UKRI Trustworthy Autonomous Systems Hub

SECURITY

KEEPING OUR AUTONOMOUS SYSTEMS SECURE IN AN UNCERTAIN WORLD

As technology advances, concerns about the safety and security of our systems increase amid ever-emerging threats. Cyber security is big business. Our personal devices require constant updating against vulnerabilities. How, therefore, do we protect our autonomous systems (AS)? How can we ensure they remain secure, especially when we are not directly involved in their operation? How do we create systems can properly assess the risks they face in different environments and respond appropriately to any issues?

For any system, security is about providing assurance that it will maintain an acceptable level of service despite any issues that might arise during operation. This is challenging enough with any technology, but with autonomous systems it is even more complex. They carry out multiple actions at once - decision-making, control, coordination and navigation - plus they operate in unpredictable environments using AI technology, which makes this is even harder.

These complex issues are among those being addressed by the UKRI Trustworthy Autonomous Systems (TAS) Programme – a £33m multi-disciplinary research programme funded as part of the Strategic Priorities Fund. Security is the focus of one of the six TAS Nodes – separate research projects examining individual aspects of trust in autonomous systems – and was the topic of one of a series of multi-disciplinary TAS workshops.

Complexities and AS Security

When it comes to security, what does it mean for our autonomous systems to be safe and secure? The answer is complex, nuanced and multi-dimensional - particularly when we factor in users, environmental variabilities and social impacts.

A good place to start is with specification: strict, specific parameters that determine the way a machine operates, reacts and learns. However, this is actually one of the hardest areas for TAS researchers to address, as Hamid Asgari from Thales UK explains: “Specification is the foundation of everything. Who is going to provide the specification? We need to verify the behaviour of the system based on the specification, to see whether it meets requirements or not.”

In effect, we are looking for a very structured security framework within a very unstructured environment. This is enormously challenging. Autonomous systems need to operate in a predictable manner, but they operate largely in environments where there is much uncertainty. We have to make assumptions about the threats and situations they might encounter and the behaviour they might display. We still have a great deal to learn in this area, with much of our existing knowledge being purely theoretical and based on simulations.

There are also issues with adapting the security protocols in existing technology to include autonomous systems. Many commercial organisations, for example, use information and control panels already in existence and may only be able to cope with some of the security requirements that an autonomous system demands. Professor Weisi Guo from Cranfield University says this poses real challenges: “Many commercial systems were not necessarily designed for AS. We are trying to come up with the correct requirements for autonomous systems - designing the right security protocols and new metrics which these systems will rely on.”

The TAS Programme has been examining the various security challenges in three key areas: usage, operations, users. Each area comprises ‘onion-style’ layers, which include security threats within the autonomous system itself, the AS in operation, human ‘user’ influences and the wider world. Research is underway into how threats can run across different layers and how this impacts the way that systems adapt and learn.

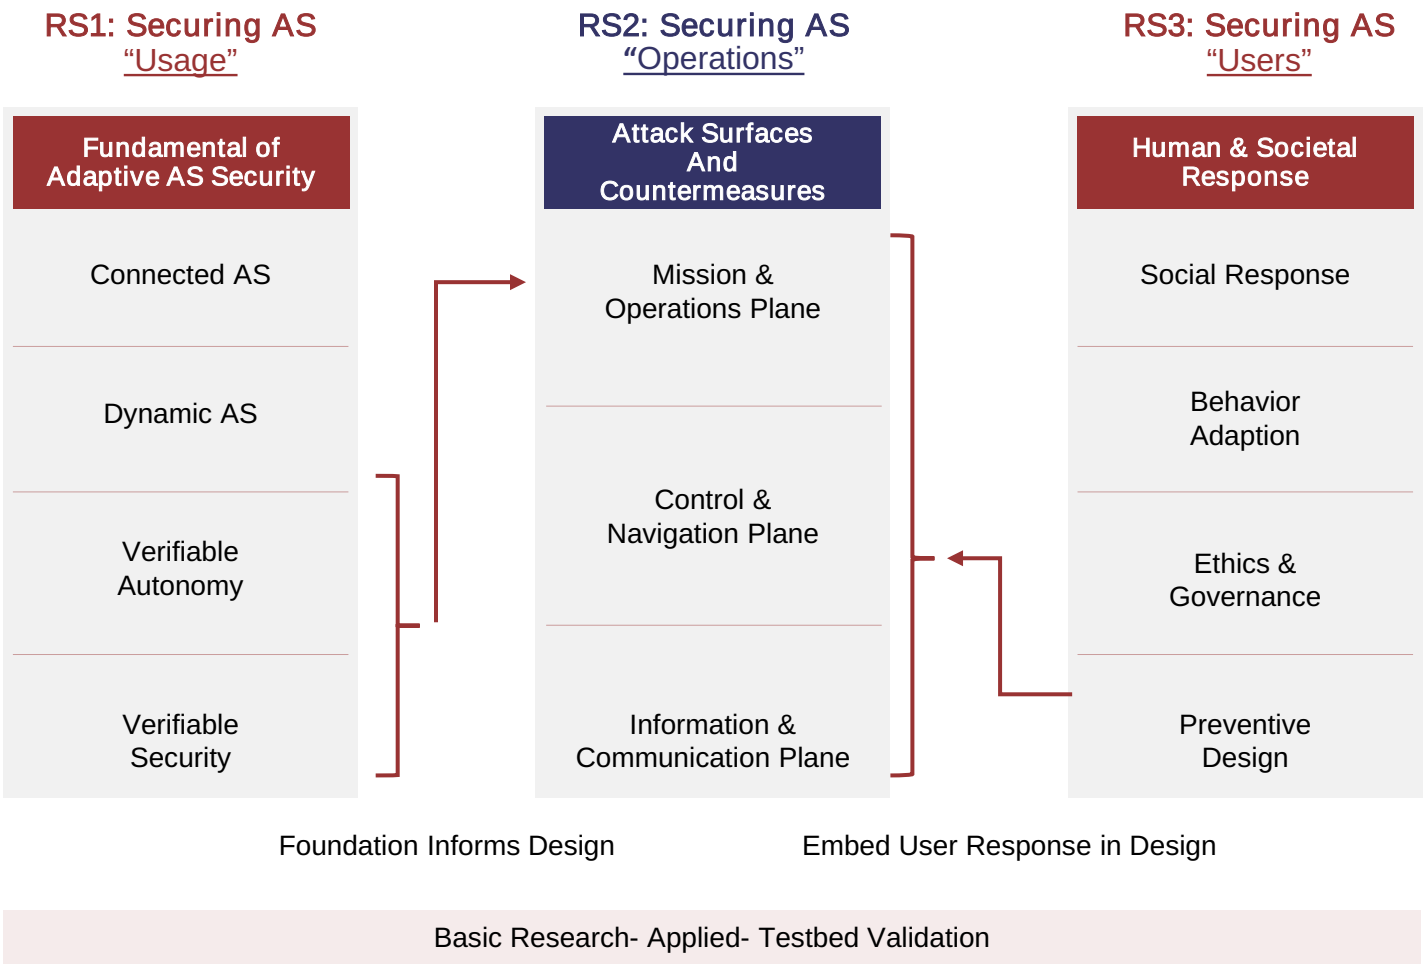

The TAS Security Node's 3 key focus areas, and the interconnected research strands (from Suri et al., 2022)

The human factor

A particular area of study for the TAS Programme is examining the human and societal impacts on our autonomous systems - the social-technical aspects. How do we secure AS from new threats if they rely on human interaction? What are the ethical, legal and social implications (ELSI)?

TAS researchers are working in partnership with stakeholder organisations to co-design methods of developing awareness of the socio-technical aspects of AS security and exploring multiple ways of measuring how humans adapt to issues that arise. Co-design is central to equipping individuals, the industry and designers with tools that they can use at different points in the decision-making process. The outcomes of this work will form a toolkit for those who design, build, own and deploy autonomous systems.

Social acceptance of autonomous systems is also important to ensure wide adoption across society and industry. As users, what do we expect from our technology? Professor Jose Such from King's College London says it is important to understand what users need and think: "These AS are not in a vacuum. They will be interacting with people, and we are interested in understanding not only how people engage with, but also perceive these systems, and how you may influence this perception."

Work is being done to gather information in this area. Crowdsourcing has been used, for example, in relation to information flows and contexts within the smart home personal assistants AS ecosystem. However, information is still sparse and there are many gaps to be filled. Ethical and social aspects of AS security are commonly a secondary consideration in commercial environments, with functional and traditional security considerations often taking precedence. The challenge is how to motivate the industry to be more open and to prioritise security.

A positive example of work currently being undertaken in this field is in the design of road infrastructure. The University of Lancaster is collaborating with Highways England on an ongoing project about autonomous vehicles (AVs) and traditional road users. They are looking at how to embed ethical, legal and social considerations into a new AS co-design process. The goal is ultimately to enable driverless cars to exchange information with the infrastructure safely and securely, to improve on-board decision-making.

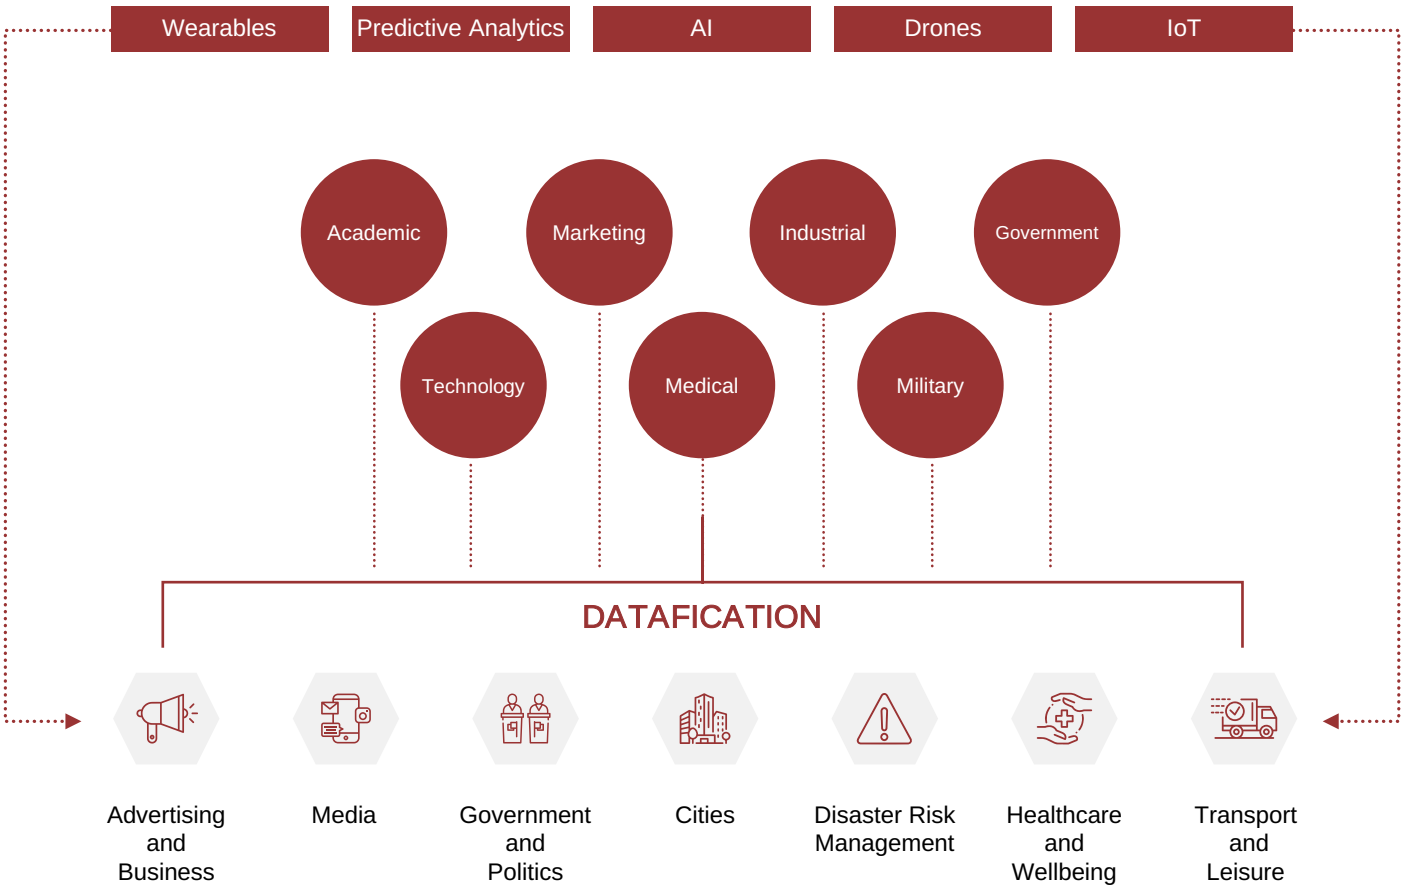

Cross sector e-society data science and AI/AS processing motivating the ELSI approach (from Buscher et al., 2018)

Incorporating arts and culture

The fields of arts and culture and security of autonomous systems seem, at first glance, to be worlds apart. In reality, however, there is a surprising amount of synergy, which is proving extremely helpful in the area of communication.

One of the challenges for the technology sector is ensuring a wide understanding of security issues among non-technical users. Using visuals and standardised, natural language to communicate this to a wider audience can be very effective. It is not enough to tell users what an autonomous system does and how they should interact with it. They need to be shown this in a way that is visually and linguistically understandable.

King's College London have been using film and arts to communicate with non-expert users. This example demonstrates how they used Nicolas Cage films to highlight technology in action and the threats that users face.

Film Title	keywords	Category
G-force	Hacking; Password cracking; Worm; Virus; iot attack	Hackers and cryptologists
Kick-Ass	Anonymity ; Tracing	
National Treasure	Steganography; Multiple-stage attack; Masquerading attack; Forging; Pseudonym; Social engineering; Invisible Ink; Hacking; Sensor attack; Biometrics; Integrity; Password guessing; Background Knowledge; Indistinguishability; Swapping; Tracing; Ottendorf cipher	
National Treasure; Book of Secrets	Playfair cipher, Password guessing; Device cloning; Hacking; Denial of service; Social engineering; Surveillance system; Codes	
Snowden	Cipher machines; Algorithms; Attacks; Malware; Surveillance; Trojan horse; Password Protection	
Wind talkers	Codes	
Con Air	Steganography; Authentication/Deception	Detectives and spies
Face/Off	Multi-factor authentication; Biometrics	
Gone in 60 seconds	Steganography; Codes; Insider attack	
Teen Titans GO! To the Movies	Pseudonyms; Multi-factor authentication; Biometrics; Masquerading attack; Social engineering	
Lord of war	Identity/Authentication; Confidentiality; Integrity; Social Engineering	Ordinary people
Spiderman: Into the Spider-Verse	Pseudonyms; Identity; Accountability	
The Humanity Bureau	Stolen Identity	
Ghost Rider	Biometrics	
	Anonymous	Allegory
The Sorcerer's Apprentice	Authentication	

An example analysis (from Vigano, 2021b) of 15 Nicolas Cage films exploring the importance of the role of his films and other arts in explaining cybersecurity to non-expert audiences

IMPACTS OF SECURITY IN THE REAL WORLD

TAS researchers have been focusing on the security of autonomous systems in two main areas: uncrewed and crewed airspace integration; and autonomous vehicles and trust among human users.

One case study they are involved with is a Heathrow Airport-led Innovate UK project, entitled Fly2Plan. The 14 partner, 15 month-long project is costing £4.6 million and aims to develop a new information-sharing model for crewed and uncrewed aerial vehicles (UAVs).

The project includes exploring secure systems using AI and other data sharing technologies for shared air and land space, air traffic management, and flight and UAV operations - including deliveries, transport, emergency response and maintenance. A secure cloud infrastructure is being adopted to replace legacy analogue systems, and more digital data and voice communication systems are being incorporated into air traffic management.

The end goal is to increase operational resilience and safety in UK airspace while reducing costs.

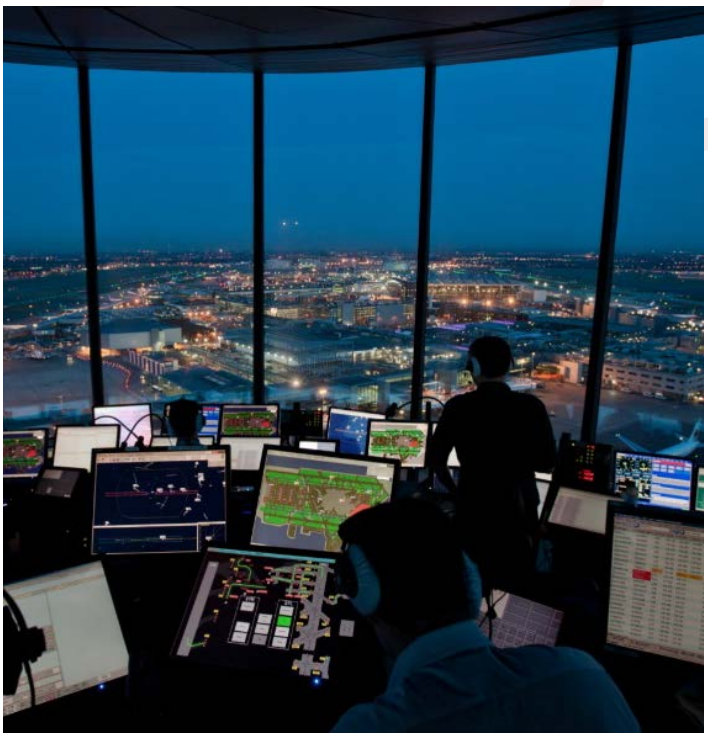

London Heathrow Air Traffic Control tower at night with, computer systems and human operators

Autonomous systems working alone

The importance of security in technology is even further amplified when humans are taken out of the loop. There is much research taking place into autonomous systems working in collaboration with their peers.

Thales are looking into various applications of multi-asset cooperative AS ‘squads’, in particular how to assess complex, real-word situations across a variety of sectors including transport, maritime, UAVs, defence and civil aviation. Researchers are looking at new approaches to identifying and analysing mission requirements in complex situations.

An example of this in marine research, where a seabed survey is taking place: an Unmanned Surface Vehicle (USV) in a shipping channel is deploying a small underwater robot (UUV), safely maintaining station and communications with a drone (UAV), before recovering the UUV.

The goal is to design robust Integrated Mission Management Systems that can supervise large squads of different autonomous systems that are working in collaboration.

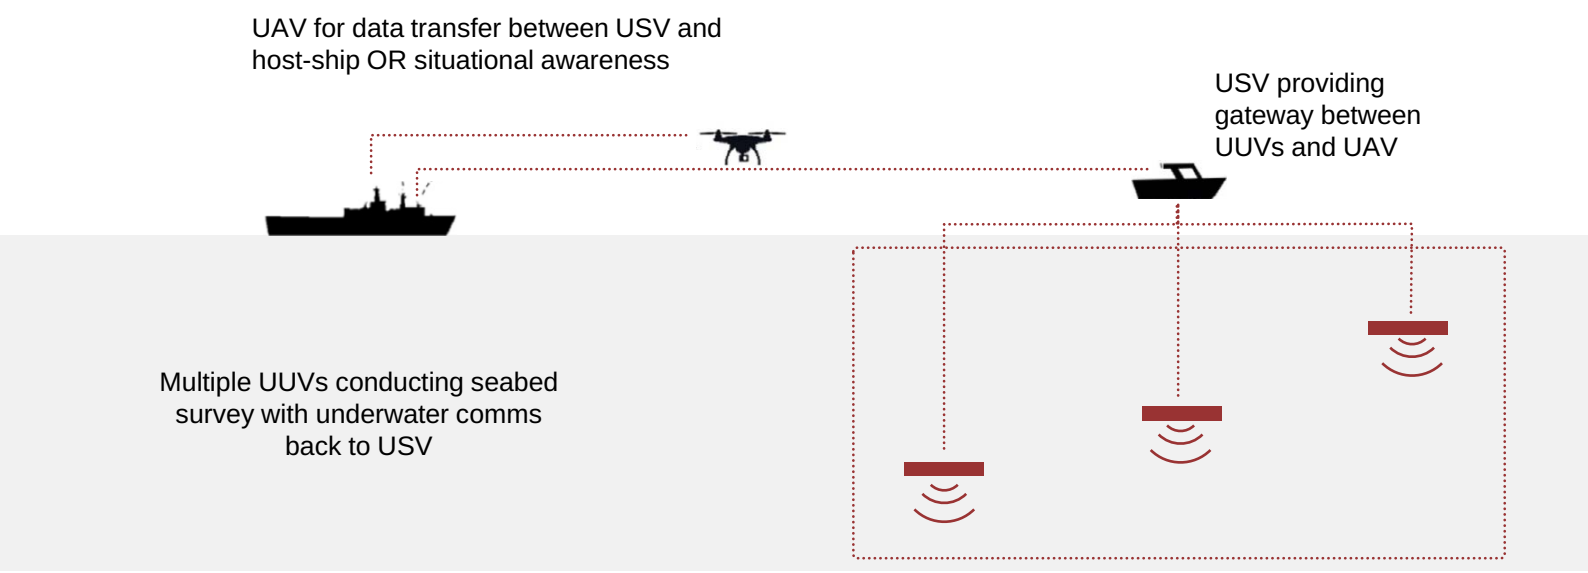

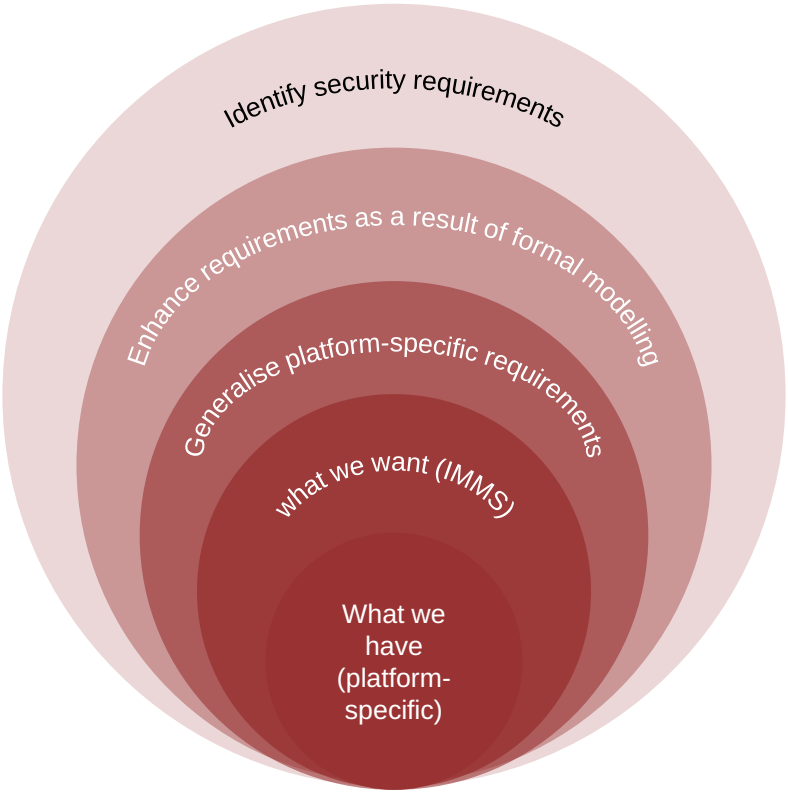

An AS squad use case scenario (top) and approach for eliciting requirements for the integrated mission management system (from Dghaym et al., 2021)

However, such situations do throw up a number of questions surrounding security. How do we monitor the autonomous systems during operations and how do we predict what they're doing if they go 'off- grid'? If threats or attacks occur, how do we maintain safety of both the systems and their environment?

TAS researchers have been developing models to better understand the security required, including scenarios around how critical machine-learning processes can get compromised. What happens if the learning process breaks down or the decision-making process for control, coordination, navigation and communication stops working? A robust security protocol for peer-to-peer machine learning is currently being worked on, able to tune itself to the type of data it is getting in that environment. Early results are very promising.

FUTURE THREATS AND AS DEVELOPMENT

With the challenges we face and the limited data available regarding security, how do we envisage the future? What is realistic for us? Can we really develop autonomous systems that can cope with any threat or situation, then react in the appropriate way? Do we see a future involving secure squads of autonomous vehicles that can operate without human intervention and deal with problems if they arise?

It is fair to say, we are on a journey.

According to some researchers, the future of AS security depends on the way we approach it. Professor Neeraj Suri from the University of Lancaster explains: “We need to change the paradigm. We intuitively think in terms of safety or security by design. Brittle security - wonderful if it holds, but terrible if it is compromised.”

The social, ethical and legal elements are also part of the journey. According to Professor Corinne May-Chahal from the University of Lancaster: “I think it is critical that we learn how to design-in ethics into AS. This will be critical for the development of AS. It's not just about AS- but about the functioning of AS in our everyday lives.”

There are questions too about future-proofing our systems. Can security ever keep up with the speed of innovation and usage? Professor Weisi Guo says this is extremely difficult: “Typically, you have a design life and then an operational life. This can be quite long. We don't know what will change in the next 10-20 years, so how do we design security for AS for the future?”

The journey continues, and we, the users of technology, are an important part of it. Often the push for autonomy comes from industry, but we may not be convinced that we need it or trust it. How do we reach a reasonable level of trust on these systems? We often have concerns around privacy and transparency. In order to move forward in a constructive way, it is important to inform and engage with the public about what we want and what we will accept. We need to be taken on the journey, and not simply observe as a bystander.

More engagement is needed with industry and the regulators too, to understand specific requirements and needs. Professor Luca Vigano from King’s College London explains: “Security has always been an add-on. Security by design is stated but not happening. People are realising that this is important, but we are still not seeing enterprise catching up with this.”

It is very clear that more work is needed. Technology is fast-moving; the ideas, visions and challenges for autonomous systems are multi-disciplinary and security is an ever-evolving field. What we do know for certain is that more teams and networks like the TAS Programme are much needed, in order for us to be able to create secure, integrated systems in the future.

REFERENCES

Suri, N., et al., 2022, UKRI Trustworthy Autonomous Systems Node on Security- TAS-S Annual Report 2020-2021. https://www.tas.ac.uk/News/node-in-security-annual-report/.

Viganò, L., 2021b. Nicolas Cage is the Center of the Cybersecurity Universe. in C Ardito, R Lanzilotti, A Malizia, A Malizia, H Petrie, A Piccinno, G Desolda & K Inkpen (eds), Human-Computer Interaction – INTERACT 2021 - 18th IFIP TC 13 International Conference, Proceedings. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 12932 LNCS, Springer-Verlag Berlin Heidelberg, pp. 14-33. https://doi.org/10.1007/978-3-030-85623-6_3

Dghaym, D., Hoang, T. S., Turnock, S., Butler, M., Downes, J., Pritchard, B., 2021. An STPA-based formal composition framework for trustworthy autonomous maritime systems. Safety Science, 136(0925-7535), https://doi.org/10.1016/j.ssci.2020.105139

Buscher, M., et al., 2018. The IsITethical? Exchange: Responsible Research and Innovation for Disaster Risk Management. Proceedings of the 15th ISCRAM Conference. ed. Boersma, K.;Tomaszewski, B. Rochester: ISCRAM, 2018. pp. 254-267.
http://idl.iscram.org/files/monikabuscher/2018/2105_MonikaBuscher_etal2018.pdf

About the Trustworthy Autonomous Systems (TAS) Hub

The TAS Hub sits at the centre of the £33M Trustworthy Autonomous Systems Programme, funded by the UKRI Strategic Priorities Fund. Its role is to coordinate and work with six research nodes to establish a collaborative platform for the UK to enable the development of socially beneficial autonomous systems that are both *trustworthy in principle and trusted in practice* by individuals, society and government. For more information, please visit the website: https://www.tas.ac.uk/.