decision making from experience-based heuristics as well as automatic non-rational 'gut thinking' intuition.

This *embodied Mind*, as Nippin calls it, acknowledges that decision making is not restricted to the brain; our entire body participates in the process.

This more rounded model of a human being helps develop more insightful narratives around why humans make decisions in a given environment and situation. For example, Nippin notes that our senses do not give us an objective view of the external environment. How we perceive the world is locked in what actions are unconsciously 'afforded' to us (this important concept he calls "Affordance").

As well as describing the problems with our approaches to accident investigation, Nippin describes a concrete solution and way forward to address these issues through his iCue methodology. This brings in the 1B3M considerations and through finding connection and tensions between the three mind perspectives, provides greater insights and opportunities for learning.

Nippin argues that as much of our decision making is unconscious, true learning comes from surfacing the unconscious. Once the iCue maps are developed, linkage between the three aspects of the *embodied Mind* questions can be framed to promote greater learning. This helps dispel the perception that an investigator is asking questions to support their own agenda; to simply find someone or something to blame. For example, changing the enquiry from "Why did you use the wrong equipment?" to "What can we learn from the crew's perception of the use of the equipment?"

If you read the book, and I encourage you to do so, then I wonder if like me you will see an alternative view of Schettino's actions. While not suggesting that Schettino's behaviour was acceptable, there is an additional narrative here: *"one where Schettino absorbed all the sins of our society – cheap design, questionable regulations, a race-to-the-bottom approach to operating standards, material consumption habits. Instead of being a superhero he allowed distress and trauma to take over his decision making. In an investigation framed as a quest for heroes and anti-heroes, such a narrative is not possible".*

Let me conclude by saying that the book is very easy to read and although 448 pages, the font is quite large and includes a reproduction of the *Nozomi* Accident Report, another highly informative accident, referenced in the book and used to illustrate the use of the iCue method. My only criticism would be in the diagrams; they are black and white and some of the text is sometimes too small to read. The Kindle version of the book is however in colour and allows diagrams to be expanded.

In summary, I don't think I would be exaggerating by suggesting that everyone should read this book! Some of the ideas are, possibly controversial, calling into question deeply held assumptions and cherished beliefs as well as the long-established cultural and regulatory rituals that seem only to provide comfort and certainty to society rather than learning. It would require seismic changes in our industry to immediately adopt all of Nippin's ideas but as an industry, we already sense there is a need for change with concepts such as "Restorative Just Culture" [3]. By giving us a much broader and profound perspective, I hope this book and its methodologies can help guide us in the right direction of travel.

References

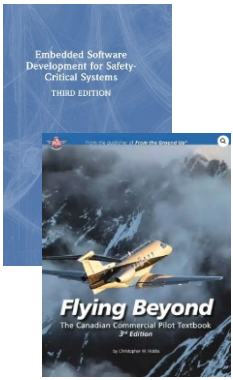
[1] Are We Learning from Accidents?: A quandary, a question and a way forward, Novellus, 25th April
[2] Costa Concordia, https://en.wikipedia.org/wiki/Costa_Concordia, accessed Dec 2025
[3] Adam Johns, Introducing a Restorative Just Culture and the Learning Review at the Docklands Light Railway, Proceedings of the Thirtieth Safety-Critical Systems Symposium, 2022

60 Seconds with ... Chris Hobbs



Chris Hobbs was first required to make a decision about shipping some potentially dangerous software in the late 1980s and has been involved with this topic since. He programs in Ada, C, Python and, more recently, Rust and particularly specialises in applying formal mathematical models to prove the correctness of algorithms and system design.

He is a member of the Canadian Mirror Committee considering version 3 of IEC 61508 and a stakeholder in the development of UL 4600. He is also a member of two SCSC working groups: one on the production of Assurance Cases and one providing guidance on the safety assurance of complex, socio-technical systems.



The third edition of his book "*Embedded Software Development for Safety-Critical Systems*" has just appeared.

Apart from software development, Chris is a pilot and the author of various aviation books including "*Flying Beyond: The Canadian Commercial Pilot Textbook*". Working with various instructors, he has also been learning the 24 songs of Schubert's Winterreise cycle for the past 24 years.

What first attracted you to working in the field of System Safety?

In the late 1980s I was working on a piece of equipment for sale to British Telecom. A technician maintaining the equipment could command it to connect test gear to the external, copper line. Before it did so, the equipment checked to see whether there was a high voltage on the external line and, if there was, wouldn't close the relays, thereby protecting the technician.

A short time before deployment, a developer reading the code noticed that, under certain unlikely circumstances, the high voltage check would be skipped. As team leader, I was under pressure to sign the equipment off: as management pointed out, a high voltage was rare and the circumstances of the test being skipped were also rare. I had none of the tools I have now to decide whether or not I should sign it off. After a diversion into high-dependability software, I returned to a focus on safety in 2009.

What aspect of your career are you most proud of?

The fact that I've never had a "career". I've had a long sequence of interesting and challenging "jobs", and have even on some occasions moved from one job to a more challenging job where the pay was less.

Tell us one interesting unusual fact about yourself!

I spent two years living in Libya when I was 11 and 12 years old – my father was a civilian engineer working at the RAF base in El Adam. This gave me the opportunity to follow dangerous pursuits (e.g. putting a match to the cordite fuses in unexploded munitions left over from the Second World War), but it also meant that I missed my Eleven-Plus examination. On returning to the UK, my father had to plead with the headmaster of Huntingdon Grammar School (Oliver Cromwell's *alma mater*) to get me a place in the second form there.

What advice would you give to yourself age 12?

When my children left home, I gave them a piece of worldly advice that should have been told to me when I was 12: "Never say you understand something when you don't".

What future changes would you like to see in the field of System Safety?

I would like to see IEC 61508, ISO 26262, EN 50716 and other standards accept the possibility of random software failure. IEC 61513 gets it right – "*Since the triggers which activate software faults are encountered at random during system operation, software failures also occur randomly.*" – but many of the standards explicitly state that all software failures are systematic. Edition 3 of IEC 61508 (due out in early 2026) has dug even deeper than Edition 2 into this hole. The misunderstanding seems to stem from inexperience with multi-threaded systems on multi-core processors, but even simple programs have been failing since the 1950s on certain (uncontrolled) inputs by users. This lack means that our safety analyses of any system containing software are fundamentally flawed.

What's your most favourite quote or motto?

"Many promising reconciliations have broken down because while both parties came prepared to forgive, neither party came prepared to be forgiven." Charles Williams

If you could learn to do anything, what would it be?

In 2002 I was offered a job with a startup company in California. I decided not to take it, but realised that if I was not going to take up that challenge, then I would need a different big challenge to compensate. With my wife as accompanist at the piano', I decided to learn the 24 songs in Schubert's Winterreise song cycle. Following weekly lessons from a variety of teachers and evening practices we have begun to explore the depth of those songs: perhaps the peak of any male singer's ambition. You may say I'm clearly a slow learner (23 years to learn 24 songs), but if I could learn to do anything, I would want to sing those 24 songs as Ian Bostridge CBE does (but in a lower key!)

I understand you train commercial pilots – what's the most important lesson for flying safely?

I have retired from the actual training, although my commercial pilot textbook is still being used (now in its 3rd edition). I think that the most important lesson is understanding the equipment (and its interface). The simulator is the place to be learning which buttons to push when approach control changes a clearance via a particular waypoint to a vectored approach, not in thick cloud with the possibility of icing! And, of course, a pilot needs to know the limitations of that equipment.

If you could be any fictional character, whom would you choose?

I assume that Oliver Mellors is excluded on grounds of good taste and so I must choose Hamlet's friend, Horatio. He seems to represent sanity, intellectual neutrality, friendship and the ability to offer difficult advice without personal bias.

How do you see the rise of Artificial Intelligence affecting the future of safety engineering?

I find it difficult to define what is meant by the term AI: in the popular imagination it seems to be becoming synonymous with Machine Learning and, in particular, with Large Language Models (LLMs). I think that, with LLMs, AI has taken a significant step backwards since the "commonsense knowledge" work (e.g., CYC and ConceptNet) of the early part of this century. The LLMs are now devouring works generated by other LLMs with no semantic knowledge of what is being produced. And I don't think that Artificial General Intelligence (AGI) is achievable.

But this leaves open questions on the safety of autonomous devices, particularly road vehicles. We have seen the accidents caused by Phil Koopman's long-tail events: the inability to recognise a woman pushing a bicycle or a woman flying through the air.

For safety-critical systems, I see the danger in AI being not with the technology, but with people's interaction with that technology. We have seen teenagers believing what AI bots are telling them and if AI is deployed to assist humans working in potentially dangerous conditions, those humans must very clearly understand the limits of the AI system.

Connect

The Newsletter and eJournal

Do you have a topic you'd like to share with the systems safety community? Perhaps an interesting area of research or project work you've been involved in, some new developments you'd like to share, or perhaps you would simply like to express your views and opinions of current issues and events. There are now two publishing vehicles for content – shorter, more informal content, can be published in the Newsletter with longer, more technical peer-reviewed material more suitable for the eJournal. If you are interested in submitting content, then get in touch with Paul Hampton for Newsletter articles: paul.hampton@scsc.uk or John Spriggs for eJournal papers: john.spriggs@scsc.uk

The SCSC Website

Visit the Club's website thescsc.org for more details of the Safety-Critical Systems Club including past newsletters, details of how to get involved in working groups and joining information for the various forthcoming events.

Facebook

Follow the Safety-Critical Systems Club on its very own Facebook page.

www.facebook.com/SafetyClubUK

X/Twitter

Follow the Safety-Critical Systems Club's X/Twitter feed for brief updates on the club and events: @SafetyClubUK

LinkedIn

You can find the club on LinkedIn. Search for the Safety-Critical Systems Club or use the following link:

www.linkedin.com/groups/3752227

Advertising

Do you have a product, service or event you would like to advertise in the Newsletter? The SCSC Newsletter can reach out to 1,000's of individuals involved in Systems Safety and so is the perfect medium for engaging with the community. For prices and further details, please get in touch with the Newsletter Editor.

SCSC Working Groups


Security Informed Safety


Assurance Cases


Service Assurance


AI / Autonomous Systems


Data Safety Initiative


Safer Complex Systems


Multicore


Safe System Architecting


Safety Culture


Ontology



Safety Futures Initiative


Systems Approach to Safety of the Environment


Safety Management & Safety Management Systems

The Safety-Critical Systems Club is committed to supporting the activities of working groups for areas of special interest to club members. The purpose of these groups is to share industry best practice, establish suitable work and research programmes, develop industry guidance documents and influence the development of standards.

Security Informed Safety



The Security Informed Safety Working Group (SISWG) aims to capture cross-domain best practice to help engineers find the 'wood through the trees' with all the different security standards, their implication and integration with safety design principles to aid the design and protection of secure safety-critical systems and systems with a safety implication.

The working group aims to produce clear and current guidance on methods to design and protect safety-related and safety-critical systems in a way that reflects prevailing and emerging best practice.

The guidance will allow safety, security and other stakeholders to navigate the different security standards, understand their applicability and their integration with safety principles, and ultimately aid the design and protection of secure safety-related and safety-critical systems.

Last year, the working group published its first piece of guidance: Co-Assurance of Safety and Security.

Lead **Stephen Bull** stephen.bull@scsc.uk

Assurance Cases



The Assurance Cases Working Group (ACWG) has been established to provide guidance on all aspects of assurance cases including construction, review and maintenance. The ACWG will:

- Be broader than safety, and will address interaction and conflict between related topics
- Address aspects such as proportionality, rationale behind the guidance, focus on risk, confidence and conformance
- Consider the role of the counter-argument and evidence and the treatment of potential bias in arguments



One of the working group's activities is the maintenance of the Goal Structuring Notation (GSN) Community standard. See scsc.uk/gsn for further details. This includes a library scsc.uk/library of relevant papers previously published by Tim Kelly, which are now hosted on the SCSC website.

In May 2021, the group published v3.0 of the standard: scsc.uk/scsc-141C

In Aug 2021, the group published v1.0 of the Assurance Case Guidance: scsc.uk/scsc-159

Lead Jane Fenn jane.fenn@baesystems.com with support from **Phil Williams** phil.williams@scsc.uk

Service Assurance

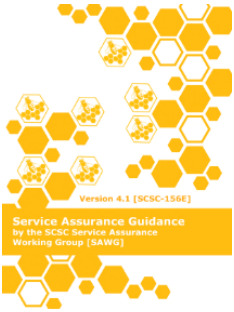


Risks presented by safety-related services are rarely explicitly recognised or addressed in current safety management practices, guidelines and standards.

It is likely that service (as distinct from system) failures have led to safety incidents and accidents, but this has not always been recognised. The Service Assurance Working Group (SAWG) has been set up to produce clear and practical guidance on how services should be managed in a safety-related context, to reflect emerging best practice.

The group published guidance v4.1 in Jan 2026: scsc.uk/scsc-156E

Lead Kevin King kevin.king@baesystems.com

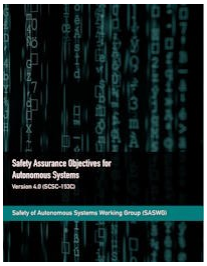


Safety of AI / Autonomous Systems



It is clear that AI and autonomous systems can introduce many new paths to accidents, and that autonomous system technologies may not be practical to analyse adequately using accepted current practice. Whilst there are differences in detail, and standards, between domains many of the underlying challenges appear similar and it is likely that common approaches to core problems will prove possible.

The SCSC Safe AI Working Group (SAIWG) aims to capture cross-domain best practice and guidance on key topics within the design, evaluation, assurance, and approval of safety systems that use or are developed using AI, bringing together emerging standards and key results from the incredible amount of research being conducted into AI safety.

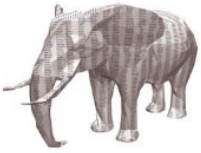


The working group was kicked-off at SSS'24 and is led by Alan Simpson. The SAIWG will conduct regular meetings, workshops, and publications to share knowledge and experience on various topics related to AI and safety systems, such as coordination of safety with other disciplines, evaluation of risk, and mapping of terminology and language.

The group builds on the earlier work of the Safety of Autonomous Systems Working Group (SASWG) that culminated in production of the 4th version of its guidance: Safety Assurance Objectives for Autonomous Systems, in Feb 2024 scsc.uk/scsc-153C

Lead Alan Simpson alan.simpson@ebeni.com

Data Safety Initiative



Data in safety-related systems is not sufficiently addressed in current safety management practices and standards.

It is acknowledged that data has been a contributing factor in several incidents and accidents to date and there is foreseeable harm that can arise from Machine Learning and Large Language Models' use of data by Artificial Intelligence (AI) systems that are subject to issues of biasing, interpretation and, arguably, falsification. There are clear business and societal benefits, in terms of reduced harm, reduced commercial liabilities and improved business efficiencies, in investigating and addressing outstanding challenges related to safety of data.

The Data Safety Initiative Working Group (DSIWG) aims to have clear guidance on how data (as distinct from the software and hardware) should be managed in a safety-related context, which will reflect emerging best practice.



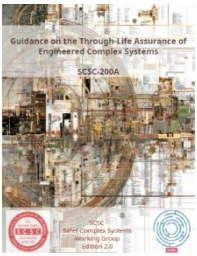
An update to the guidance (v4.0) was published in Jan 2026: scsc.uk/scsc-127K This was a major release for the guidance and saw the original document reorganised now into three more management parts that will be upissued independently.

Lead Mike Parsons mike.parsons@scsc.uk

Safer Complex Systems

The Safer Complex Systems Working Group (SCSWG) builds on the IET/RAE work already done in this area. It is recognised that the RAE work is ongoing and collaboration is encouraged. The group's mission is to produce practical guidance on developing, managing and assuring complex systems throughout their lifecycle (so as to achieve and justify their safety). The following provides a problem statement:

- It is acknowledged that complex systems are becoming more prevalent with more opportunities to cause harm
- There are also new complex systems arising from using combinations of existing systems and services which are then used for safety purposes
- Complex systems are not sufficiently addressed in current safety management practices and standards
- In particular, complex interactions and emergent behaviours are not currently assessed and managed sufficiently
- There could be benefit in developing new analysis, tools and techniques to manage complex system risks
- There are clear business and societal benefits, in terms of reduced harm, reduced liabilities and improved business efficiencies, in improved management of complex systems risk

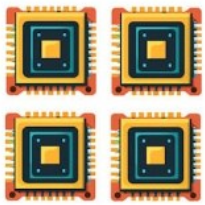


The group produced the second edition of Guidance on the Through-Life Assurance of Engineered Complex Systems (scsc.uk/scsc-200A) in January 2026. Note that the title now includes the word "Engineered" to be specific about the types of system in scope.

Contact the group lead if you would like to attend future meetings or find out more.

Acting Lead **Mike Parsons** mike.parsons@scsc.uk

Multi- and Manycore Safety



It is becoming harder and harder to source single-core devices and there is a growing need for increased processing capability with a smaller physical footprint in all applications. Devices with multiple cores can perform many processes at once, meaning it is difficult to establish (with sufficient evidence) whether or not these processes can be relied upon for safety-related purposes.

Parallel processes need to access the same shared resources, including memory, cache and external interfaces, so they may contend for the same resources. Resource contention is a source of interference which can prevent or disrupt completion of the processes, meaning it is difficult to know with a defined uncertainty the maximum time each process will take to complete (Worst Case Execution Time, WCET) or whether the data stored in shared memory has been altered by other processes.

The Multi- and Manycore Safety Working Group (MCWG) has been established to explore the future ways of assuring the safety of multi- and manycore implementations.

Lead **Lee Jacques** Lee.Jacques@leonardocompany.com

Safe System Architecting



The SCSC Safe System Architecting Working Group (SSAWG) is a joint initiative between the INCOSE UK Architecture Working Group and the Safety Critical Systems Club, with a vision of addressing systems safety through the promotion of choosing, using, and developing appropriate architecture. Two focal areas are defined as:

- Safety-driven architecting – what safety drivers are likely to exist when architecting is being considered
- Architecting system safety – what architectural factors may realise, enable, support, or preclude the achieving of safety considerations

The groups Mission Statement is:

- To produce practical guidance on addressing system safety during system architecting.
- To encourage and facilitate collaboration between INCOSE and the club.

The working group meets every 6 weeks online via Teams. Note that the INCOSE UK AWG meets every quarter.

If you would like to be involved in the group, please contact the co-chairs.

Co-chairs: **Siyuan Ji** s.ji@lboro.ac.uk and **Jane Fenn** jane.fenn@baesystems.com

SCSC Safety Culture



The Safety Culture Working Group (SCWG) has been established to provide guidance on creating and maintaining an effective safety culture. The group seeks to improve safety culture in safety-critical organisations focussed on product and functional safety, by sharing examples and latest approaches collated from real-life case studies.

Meetings provide an opportunity to discuss any particular aspects attendees are interested in taking forward, and to help set future directions for the group.

In Dec 2023, the group published a position paper for assessing and managing safety culture. scsc.uk/r2995

A seminar was held 19th June 2025 in London and focussed on how Safety Culture practices need to change to accommodate Artificial Intelligence. SCSC members can access the seminar presentations and videos here: scsc.uk/e1156

Lead **Anne Seldon** anne.seldon@wae.com

Ontology



The Ontology Working Group (OWG) develops ontologies that will form the basis of SCSC guidance, as well as having wider industrial and academic applications.

During system development, and especially with Model-Based Systems Engineering (MBSE), it is essential to build a model of the world that the system will inhabit, and this involves developing and structuring concepts, terms and their relationships. An ontology is an explicit specification of such a conceptualization and provides the foundation for MBSE and architectural descriptions in general. Ontologies are also increasingly important, if not essential, for explainable Artificial Intelligence (AI), enterprise knowledge systems and the standards that underpin them.

The OWG is currently working on the definition of an ontology of risk and value for application in guidance for risk-based decision making – notably safety and security. The framework for modelling is the Unified Foundation Ontology (UFO) and OntoUML, a domain-specific language (DSL) for modelling in UML tailored for ontological modelling based on UFO.

The OWG has also embarked on a series of SCSC Newsletter articles to gently explain ontological concepts and notations. The aim is to provide an accessible introduction to the domain so readers can be more informed and to help demystify a topic that may have a perception of being quite academic. The third article in the series can be found in this edition.

Lead Dave Banham ontology@scsc.uk

The Safety Futures Initiative

The Safety Futures working group meets once a month to bring together fresh perspectives, innovative ideas and insights.

The Safety Futures is diligently working on developing a comprehensive roadmap to guide new safety professionals in exploring career paths across various industries. To supplement this, they have an upcoming initiative to create a University 'League' Table to rank universities based on their safety engineering courses and the career pathways they support.

They are also looking to start a reverse mentoring programme where we can match more experienced people with new professionals for information exchange.

The group encourages anyone to get involved to help shape the future of safety engineering careers.

Lead Giles Howard giles.howard@scsc.uk

Systems Approach to Safety of the Environment



The Systems Approach to Safety of the Environment Working Group (SASEWG) is a new group intending to apply Systems Safety practices to systems that are embedded within the natural environment, while focussing on that environment.

The group aims to produce clear guidance on how engineered systems should be developed and managed throughout their entire lifecycle so as to preserve, protect and enhance the environment.

Please get in touch with the working group lead if you would like to join or find out more about this group.

Lead **James Inge** james.inge@scsc.uk

Safety Management and Safety Management Systems



This new working group aiming to cover Safety Management and also Safety Management Systems (SMSWG) is now in full swing with almost 60 members and held its fourth working group meeting on 15th December 2025.

Safety Management can be a difficult challenge with many different stakeholders with their own objectives pulling in different directions, perhaps one might draw parallels with the equally challenging task of herding cats! This analogy has given rise to the group's icon as shown.

Being an effective Safety Manager requires a balance of technical knowledge, regulatory understanding, soft skills, and leadership. Yet many are asked to lead safety without prior experience or clear support. This group will explore challenges such as overcoming organisational resistance, gaining senior buy-in, building trust across the workforce, and creating a positive safety culture where reporting is safe and encouraged. The group will also explore regulatory frameworks, going beyond compliance to fully understand the intent behind the rules.

The group is led by Si Hays so get in touch with him if you want to find out more about the group.

Lead **Si Hays** si.hays@scsc.uk

SCSC Membership

The SCSC provides a range of services to the System Safety community including seminars, tutorials, leadership events, specialist topic working groups, the annual symposium and a comprehensive body of publications. Membership brings many valuable benefits such as free access to online events, the SCSC Newsletter and access to presentations and other resources from events.

Individual Membership

To become an individual member of the SCSC please register on the SCSC website using the icon at the top right of any page and select "Register". Complete and save your account registration and then verify your email address. Once registered and logged in click the link "why not join the SCSC..." inviting you to become a member at the top right of the page or select "Pay membership" from the icon.

Individual membership can be paid online using a credit/debit card through our secure payment partner Realex Global Payments or contact Alex King for other payment methods. For student or retired member rates please contact Alex King to get your account status changed.

Corporate Membership

Your company contact with the SCSC should arrange the membership and any renewals for your organisation. To join as a member covered by a corporate membership, register as per the instructions for an individual member and then contact Alex King to confirm your affiliation.

Renewing Membership

You should be notified by email when your membership is almost expired or shortly after it has expired. These notifications will contain a link to the online renewal page or you will be able to renew when logging onto the website through the 'click to renew' link.

Membership Fees

The following fees are applicable from January 2026 for new and renewing members:

- 1 year Individual Membership: £159
- 2 year Membership: 7% discount: £295
- 3 year Membership: 17% discount: £395
- 1 year SFI Membership: FREE for first year, £35 for years 2 & 3
- 1 year Membership, retired member rate: £35
- For Corporate Membership discounts contact Alex King

A one-month Publication Pass is also available for £15. This allows access to all SCSC website publications in a particular calendar month.

Contact Alex King using office@scsc.uk

The SCSC Steering Group

Current Members


Stephen Bull
stephen.bull@scsc.uk


Carmen Carlan
carmen.carlan@scsc.uk


Dewi Daniels
dewi.daniels@scsc.uk


Dai Davis
dai.davis@scsc.uk


Jane Fenn
jane.fenn@scsc.uk


Paul Hampton
paul.hampton@scsc.uk


Giles Howard
giles.howard@scsc.uk


James Inge
james.inge@scsc.uk


Brian Jepson
brian.jepson@scsc.uk


Graham Jolliffe
graham.jolliffe@scsc.uk


Alex King
alex.king@scsc.uk


Kate McDougall
kate.mcdougall@scsc.uk


Mark Nicholson
mark.nicholson@scsc.uk


Yvonne Oakshott
yvonne.oakshott@scsc.uk


Wendy Owen
wendy.owen@scsc.uk


Mike Parsons
mike.parsons@scsc.uk


Davy Pissoort
davy.pissoort@scsc.uk


Karin Rudolph
karin.rudolph@scsc.uk


John Spriggs
john.spriggs@scsc.uk


Sean White
sean.white@scsc.uk

Note that honorary SCSC membership has replaced honorary Steering Group membership.

Club Positions

The current and previous (marked in italics) holders of club positions are as follows:

Managing Director
Mike Parsons 2019-

Tim Kelly 2016-2019
Tom Anderson 1991-2016

Steering Group Chair
Dewi Daniels 2024-

Roger Rivett 2019-2024
Graham Jolliffe 2014-2019
Brian Jepson 2007-2014
Bob Malcolm 1991-2007

Programme & Events Coordinator
Mike Parsons 2014-

Chris Dale 2008-2014
Felix Redmill 1991-2008

Manager
Alex King 2019-

Honorary Solicitor
Dai Davis 2022-

Newsletter Editor
Paul Hampton 2019-

Katrina Attwood 2016-2019
Felix Redmill 1991-2016

University of York Coordinator
Mark Nicholson 2019-

Website Editor
Paul Hampton 2025-

Brian Jepson 2004-

eJournal Editor
John Spriggs 2021-

Administrator
Alex King 2016-

Joan Atkinson 1991-2016

Diversity, Equity and Inclusion (DE&I) Lead
TBC

Louise Harney 2024-2025
Wendy Owen 2023-2024

Safety Futures Initiative Leads
Giles Howard 2025-

Khadijah Khatun 2024-2025
Zoe Garstang 2019-2024
Nikita Johnson 2019-2021, 2023-2024

Calendar

January

M	T	W	T	F	S	S
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	

February

M	T	W	T	F	S	S
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	

March

M	T	W	T	F	S	S
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31					

April

M	T	W	T	F	S	S
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30			

May

M	T	W	T	F	S	S
				1	2	3
4	5	6	7	8	9	10
11	12	13	14	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	31

June

M	T	W	T	F	S	S
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30					

July

M	T	W	T	F	S	S
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	31		

August

M	T	W	T	F	S	S
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31						

September

M	T	W	T	F	S	S
	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30				

October

M	T	W	T	F	S	S
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	

November

M	T	W	T	F	S	S
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30						

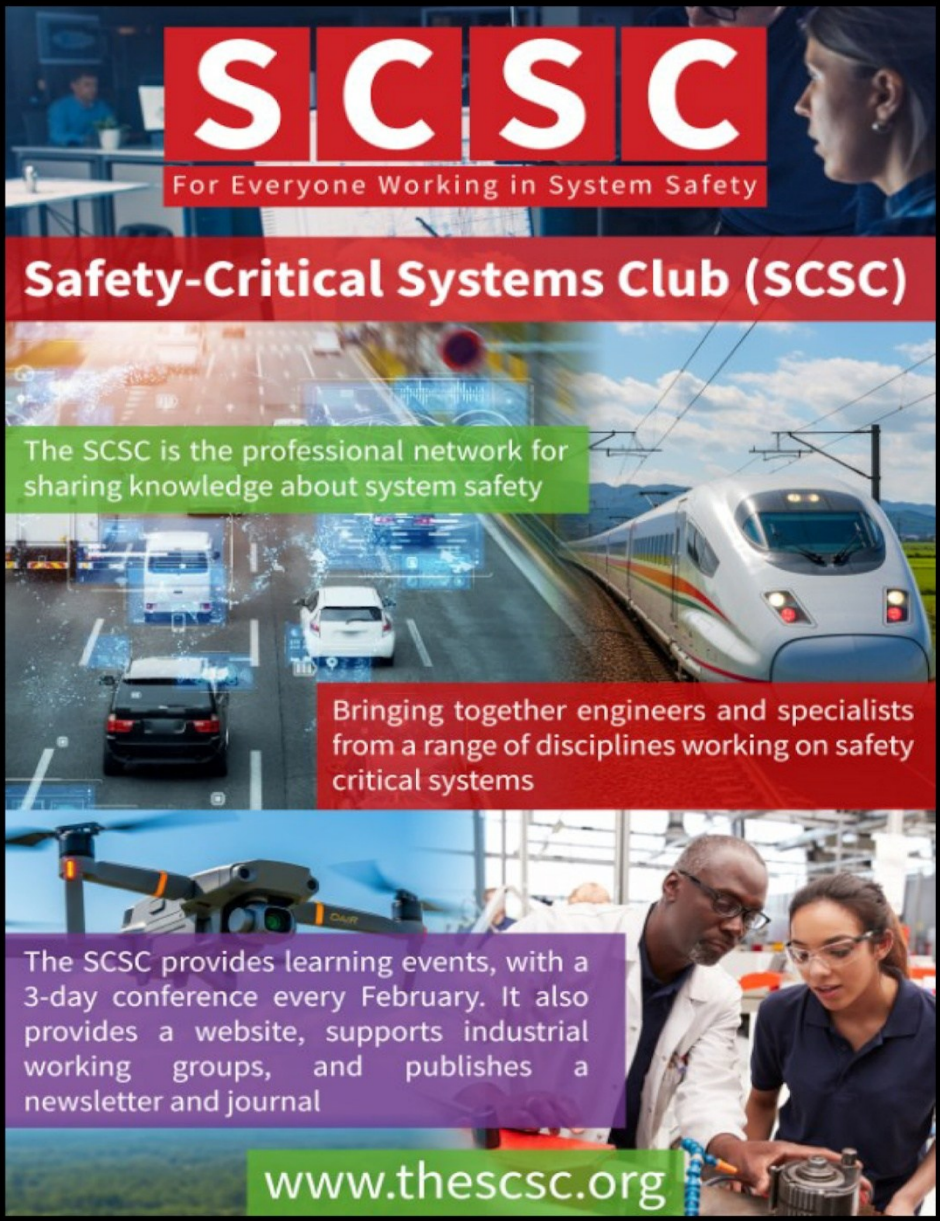
December

M	T	W	T	F	S	S
	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31			

Events Diary



27 Jan 2026 Working Group Meeting **DSIWG#96** **Online** scsc.uk/events-diary/working-groups/dsiwg-96	**29 Jan 2026** Working Group Meeting **SCSWG#24** **Online** scsc.uk/events-diary/working-groups/scswg-24	**4 Feb 2026** Working Group Meeting **SSWAG Update** **Online** scsc.uk/events-diary/working-groups/sswag-update	**10-12 Feb 2026** SCSC Symposium **Safety Critical Systems Symposium SSS'26** **York, UK** scsc.uk/sss'26
24 Feb 2026 Working Group Meeting **OWG#102** **Online** scsc.uk/events-diary/working-groups/owg-102	**25 Mar 2026** Working Group Meeting **SAWG#61** **London, UK** scsc.uk/events-diary/working-groups/sawg-61	**23 Apr 2026** SCSC Seminar **Soft Skills for Safety Practitioners** **London, UK** scsc.uk/e5055	**22-25 Sep 2026** The 45th International Conference on Computer Safety, Reliability and Security (SafeComp 2026) **Universitat Politècnica de València (UPV), Spain** www.his-conference.co.uk



Supporting Development of Safer Systems
The Systems Safety SCSC Community and Club
Safety Critical Systems Club