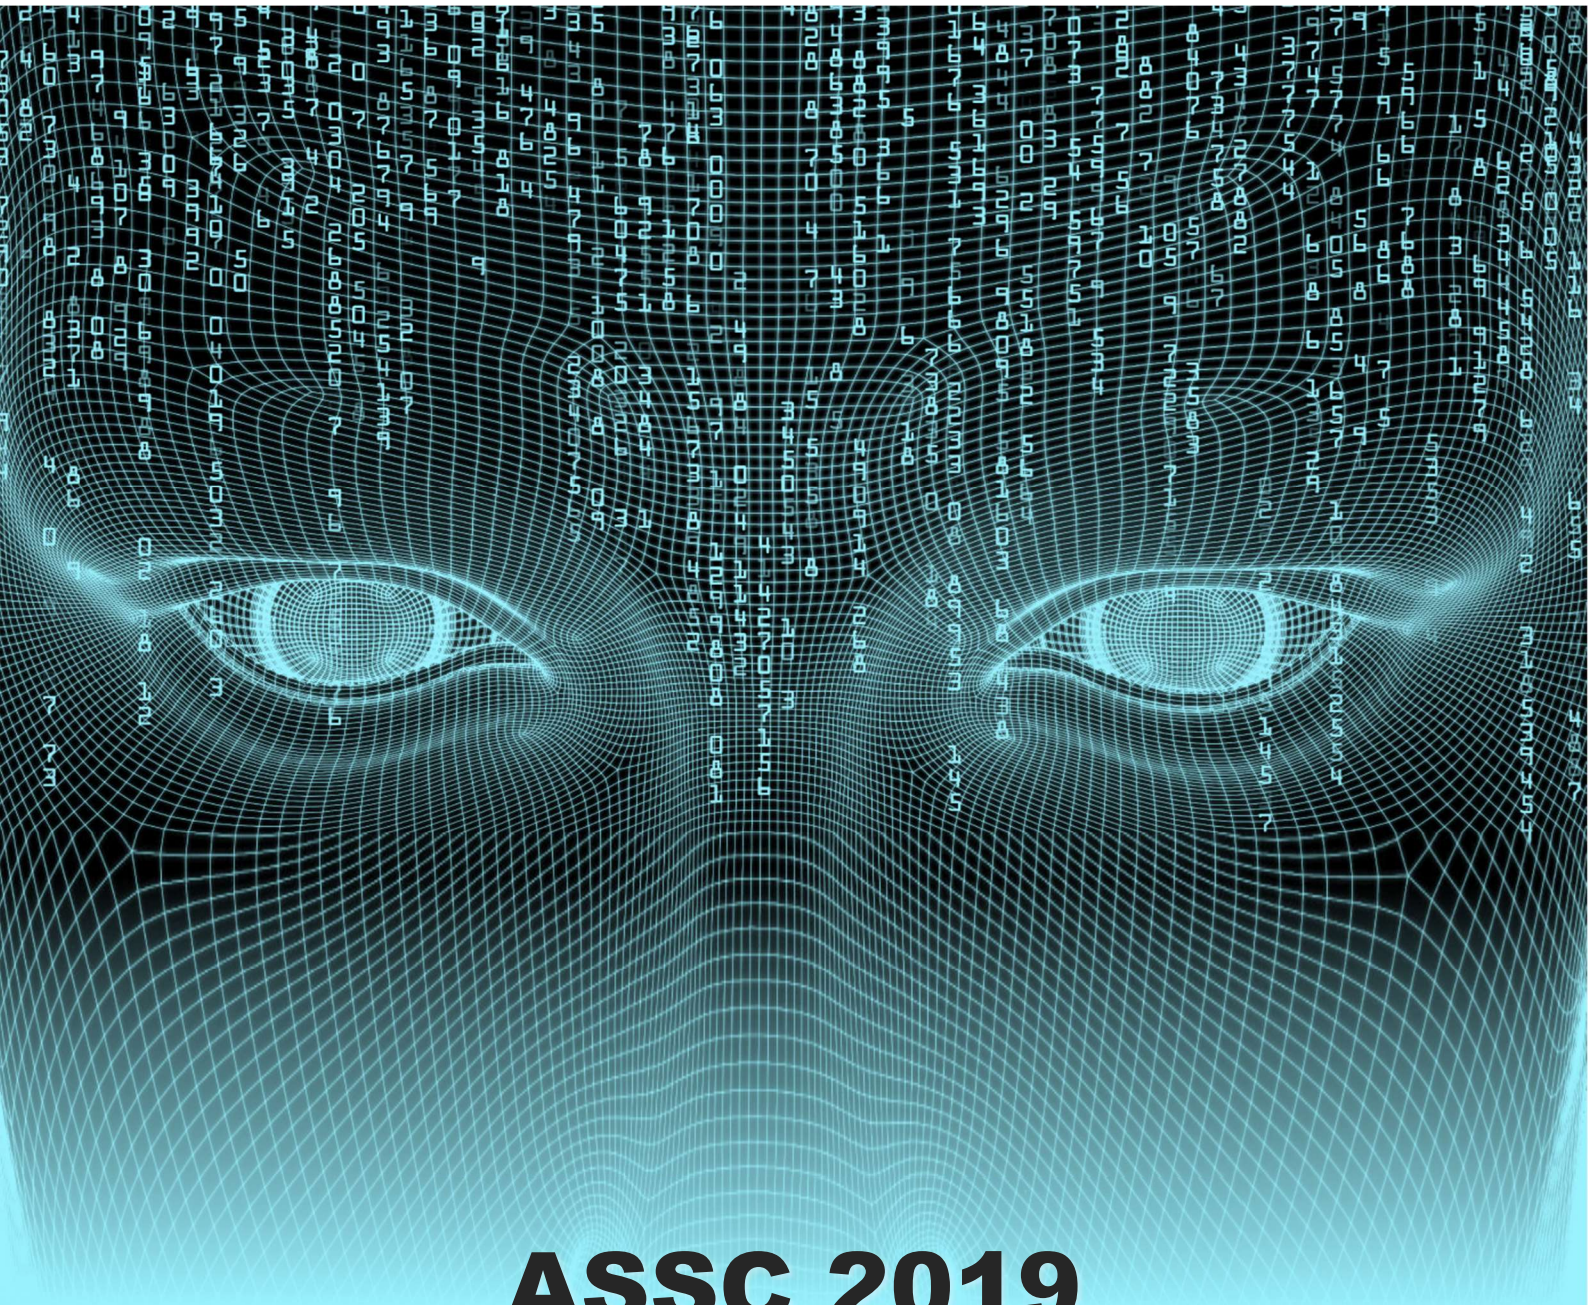

ASSC 2019

Conference Program

How can we be Artificially Safe?

Brought to you by

Australian Safety Critical Systems Association

A National Special Interest Group of acs

SPONSORS FOR ASSC2018

Australian Government

Department of Defence

RGB ASSURANCE

Nova Systems

Experience Knowledge Independence

www.NovaSystems.com

Dedicated Systems

Tutorial Abstracts

	Tutorial Day 22 May 19	Conference Day One 23 May 19	Conference Day Two 24 May 19
8.00	**Registrations**	**Registrations Open, Tea/Coffee**	
8.30	Jin-Song Dong, Zhe Hou, Hadrien Bride Griffith University *Machine learning for dependable decision-making*		
8.45		**Welcome & Introductions**	**Welcome & Introductions**
9.00		Keynote: Prof Phil Koopman Carnegie Mellon University *Autonomous Vehicle Safety and Perception Robustness Testing*	Keynote: Pam Melroy Director of Space Technology and Policy *Acceptable Risk in Human Space Flight – An Astronaut's Perspective*
9.45		Paper: Eryn Grant Acmena Group Pty Ltd *More than meets the AI: can systems thinking leading indicators assist proactive safety in artificially intelligent systems?*	Presentation: Dr. Holger Becht RGB Assurance *The Ironies of Automation with AI*
10:15		BREAK – Tea/Coffee	BREAK – Tea/Coffee
10.30		Keynote: Kelvin Ross Chairman KJR *Can We Trust Artificial Intelligence?*	Keynote: Ganesh J. Pai *NASA Research Park, SGT Inc* *Dynamic Assurance Cases: A Pathway Towards Trusted Autonomy*
11.00	BREAK – Tea/Coffee		
11.15	David Ward *Safety of the intended functionality (SOTIF; ISO/PAS 21448) in road vehicle automation*	Presentation: Ben Merange RGB Assurance *Machine Learning for Rail Safety Incident Classification*	Presentation: Achim Washington RMIT *Research Summary – Application of Bayesian methods and normative decision theory to aviation safety regulatory process*
11:45		Sponsor Presentation: Dedicated Systems	
12.00		*LUNCH*	*LUNCH*
13.00		Invited Speaker: Bijan Elahi Medtronic Technical Fellow *Is Artificial Intelligence in healthcare doomed, or destined for greatness?*	Keynote: David Ward *Horiba-Mira* *Standards for road vehicle automation: Are we nearly there yet?*
13:45	*LUNCH*	Paper: Kevin Anderson Systra Scott Lister *Functional Safety Assessment of Train Management Control System*	Presentation: Dr. Tim McComb RGB Assurance *Artificial Intelligence and Safety Standard Compliance: Challenges*

Tutorial Abstracts

14.15	Phil Koopman *Introduction to Critical Systems & Automotive Software Safety Issues*	BREAK – Tea/Coffee	BREAK – Tea/Coffee
14.45		Presentation: Drew Rae *Griffith University* Safety Assurance for Artificial Intelligence is Futile (and that's okay)	*Presentation:* Jon Sciortino *Nova Systems* *Autonomy: Is It Really As Safe As We Think? – Practical Examples from the Real World*
15.30		Paper: Graham Hjort 4Tel Pty Ltd *Next-Generation Rail Systems using Artificial Intelligence and Machine Learning*	Keynote Panel Session *"How can we be artificially safe"*
16.15		*CLOSE*	CLOSE
17:00	*CLOSE*		

Tutorial Abstracts

TUTORIAL DAY

WEDNESDAY 22ND MAY 08:30 – 17:00

Machine learning for dependable decision-making

Machine learning (ML) has been very successful in prediction, classification, regression, anomaly detection and other forms of data analytics. ML is becoming an integrated part of automated decision-making for critical systems. However, most existing ML techniques are used as black-boxes and they do not provide a high level of trustworthiness. In fact, there have been numerous cases where ML-based applications failed and caused tremendous damage. To address the security and safety concerns of critical systems, advances in trust-related aspects of machine learning are important. Explainable artificial intelligence (XAI) is one way to improve trustworthiness and it has attracted much attention recently. Machine learning approaches that are capable of explaining the rationale behind the predictions are more relatable and transparent. Another way to improve trustworthiness is to develop ML techniques that can produce auditable predictive models. The verification of these models provides formal guarantee that the models are correct, safe and secure with respect to user's requirements.

In this tutorial we will cover recent developments in the domain of transparent and auditable machine learning techniques. We will introduce our latest research combining advanced machine learning, high performance computing and automated reasoning techniques. We will also present the fruit of our work: Silas -- a state of the art ML toolkit for dependable machine learning.

Zhe Hou, Hadrien Bride & Jim Song Dong - Griffith University

Zhe Hou obtained PhD from The Australian National University in 2015. He has developed various techniques and tools for verifying computer programs. In 2015, he joined Nanyang Technological University to undertake the project “Securify” (5M USD) with Singapore Defence Science Organisation. He was the main researcher responsible for developing formal models for the SPARCv8 instruction set architecture and weak memory model, and verifying information-flow security for the formal models. He joined Griffith University in 2017 on a DST funded project to develop trusted autonomous systems in collaboration with Australia Defence Science and Technology Group. Zhe's most recent research concerns the explainability and trustworthiness of machine learning algorithms. He is helping Dependable Intelligence develop a new machine learning tool named Silas with a focus on explainability and formal verification of prediction models.

Tutorial Abstracts

Safety of the intended functionality (SOTIF; ISO/PAS 21448) in road vehicle automation

Since original publication of the road vehicle functional safety standard ISO 26262 in 2011 it was quickly noted that due to wider system safety issues not being in scope, guidance was needed on addressing additional factors which could influence safe operation of automated driving features. The concept of SOTIF was originally conceived to address failures in driver assistance systems (ADAS or Level 1 / Level 2 automation) associated with sensor performance limitations, for example "false positive" triggering of automatic emergency braking caused by a vehicle radar acquiring an incorrect target. The first iteration of this guidance was recently published as ISO/PAS 21448.
In this tutorial we will examine

- Background and context of automated driving features in road vehicles
- Functional safety, SOTIF and the wider system safety context
- Brief overview of ISO/PAS 21448
- SOTIF "area" concept for managing complexity and unknowns of automated driving scenarios
- Introduction to SOTIF approach to different "areas" of scenarios
- Case studies
- Conclusions and future outlook towards Level 3 and higher levels of automation.

David Ward - Horiba Mira Ltd

Dr David Ward is General Manager, Functional Safety at HORIBA MIRA Limited, a leading independent provider of automotive engineering services. Dr Ward is a recognized international expert in automotive functional safety with 25 years' experience in safety and reliability of embedded electronic systems in automotive and other industries. Within his role at HORIBA MIRA, David is responsible for training, consultancy and independent safety assessment in the functional safety standard ISO 26262 and other related standards. He is involved in automotive cybersecurity as well as technology development in functional safety, connected and autonomous vehicles, and vehicle electrification.
He is the UK Principal Expert to ISO/TC22/SC32/WG8 "Road vehicles – Functional safety" which has developed ISO 26262 and ISO/PAS 21448 "safety of the intended functionality"; a member of the ISO/SAE joint working group developing ISO/SAE 21434 "Road vehicles – Cybersecurity engineering"; as well as contributing extensively to the UK's MISRA initiative.
David was presented with the IMechE Prestige Award for Risk Reduction in Mechanical Engineering following 20 years of work leading automotive industry efforts to develop international safety standards. These efforts began with MISRA (The Motor Industry Software Reliability Association) and have continued with the development of the international standard ISO 26262, which Dr Ward and his team at MIRA continue to influence as work considers how to extend its scope to highly automated vehicles.
David is also Visiting Professor of Functional Safety at Coventry University, UK and RAE Visiting Professor of Industrial Design at the University of Leicester, UK.

Tutorial Abstracts

Introduction to Critical Systems & Automotive Software Safety Issues

Over the past two decades the automotive industry has dramatically increased the use of life-critical computer based control systems. However, because there are no regulatory requirements that mandate the use of software safety standards and good practices, the results have been uneven. This tutorial will discuss a case study of vehicles that did not conform to many accepted safety practices and how that eventually lead to an adverse legal verdict and costs of well over a billion dollars. That case study motivates a discussion of critical system practices and safety requirement techniques that are applicable to both conventional vehicles and autonomous vehicles. Finally, a discussion of the role of attribution to driver error reveals that autonomous vehicles will not only change how people drive, but also require a significant overhaul of the US automotive safety regulatory system. Tutorial modules are selected from material in a graduate-level course on embedded system safety taught at Carnegie Mellon University, and include:

- A first-hand account of what really happened in the Toyota Unintended Acceleration cases (both legal & technical)
- Critical system assurance principles applied to large automotive fleets
- A safety envelope-based approach to safety requirements
- Historical perspective on blaming the driver and how that will affect future regulation in potentially surprising ways

Prof. Philip Koopman - Carnegie-Mellon University

Prof. Philip Koopman started working on autonomous vehicle safety over 20 years ago with the Carnegie Mellon University NavLab project. Recently he has done stress testing and run time monitoring of robots and autonomous vehicles. He currently works on technical, policy, and regulation issues regarding self-driving car safety and perception validation. Other areas of interest include software safety in a wide variety of industrial applications, robustness testing, and embedded system software quality. His pre-university career includes experience as a US Navy submarine officer, embedded CPU designer at Harris Semiconductor, and embedded system architect at United Technologies. He is co-founder of Edge Case Research, which provides tools and services for autonomous vehicle testing and safety validation.

Conference Abstracts

CONFERENCE DAY 0NE

Thursday 23rd May 08:30 – 17:00

Keynote

Autonomous Vehicle Safety and Perception Robustness Testing

Prof. Philip Koopman, Carnegie Mellon University and Edge Case Research

Making self-driving cars safe will require a combination of techniques. ISO 26262 and the draft SOTIF standards will help with vehicle control and trajectory stages of the autonomy pipeline. Planning might be made safe using a doer/checker architectural pattern that uses deterministic safety envelope enforcement of non-deterministic planning algorithms. Machine-learning based perception validation will be more problematic. We discuss the issue of perception edge cases, including the potentially heavy-tail distribution of object types and brittleness to slight variations in images. Our Hologram tool injects modest amounts of noise to cause perception failures, identifying brittle aspects of perception algorithms. More importantly, in practice it is able to identify context-dependent perception failures (e.g., false negatives) in unlabelled video that reveal systematic perception defects.

Biography

Prof. Philip Koopman is a faculty member at the Carnegie Mellon University ECE department, with additional affiliations with the Institute for Software Research and the Robotics Institute. He leads research on safe and secure embedded systems and teaches cost-effective embedded system design techniques.

He has over 20 years of experience with autonomous vehicle safety, dating back to the CMU Navlab team and the Automated Highway Systems (AHS) program. His most recent projects include using stress testing and run time monitoring to ensure safety for a variety of vehicle and robotic applications for the research, industry, and defense sectors. He has additional experience with automotive and industrial functional safety, including testifying as an expert in vehicle safety class action litigation and consulting to NHTSA.

He is co-founder of Edge Case Research, which provides tools and services for autonomous vehicle testing and safety validation. His pre-university career includes experience as a US Navy submarine officer, embedded CPU designer at Harris Semiconductor, and embedded system architect at United Technologies. He is a Senior Member of IEEE, a Senior Member of ACM, and a member of SAE.

http://www.ece.cmu.edu/~koopman

More than meets the AI: can systems thinking leading indicators assist proactive safety in artificially intelligent systems?

Eryn Grant

Acmena Pty Ltd

Whilst accident analysis is an accepted approach to safety management, it is reactive in nature and often requires the occurrence of adverse events to learn appropriately. For the introduction of Artificial Intelligence (AI) systems, this is problematic, as learning may well rely on ‘things going wrong’ to build on future system safety. AI represents some significant complexity; and the risk of unintended consequences and system behaviours makes it difficult to model and anticipate how safety issues might arise. This paper suggests that by integrating what is currently known about accident causation with AI technologies, it may assist AI to achieve optimal learning without adversity.

A proactive approach to safety using AI may be achieved by monitoring ‘normal performance’. However, there are few methods supported by safety science that provide specific support on the

identification of the conditions that could create accidents and safety compromising events that integrate whole systems properties without focusing on end errors. The research presented in this paper is a response to this capability gap. It presents the core tenets of accident causation that may be applied as performance indicators to proactively assess complex work systems normal performance and provide learning opportunities for AI. Implications for future research are also discussed.

Biography
Eryn Grant is a Human Factors consultant at Acmena in Brisbane, Australia working mainly within the rail sector. Eryn's PhD research within the Centre of Human Factors and Sociotechnical Systems at the University of the Sunshine Coast examined the use of a combined systems ergonomics approach to assess complex work systems across multiple safety critical domains. In doing so, she identified fifteen core tenets of accident causation that she believes are the key to proactive safety management. Her PhD thesis is currently under submission.

Machine Learning for Rail Safety Incident Classification

Ben Merange – RGB Assurance

Risk models are an essential tool to inform safety-related decisions. Data-driven models require large-scale analysis of historical data, which presents both a significant initial investment and substantial ongoing maintenance costs. In this presentation, we discuss our recent applications of machine learning for risk modelling, which involved classifying rail safety incidents based on textual descriptions. This led to a 40% time reduction in modelling effort. Safety-related textual datasets are plentiful, but often seem unwieldy. They can, however, be managed using neural networks to provide valuable insight and aid quantitative models/assessments. We briefly cover the fundamentals of neural networks, discuss a sample example, then explore our safety-related applications. In particular, we discuss the surprising relationship between model complexity and effectiveness, the ability to integrate machine learning into existing toolsets, and techniques to build trust in the technology.

Biography
Ben Merange is a Systems Engineer at RGB Assurance with a background in mechatronics, specialising in safety-critical automated systems. Ben has a keen interest in machine learning and computer vision, having worked on multiple projects that incorporate these technologies.

Keynote

Is AI in healthcare doomed, or destined for greatness?

Bijan Elahi - Medtronic

AI is coming and there is no stopping it. The momentum for AI usage in the healthcare industry is building up and likely to exceed an annual growth rate of 40% by 2021. The biggest benefits of AI are expected to be in the areas of: robotic surgery, radiological image analysis, diagnostics, and nursing. Fully automatic and autonomous medical systems are already released and being used, and nurses and doctors have started adopting the technology to reduce manual work, and to provide more accurate service and impactful interventions to patients. An analysis by Accenture says that key AI applications in healthcare can create $150 billion dollars in annual savings for the US healthcare economy by 2026.
While in other technology areas application of AI is geared primarily at performance, for healthcare applications the primary focus should be on safety. The benefits of AI are so tantalizing that they may overshadow the risks associated with the technology. Ensuring safety of AI is of paramount importance, not only to patients, but also to AI itself. Only a few catastrophic outcomes could spell doom for AI.

It is widely accepted that 100% safety is unachievable. Risk management strives to determine the level of risk and balance it against the potential benefits.

Biography
Award winning, international educator, consultant and author Bijan Elahi has worked in risk management for medical devices for over 25 years at the largest medical device companies in the world, as well as small startups. He is currently employed at Medtronic as a Technical Fellow where he serves as the corporate expert on safety risk management of medical devices. In this capacity, he offers education and consulting on risk management to all Medtronic business units, worldwide. Bijan is also a lecturer at Delft University of Technology, and Eindhoven University of Technology in the Netherlands, where he teaches risk management to doctoral students in engineering. Bijan is a frequently invited speaker at professional conferences, and is also a contributor to ISO 14971, the international standard on the application risk management to medical devices. He is the author of the book Safety Risk Management for Medical Devices.

Functional Safety Assessment of Train Management Control System

Kevin Anderson
Systra Scott Lister

The NSW Country Regional Network (CRN) was established in 2011. The legacy Train Management Control System (TMCS) was supported by a safety-related GPS Watchdog and the combination (TMACS) initially certified in 2000 covering train orders principally for the Parkes –Broken Hill Indian Pacific route and from 2011 rolled out across all relevant CRN territory.
Obviously, in the more than twenty years since TMACS was conceived and executed, technologies, techniques and assurance methods have changed markedly At each step tests of 'due diligence' have been applied to maintain system safety assurance.

Biography
Kevin Anderson has some 50 years' experience in engineering investigation and safety assurance. His two main fields are systems assurance and risk management and his process industry, railway and aviation safety experience is at the level of a Subject Matter Expert (SME).
Kevin has provided facilitation and independent certification services to high risk, safety-critical and vital industries including aviation, defence, fire and life safety, oil and gas, pipelines, process industry including, railways above and below ground, safety-critical/vital systems, and tunnel traffic and operations control systems.
He notes that there are significant safety assurance synergies in cross discipline approaches across industries and is familiar with Regulators such as Office of the National Rail Safety Regulator (ONRSR), Safety Management System (SMS) and relevant standards such as AS 61508 – Functional Safety for etc….
Kevin has delivered numerous courses in risk management and system safety assurance, co-authored the Engineers Australia Safety Case Guidelines in 2002, published an 'Introductory Text on Risk and Reliability' and written over 40 papers. He has managed over 600 engagements, more than 180 in rail, supported by extensive analysis and interpretation of incident records

Safety Assurance for Artificial Intelligence is Futile (and that's okay)

Drew Rae
School of Humanities, Languages and Social Science, Griffith University

Our current standards, regulations and guidelines for safety demand a level of safety assurance that cannot be met for an autonomous system. There are three insurmountable forces standing in our way. Firstly, there is the problem of illegibility. No regulating authority can maintain up-to-date and complete knowledge of the systems they are approving. Secondly, there is the problem of intractability. The size and inter-connectedness of systems is growing faster than the ability of analysis tools to cope with complexity. Thirdly, there is the problem of incomprehensibility. Not even the people most familiar with intelligent algorithms and autonomous systems can predict how they will behave.
So is it time to panic? Not remotely. It is time to reflect on why we conduct safety assurance in the first place, and whether it has ever been able to meet our expectations. By taking illegibility, intractability and incomprehensibility to their extremes, the prospect of Artificial Intelligence provides a safe speculative opportunity to confront the foundational challenges and limitations of Safety Engineering. In this talk, illustrated by disasters from the 20th, 21st and 22nd centuries, along with a liberal sprinkling of Wile-E-Coyote cartoons, Dr Drew Rae will examine how civilisation has already met, failed and survived all of the problems that AI brings for Safety Assurance.

Biography
Dr Drew Rae is a Senior Lecturer in the Safety Science Innovation Lab at Griffith University, where he teaches courses on safety engineering, and supervises early-career researchers. Drew's own research brings a critical cross-disciplinary approach to the examination of myths, rituals and bad habits that surround safety practice. Drew's recent publications challenge the common assumption that risk assessments and incident investigations lead to safer work. He suggests alternatives based around a better understanding of the constraints that prevent safe innovation, and the resources that support successful work. Drew presents the DisasterCast podcast and is Associate Editor for the journal - Safety Science.

Next-Generation Rail Systems using Artificial Intelligence and Machine Learning

Graham Hjort General Manager of Control Systems at 4Tel
The theme for ASSC2019 is "How can be ARTIFICALLY SAFE?", and asks the question of how artificial intelligence can be used to provide improved safety. Leveraging from the rapidly evolving autonomous car industry this paper will explore how the rail industry could improve safety, efficiency and cost-effectiveness through the application of artificial intelligence solutions to assist the train driver. The potential benefits will be explored and examined in relation to the Australia's open access railway environment and existing conventions for delivery of improved safety. If the car industry is provide any guidance on how artificial intelligence will impact the rail industry the lesson is simple. Artificial intelligence will play a role in future technologies, including those with the potential to impact safety. The question for the industry is not "will it apply?", but simply "how will it be used?". This is just one example.

Biography
Graham is responsible for the management of train control technology for the John Holland Country Regional Network. He has over 25 years' experience with signalling and train control systems. He has overseen the technology transformation for the Country Regional Network to now provide industry leading safety and efficiency. With his colleagues at 4Tel and partnership with Newcastle University he is now exploring further safety and efficiency improvements through the application of AI.

CONFERENCE DAY TWO

Friday 24th May 08:30 – 16:00

Keynote

Acceptable Risk in Human Space Flight – An Astronaut's Perspective

Pamela A. Melroy
Director, Space Technology and Policy, Nova Systems

Humans have been traveling to space for over 50 years, but until recently human space flight has been the purview of governments and extensively trained government astronauts. With recent developments in commercial space, now private companies are seeking to enable anyone with the capacity to pay for a seat to have the experience of space flight. NASA and other agencies have had a specific view about acceptable risk for decades; how will acceptable risk play out in the commercial space transportation? As a former astronaut and later a commercial space regulator, I will discuss the technical and cultural considerations around determining acceptable risk.

Biography
Pam Melroy is a retired Air Force test pilot and former NASA astronaut and Space Shuttle commander.

She was commissioned in the United States Air Force and served as a KC-10 copilot, aircraft commander, and instructor pilot. Melroy is a veteran of Operation Just Cause and Operation Desert Shield/Desert Storm, with over 200 combat and combat support hours. She went on to attend the Air Force Test Pilot School at Edwards Air Force Base, California. Upon her graduation, she was assigned to the C-17 Combined Test Force, where she served as a test pilot until her selection for the Astronaut Program. She has logged more than 6,000 hours flight time in more than 50 different aircraft.

The Ironies of Automation with AI

RGB Assurance
Simon Connelly & Dr Holger Becht
Lisanne Bainbridge presented the ironies of automation 35 years ago and identified some possible solutions. Bainbridge argued that automation may expand rather than eliminate problems with human operators. Society is becoming more and more dependent on automation, including the use of artificial intelligence. However even highly automated systems require human operators for supervision, adjustments, maintenance, etc. This talk reflects on the original ironies of automation and discusses where and how the ironies persist with AI technology, and if new ironies have emerged with the use of AI. We then discuss the implications on the role of the human operator of these automated systems.

Biographies
Simon Connelly is a Principal System Safety Consultant with RGB Assurance, with over 18 years experience in System Safety and RAM engineering, with 12 years specifically in the rail industry. He has broad experience in assurance of rail safety systems, including signalling and train control and protection systems. A particular focus of Simon's work has been in considering risks and hazards exposed at the boundaries of systems, including interfaces between systems of systems and human machine interfaces. In addition to his work in safety systems, Simon is interested in the application of assurance techniques more broadly to the systems engineering lifecycle, and how increased rigour early in a project can reduce risk to project delivery at later lifecycle phases.

Dr Becht is a career systems assurance professional with twenty years' post-graduation experience. At RGB Assurance, he has provided safety engineering, RAM, and human factors services for many projects in railway, defence, aerospace, and air traffic management spanning the Asia-Pacific region (covering Australia, Malaysia, as well as a number of jobs in India and sub-Saharan Africa). Holger possesses a PhD in software engineering of high integrity systems and a Masters in System Safety and Human Factors. Holger is a nationally- and internationally-recognized expert in the field of system safety and systems assurance and has an extensive list of publications.

Keynote

Dynamic Assurance Cases: A Pathway Towards Trusted Autonomy

Dr. Ganesh Pai Senior Research Engineer with SGT, Inc. @ NASA Ames Research Center

Over the past decade, easier accessibility to large quantities of well-labelled data, together with advances in the capabilities of graphics processing units—aided by Moore's law—has contributed to ground-breaking advances in machine learning (ML). That, in turn, has facilitated an increasing use of so-called learning-enabled components (LECs) in safety and mission-critical applications, e.g., deep neural networks used for perception in self-driving road vehicles. Urban Air Mobility is an emerging concept that promises to revolutionize air transportation through greater autonomy—i.e., by leveraging LECs—coupled, in part, with innovations in unmanned aircraft systems (UASs). The pace of innovation in ML and LEC technologies currently far outstrips that of the applicable regulatory and standardization efforts to create the bases against which it would be established that the resulting systems can be relied upon. Risk-based approaches to engendering trust, in the form of argument-based safety cases, have shown promise for the assurance and subsequent operational approval of novel systems. However, LECs pose particular challenges for certification, as does the gap between the state of the art in safety assurance, and how aviation systems are certificated in practice. Towards straddling this gap, we are developing the dynamic assurance case (DAC) concept as a model-based, multifaceted approach to the assurance of LEC-based systems. Our vision is one of a rich, expressive, and formally-founded framework, going well beyond how argument-based safety cases are currently developed. In particular, besides recording assurance rationale in a modular fashion, DACs: i) capture assurance policies and a conforming assurance architecture, ii) provide a framework for assurance quantification, and iii) also supply the means to admit design-time verification and validation (V&V) evidence, along with run-time evidence from operational monitoring. We are co-developing these assurance technologies and the supporting tool infrastructure, to be demonstrated on both aviation and maritime autonomous systems platforms.

Biography

Dr. Ganesh Pai is a Senior Research Engineer with SGT, Inc., a KBRWyle business unit, and a contractor member of the scientific staff in the Intelligent Systems Division at the National Aeronautics and Space Administration (NASA), Ames Research Center, California. His research addresses the broad area of safety and mission assurance, as applied to aerospace systems and software, while his professional practice has supported the safe engineering and operations of Unmanned Aircraft Systems (UAS). He was a principal member of the team that created the safety case for a ground-based detect and avoid solution that demonstrated the capability to conduct safe beyond visual line of sight UAS operations in civil airspace, an achievement for which he was recognized by a 2014 NASA honor award. More recently, his research focus has expanded to include dependability analysis and assurance technologies for assured autonomy to support both NASA's Airspace Operations and Safety Program, and the Quantifiable Assurance Cases for Trusted Autonomy (QUASAR) project funded by the US Defense Advanced Research Projects Agency

(DARPA), on which he is co-investigator. Dr. Pai holds a doctorate degree in Computer Engineering, and a Master of Science degree in Electrical Engineering, both from the University of Virginia. He has authored more than 40 articles spanning the broad areas of systems and software engineering, with a focus on dependability and safety. He has also served on the program committees of numerous workshops and conferences in those areas, including as co-chair of the ongoing workshop series on Assurance Cases for Software-intensive Systems (ASSURE). He is a senior member of the IEEE and the AIAA, and a member of the IEEE Computer Society, and Eta Kappa Nu, the international honor society of the IEEE.

Research Award Summary –
Application of Bayesian methods and normative decision theory to aviation safety regulatory process

Achim Washington
PhD Candidate, Aerospace School of Engineering, RMIT

The aviation sector is faced with a novel array of new airspace users including Urban Air Mobility (UAM) concepts, personal air mobility vehicles, reusable space launch vehicles, and Unmanned Aircraft Systems (UAS). Focusing on UAS, there is much effort being directed towards the development of safety regulations for this industry. National Aviation Authorities (NAA) have advocated the adoption of a risk-based approach to the development of regulations, whereby regulations are driven by the outcomes of a systematic process to assess and manage identified safety risks. The overall aim of this research is to improve regulatory outcomes under the new paradigm of risk-based regulation (not just risk-based rule-making but also risk-based compliance assessment and compliance finding), through providing a conceptual framework for the rational, transparent and systematic treatment of uncertainty in the risk assessment and regulatory decision-making processes. The research proposes the application of Bayesian methods and normative decision theory to the aviation safety regulatory process. System Safety Regulations (SSR), commonly referred to as “Part 1309” regulations, for UAS are used as a case study. It is posited that the general theoretical approach proposed can improve the objectivity, consistency, and transparency of current aviation regulatory processes. The generalised approaches presented in this research enable the adoption of risk-based rule-making for new aviation sectors and provides the theoretical basis for risk-based compliance; a paradigm shift in how aviation safety regulators approach risk-based regulation.

Biography
My name is Achim Washington and I completed my Bachelor’s degree in Aerospace engineering at the Royal Melbourne Institute of Technology (RMIT) in 2014 where I graduated with first class honours. I was the recipient of a number of awards during the course of this degree, including the Vice Chancellor’s list for Academic Excellence (2011), Boeing Achievement Award (2011), Golden Key International Honour Society (2012), Wackett Centre Research Scholarship (2013), and the QinetiQ Aero Structures Award (2014). Currently, I am a final year PhD student at RMIT University. The title for my thesis is “Risk-based Regulation of Unmanned Aircraft Systems”. This work essentially has to do with the risk and uncertainty surrounding unmanned aircraft systems and trying to develop a means to evaluate and manage this risk. During the course of my PhD candidature, I once again was able to secure a number of awards and research grants. Most recently I was shortlisted for the Avalon Young Innovator Award for 2019. Over the last two years I have also worked as a casual tutor RMIT university, where I have given a number of lectures on engineering risk management in the aviation industry. Finally, I have also undertaken research internships at RMIT university and the NASA Langley research centre, which have further helped strengthen my research efforts.

Conference Abstracts

Keynote

Standards for road vehicle automation: Are we nearly there yet?

David Ward - Horiba Mira Ltd

Since its original publication in 2011, ISO 26262 has become established as the state of the art in developing safety-related electronic systems in road vehicles. However since the standard was developed against the assumption a driver is present and in final control, this has raised questions about whether the standard should apply for developing highly automated driving functions. This presentation will examine what progress has been made in Edition 2 of the standard and related activities and give pointers for areas where further development is needed.

Biography

Dr David Ward is General Manager, Functional Safety at HORIBA MIRA Limited, a leading independent provider of automotive engineering services. Dr Ward is a recognized international expert in automotive functional safety with 25 years' experience in safety and reliability of embedded electronic systems in automotive and other industries. Within his role at HORIBA MIRA, David is responsible for training, consultancy and independent safety assessment in the functional safety standard ISO 26262 and other related standards. He is involved in automotive cybersecurity as well as technology development in functional safety, connected and autonomous vehicles, and vehicle electrification.

He is the UK Principal Expert to ISO/TC22/SC32/WG8 “Road vehicles – Functional safety” which has developed ISO 26262 and ISO/PAS 21448 “safety of the intended functionality”; a member of the ISO/SAE joint working group developing ISO/SAE 21434 “Road vehicles – Cybersecurity engineering”; as well as contributing extensively to the UK’s MISRA initiative.

David was presented with the IMechE Prestige Award for Risk Reduction in Mechanical Engineering following 20 years of work leading automotive industry efforts to develop international safety standards. These efforts began with MISRA (The Motor Industry Software Reliability Association) and have continued with the development of the international standard ISO 26262, which Dr Ward and his team at MIRA continue to influence as work considers how to extend its scope to highly automated vehicles.

David is also Visiting Professor of Functional Safety at Coventry University, UK and RAE Visiting Professor of Industrial Design at the University of Leicester, UK.

Artificial Intelligence and Safety Standard Compliance: Challenges

Dr. Tim McComb
RGB Assurance

Safety Standards for software assurance recommend against the use of AI in some contexts, some specific examples from IEC 61508 and CENELEC being on-line learning (dynamic reconfiguration) and prediction of faults from trends (fault forecasting). The basis for this is the difficulty AI presents in verifying correctness and predictability of behaviours, as well as the difficulty in comprehending the design/parameterisation of the algorithms.

In this talk, we’ll discuss different general types of AI technologies, and some of the challenges that would eventuate from attempting to demonstrate compliance for each one of these technologies. We will also explore some options for integration of AI within a software system, and how confidence in AI algorithms (e.g. through training and validation) might be used to bolster a safety argument, possibly in conjunction with other safety mechanisms. We hope this talk will trigger some discussion about the feasibility (or not!) of successfully arguing compliance against the standards for systems with AI components.

Biography
Dr Tim McComb is a Principal Software Safety Engineer at RGB Assurance. Tim has a keen interest in high-integrity software analysis and development and is currently working with clients in the aviation and rail sectors to help them achieve practical and effective assurance outcomes.

Autonomy: Is It Really As Safe As We Think? Practical Examples from the Real World

Jon Sciortino & Martin Shadbolt
Nova Systems

The rise of autonomy is gaining greater airplay in the media and in industry, as companies grapple with whether or not to deploy autonomous systems. In many cases, the decision to use autonomy is based on the need to improve safety and, if possible, enhance productivity. But are these systems as safe as we think they are? We use a number of public-domain autonomy incidents to frame the problem and examine the claims of safety improvements.

Biographies
Jon Sciortino is a systems engineering specialist with over 25 years' experience in the delivery and assessment of high-level capability assurance, system safety and test & evaluation expertise. Jon works internationally within both commercial and Defence sectors to assist clients with systems engineering, integration and requirements management, program analysis, harnessing innovation opportunities and the development of comprehensive test & evaluation methodologies across the business. Jon is highly sought after by organisations wanting to access his expertise and insight, particularly in complex projects and where a systems perspective is required. Jon is passionate about helping clients to solve the technologically challenging problems that really matter. He is also deeply committed to applying his wealth of experience to make a difference with clients. Jon holds degree qualifications in electrical engineering and business administration, is a Member of Engineers Australia and lives in Perth, Western Australia.

Martin Shadbolt is a Senior Systems Engineer with Nova Systems. He was awarded a Bachelor of Electronic Engineering for Curtin University of Technology in 1987 and is a Fellow of the Institute of Engineers Australia. Martin was an avionics engineer in the Royal Australian Air Force, responsible for the operation, maintenance, sustainment and enhancement of F/A-18 aircraft. He has over 25 years' experience in engineering, designing, integrating, testing and managing complex integrated systems in a wide variety of military and commercial environments. This presentation draws on his experience in systems and safety assurance for complex autonomous systems for aviation and mining.

CONFERENCE MIXER / DINNER VENUE

Bar Pacino, upstairs @ 175 Eagle St, Brisbane, from 7pm

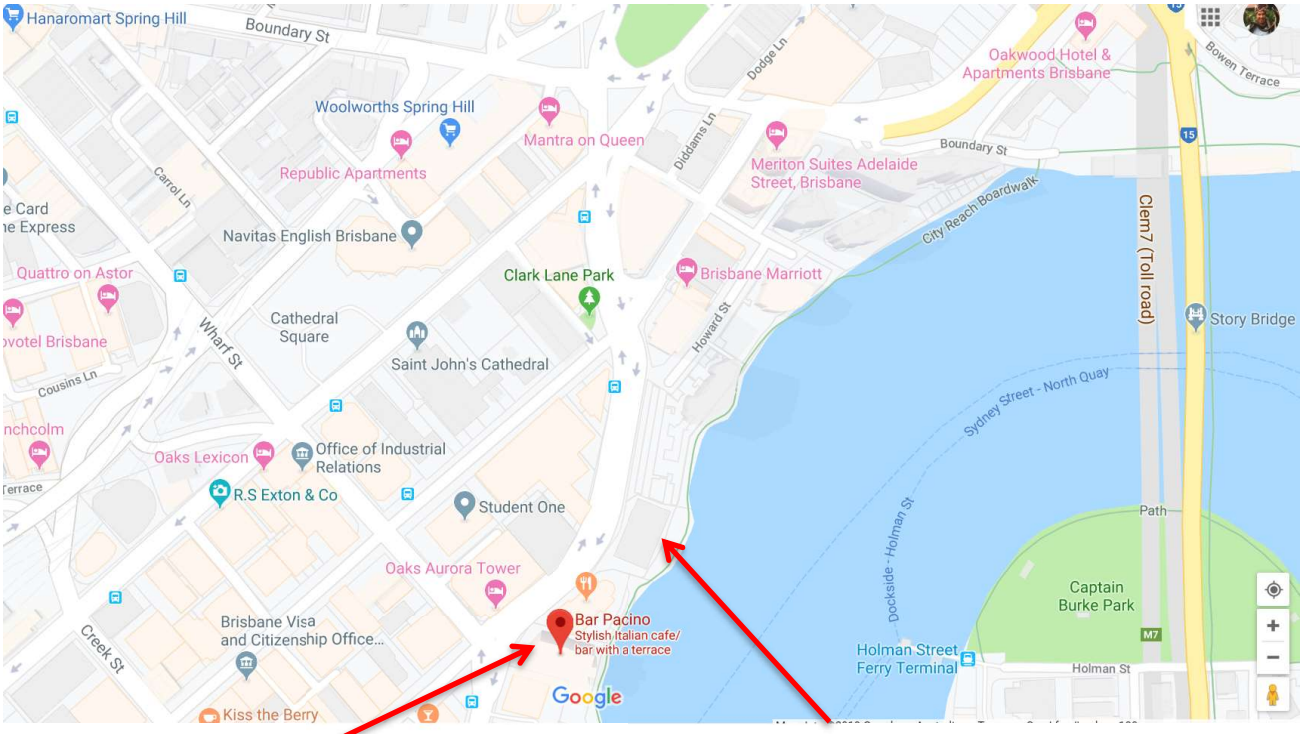

Eat here

Listen here

CONFERENCE COMMITTEE

Organising Committee

Chairperson
BJ Martin

Sponsorship
BJ Martin

Registration
George Nikandros

Communications
Dr Luke Wildman

Merchandising
Edmund Kienast

Facility & Operations
George Nikandros

Website
Dr Clive Boughton

Technical Committee

Chairperson
WGCDR Derek Reinhardt

Members
Dr Derek Reinhardt (CASG)
Paul Caseley (DSTL)
Simon Connelly (RGB)
George Nikandros (Qld Rail)
Ashley Taylor (DASA)