Department: Head
Editor: Name, xxxx@email

Trustworthy Autonomous Systems through Verifiability

Mohammad Reza Mousavi
King's College London, UK; https://nms.kcl.ac.uk/mohammad.mousavi

Ana Cavalcanti
University of York, UK; https://www-users.cs.york.ac.uk/~alcc

Michael Fisher
University of Manchester, UK; https://web.cs.manchester.ac.uk/~michael

Louise Dennis
University of Manchester, UK; http://www.cs.man.ac.uk/~dennisl/

Rob Hierons
University of Sheffield, UK; https://www.sheffield.ac.uk/dcs/people/academic/rob-hierons

Bilal Kaddouh
University of Leeds, UK; https://eps.leeds.ac.uk/mechanical-engineering/staff/1003/dr-bilal-kaddouh

Effie Lai-Chong Law
University of Durham, UK; https://www.durham.ac.uk/staff/lai-chong-law/

Rob Richardson
University of Leeds, UK; https://eps.leeds.ac.uk/staff/173/robert_richardson

Jan Oliver Ringert
King's College London, UK; https://www.kcl.ac.uk/people/jan-oliver-ringert

Ivan Tyukin
Kings College London, UK;
https://www2.le.ac.uk/departments/mathematics/extranet/staff-material/staff-profiles/it37

Jim Woodcock
University of York, UK; https://www-users.cs.york.ac.uk/jim/

***Abstract*—Autonomous systems have the promise to address many of our societal challenges in a variety of areas: healthcare, climate, and economic growth, are a few examples. To realise this potential, these systems need to be trustworthy. In this paper, we describe research carried out by a UK consortium to address a central issue in establishing trustworthiness: verifiability. We explain the issues for verification that arise due to autonomy: concerns with beneficiality as well as reliability, heterogeneous artefacts and techniques, multi-disciplinary stakeholders. We also describe our vision for tackling these issues, and our progress so far.**

AUTONOMOUS SYSTEMS can make decisions, and even take actions, independent of human control or intervention. Such systems promise to improve our lives; driverless trains and robotic cleaners are examples of autonomous systems that are already among us and work well within confined environments. To make the most of their potential and gain justified public acceptance, such systems need to be trustworthy in all scenarios. We must now work to ensure developers can design trustworthy autonomous systems for dynamic and open environments, and can provide evidence of the trustworthiness of these systems.

The defining feature of autonomous systems is, unsurprisingly, autonomy: their ability to make decisions. This general description allows for a range of levels of autonomy, depending on who or what retains control. Levels recognised across sectors are often based on the PACT categorisation developed in aerospace [4] or the the subsequent SAE levels from the automotive sector [19]. Here, levels range from "1", essentially capturing human control, all the way up to "5", wherein the system itself makes all decisions and can take actions.

Although most deployed systems can be categorised at lower PACT levels, with human operators maintaining a significant level of control (and legal responsibility), the potential applications of fully autonomous systems (level 5) can be of enormous socio-economic benefit. In producing systems with higher levels of autonomy, developers are likely to start from systems for specific use-cases and operational design domains (such as, motorway driving for vehicles) and include more use-cases gradually, as the technology matures and trust is established.

However, it remains a challenge to proceed with *fully* autonomous systems in many use-cases. While this is partly due to the immaturity of technologies or the unknown added value of autonomy in some use-cases, we believe it is more fundamentally concerned with a lack of "trust" in these systems among their users.

In this article, we discuss how we might improve "trustworthiness" of autonomous systems, and how *verifiability* can be a central part of this. We describe a new collaborative research activity in the UK to tackle the complex, heterogeneous challenges of autonomous systems verification as part of their design and deployment. While we mostly focus on this initiative in the UK, there are a number of similar initiatives, e.g., in Australia,[1] Germany,[2] and the USA. [3]

Trustworthy Autonomous Systems

Trust in a system is defined as the belief or attitude that the system is helpful and beneficial in achieving the user's goal, particularly in uncertain and risky situations [20]. In traditional cyber-physical systems, "trustworthiness" often equates to reliability. We are more likely to trust some system if it works reliably. Once we move to *autonomous* systems, which can make their own decisions and take their own actions, more issues come into play. We also want to know that the system's decisions are for our benefit. This aspect, termed 'beneficiality' in [8], concerns not just what a system does, but *why* it does it. Is it working for our benefit? Is it trying to help, rather than hinder, us? What does it intend? This aspect of beneficiality might quickly become *more* important than reliability.

Example:
Recall the famous 1984 movie "The Terminator" wherein a robot appears to have few qualms about hurting humans. Our trust in such a robot is drastically reduced by its sinister intent; reliability barely comes into it. Indeed, with such sinister intent, we would prefer the robot to be unreliable! Only once we can be certain about the beneficial nature of an autonomous robot, do we want it be as reliable as possible [23].

Aside:
We have also investigated real-world examples of autonomous systems where a high level of trust is *not* justified due to their non-transparent violation of beneficiality, e.g., by polluting the environment more than legally allowed [3].

[1]Trusted Autonomous Systems: https://tasdcrc.com.au

[2]Perspicuous computing: https://www.perspicuous-computing.science/

[3]AI Safety: http://aisafety.stanford.edu/, Assured Autonomy (Computing Community Consortium): https://cra.org/ccc/visioning/visioning-activities/2019-activities/assured-autonomy, Inst. Assured Autonomy: https://iaa.jhu.edu/, and Good Systems: https://bridgingbarriers.utexas.edu/good-systems

Published by the IEEE Computer Society

Although "trust" is itself subjective, being confident about both reliability and beneficiality is important. And, as we know from decades of research and practice, confidence in software systems is related to the strength of verification we can carry out on the software. In the example above, if we can *prove* that a robot always works both beneficially and reliably then we are more likely to trust it. Although there are many other issues at play, the verifiability of these key aspects provides important input into trustworthiness.

In the UK, a £33M programme of inter-linked projects is addressing issues related to "Trustworthy Autonomous Systems". The projects comprise large "Nodes" tackling key areas, linked together by a coordination, community-building, and engagement Hub[4]. While there are many interesting and important Nodes, for example, concerned with Resilience[5] or Security[6], in this paper we focus on the work of the *Verifiability* Node[7] and how it is tackling the verification of reliability and beneficiality in autonomous systems.

Heterogeneous Verification is Essential

Autonomy is not a binary notion and may be introduced in different levels to various systems and application scenarios. However, a key aspect is how (and why) decisions are made within our systems. This can be very different across automatic systems, where decisions might be pre-coded, adaptive systems, where decisions might appear from environmental interactions and feedback, or fully autonomous systems, in which decisions may be made in line with internal aims and goals taking into account the changing context. For each of these levels of autonomy and mechanisms of its implementation, different verification techniques may be applicable.

Verification techniques range across formal and empirical, and static and dynamic. These comprise logical specification and verification [21], dynamic testing, including model-based methods [2], [22], simulation-based testing [5], runtime verification [3], [12], and stochastic methods [28]. While there are many options, it has become clear that we cannot, and should not, rely on one approach and that a *heterogeneous*, or *corroborative* collection of verification approaches is needed [13], [18], [25].

This is just what the Verifiability Node aims to provide, together with the semantic foundations to design and justify combinations of these heterogeneous concepts and techniques, and with applications that highlight the breadth of verification issues across autonomous systems.

Bringing it all Together

In the **Verifiability Node**, our vision is to carry out foundational research to enable the possibility of having a *verified autonomy store*. Autonomous systems and the components for autonomy in such a store go through rigorous and extensive verification upon submission and throughout their evolution. Having passed submission checks, components and systems are made available in a package providing the software, models for design, for compatible platforms and environments, properties, and verification evidence. The store also provides automated facilities for verification of updates to models (to include new algorithms, platforms, and environments) and components (to cater for adaptive and evolving behaviours, and for changed or extended functionality) and for incorporating new verification evidence such as deployment test results. Verification covers components and their variability and evolution, their interoperability, and system-level properties for component compositions. Properties can pertain to reactive, real-time, intentional, adaptive and uncertain aspects of platforms and environments at all levels of abstraction, from planning and decision-making all the way to hardware and physical control. In such a setting, users can have widespread access to trustworthy systems, and developers to affordable and trustworthy components. Such a store will enable reuse and reduces the prohibitively high costs for ad-hoc verification .

In order to achieve this, we need integrated coverage of everything from models of physical components to low-level control algorithms to higher-level software to services and user interactions. A single universal modelling language, verification tool or technique is not feasible or desirable, yet, we must be able to verify different aspects of these systems and how they operate to-

[4] https://www.tas.ac.uk
[5] https://resilience.tas.ac.uk
[6] https://security.tas.ac.uk
[7] https://verifiability.org

gether to enable trust. Our long-term goal is thus to develop a *unifying framework that integrates and coalesces rigorous verification techniques of autonomous systems to quickly and easily verify complex autonomous systems*.

The activities in the Node can be categorised into the three areas below.

1) **Foundational Aspects** giving the formal and practical links between the different notations required, the different semantics used, and the different tools and techniques utilised.
2) **Verification Techniques** across the different aspects, and styles, of autonomous systems and autonomous components: verifying cyber-physical systems; verifying sub-symbolic AI (for instance, deep learning); and verifying symbolic AI layers, via both static and dynamic techniques.
3) **Bridging the Gap** to real-world autonomous systems and human-robot interactions, ranging across UAVs, service robots, chatbots, human-robot teams, etc., to deal with the reality gap.

Figure 1 provides an overview of the structure of the Verifiability Node work plan. Work packages 1-3 are concerned with foundational aspects, work packages 5-7 address verification techniques, and work packages 4 and 8, as well as the two cross-cutting strands, focus on "bridging the gap".

Particularly important for collaboration across activities are common case studies in Strand 1 that allow all the different research avenues to coalesce. We have been developing common case studies across the various work packages of the Node, for example in the areas of disaster management (a firefighting drone, to be extended with connectivity and interaction mechanisms) and assistive care (a dressing robot). Figure 2 depicts an image of our firefighting drone case study. Figure 3 depicts the robotic arm of our assistive dressing case study.

In addition to carrying out fundamental research, the Node is engaging with various stakeholders in the crosscutting Strand 2 to build a community through the various organised events and the published policy- and popular science papers, all advertised on the Node website. [8]

[8]https://verifiability.org

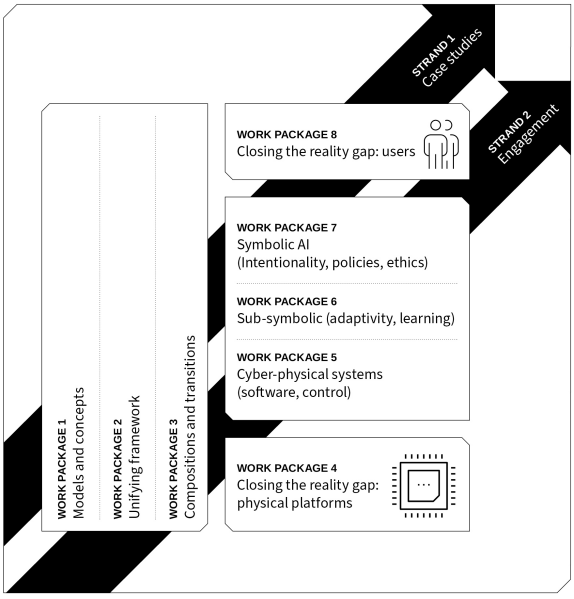

Figure 1. A Schematic View of the Verifiability Node Research Program

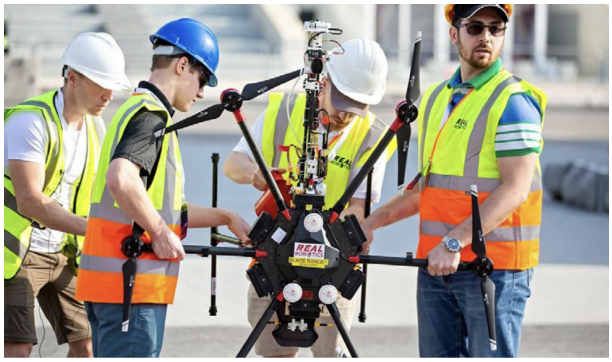

Figure 2. The firefighting drone at the Verifiability Node

Verifiability Node — Current status

The **Verifiability Node** was established in November 2020 and has already achieved several significant results. These include identifying language and notational abstractions across various domains, studying and identifying the basic building blocks of a semantic framework, and defining algorithmic abstractions, refinements, and translations across various sub-domains in the unifying framework. A detailed description of these results can be found in the Node Annual Report. [9] We highlight below a few of these results.

[9]https://verifiability.org/annual-reports/

Figure 3. The assistive dressing robot [7]

- Designed the first generation of languages to define properties for verification, operational requirements, and mappings between platform-independent and platform models [6].
- Formalised heterogeneous semantics, using our Unifying Theories of Programming (UTP) [27] and implementing this in the theorem proving framework Isabelle/UTP [17].
- Designed a compositional framework for heterogeneous specifications; we took a bottom-up approach by developing a composition of various models for the assistive dressing case study.
- Accommodated variability in learning and analysing behavioural models of autonomous systems [9]. Used AI (in particular reinforcement learning) to increase efficiency of verification strategies [24].
- Developed a runtime monitoring algorithm to search for anomalies in the state space of the system [3], as well as a general runtime monitoring framework for autonomous systems [16], [15].
- Formally verified human-level rules for autonomous systems [1] and ethical concerns in autonomous systems [10].

Verifiability Node — What Next?

The **Verifiability Node** will continue to work in all fronts above.

For example, the semantics of new notations is being fully formalised and implemented to support automatic generation and one of our next steps along this line is to mechanise relevant semantics in Isabelle/UTP. We are also applying these semantic ideas to modelling uncertainty both in case studies and more widely, in modelling digital twins. In addition, our framework for verifying 'autonomous' decision-making [11], based on verifiable agents, is being developed and expanded to handle the diversity of components.

Within the Trustworthy Autonomous Systems programme, we are collaborating with the Resilience Node on techniques for describing uncertainty in modelling autonomous systems, are collaborating with the Security Node on targetting verification to areas identified by security threat analysis, and aim to expand our collaboration further across other aspects of the programme. We will be using formal modelling and verification tools in modelling human behaviour and interaction patterns.

More widely we are keen to collaborate with teams, across academia, industry and policy, interested in working on common themes. There are existing and emerging standards such as the ANSI/UL 4600 standard for Safety for the Evaluation of Autonomous Vehicles, IEEE P7001 Standard for Transparency of Autonomous Systems [26], and IEEE P7009 Standard on Failsafe Design of Autonomous Systems [14]. The Verifiability Node has been involved in the design of the latter two standards and is currently engaging in a number of other standardisation initiatives.

Details of how to get involved can again be found at the Verifiability Node website, https://verifiability.org.

CONCLUSION

Issues around trust in technology are not new. Throughout the ages, we have had to find ways to learn to trust new tools that can benefit us. However, the issue of trustworthiness of *autonomous* systems brings new challenges. As autonomous systems essentially make their own decisions, independent of us, then our trust in these systems is not solely related to their reliability but to whether they will make the *right* decisions, even in complex and unpredictable situations. *Verifiability* has a key role not only in assessing reliability but also in establishing *beneficiality*; that systems will make decisions beneficial to us.

In this article, we described a large, multidisciplinary project focusing on the issue of trustworthiness in autonomous systems, identifying both its challenges and the results obtained so far. The vision of this "Verifiability Node" is to enhance trustworthiness through a unifying verification framework allowing for heterogeneous models, techniques and views to be analysed in tandem. This leads to holistic and wide-reaching verdicts. Our vision is that such a unified and holistic approach to verifiability will fundamentally change our approach to the verification of autonomous systems and will lead to systems that are by their construction worthy of our trust.

Our framework supports the inherent heterogeneity of autonomous systems and allows domain experts to specify their concerns in domain-specific models. The Verifiability framework takes care of connecting these models and providing holistic verification results, which are also projected back to the respective domains. Distinctive in our long-term vision is the integrated coverage of everything from models of physical components to low-level control algorithms to higher-level software to services and user interactions. To realise this vision, we closely collaborate with some of the other ongoing initiatives around the world (listed at the end of the Introduction) as well as with policy-making and standardisation bodies.

ACKNOWLEDGMENT

The work reported here is funded by the UKRI TAS Verifiability Node EP/V026801/2, Royal Academy of Engineering, and UK EPSRC.

REFERENCES

1. G. V. Alves, L. A. Dennis, and M. Fisher. A Double-Level Model Checking Approach for an Agent-Based Autonomous Vehicle and Road Junction Regulations. *Journal of Sensor and Actuator Networks*, 10(3), 2021.
2. H. L. S. Araujo, T. Hoenselaar, M. R. Mousavi, and A. V. Vinel. Connected Automated Driving: A Model-Based Approach to the Analysis of Basic Awareness Services. In *Proc. 31st International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*. IEEE, 2020.
3. S. Biewer, R. Dimitrova, M. Fries, M. Gazda, T. Heinze, H. Hermanns, and M. R. Mousavi. Conformance Relations and Hyperproperties for Doping Detection in Time and Space. *Logical Methods in Computer Science*, 18, 2022.
4. M. C. Bonner, R. M. Taylor, and C. A. Miller. Tasking Interface Manager: Affording Pilot Control of Adaptive Automation and Aiding. In *Contemporary Ergonomics 2000*. CRC Press, 2004.
5. A. Cavalcanti, A. Sampaio, A. Miyazawa, P. Ribeiro, M. C. Filho, A. Didier, W. Li, and J. Timmis. Verified Simulation for Robotics. *Science of Computer Programming*, 174, 2019.
6. A. L. C. Cavalcanti, J. Baxter, and G. Carvalho. RoboWorld: Where Can My Robot Work? In *Proc. Software Engineering and Formal Methods (SEFM)*, Lecture Notes in Computer Science. Springer, 2021.
7. G. Chance, A. Jevtić, P. Caleb-Solly, and S. Dogramadzi. A Quantitative Analysis of Dressing Dynamics for Robotic Dressing Assistance. *Frontiers in Robotics and AI*, 4, 2017.
8. R. Chatila, V. Dignum, M. Fisher, F. Giannotti, K. Morik, S. Russell, and K. Yeung. Trustworthy AI. In *Reflections on Artificial Intelligence for Humanity*. Springer, 2021.
9. C. D. N. Damasceno, M. R. Mousavi, and A. da Silva Simão. Learning by Sampling: Learning Behavioral Family Models from Software Product Lines. *Empirical Software Engineering*, 26(1), 2021.
10. L. A. Dennis, M. M. Bentzen, F. Lindner, and M. Fisher. Verifiable Machine Ethics in Changing Contexts. In *Proc. 35th Conference on Artificial Intelligence (AAAI)*. AAAI Press, 2021.
11. L. A. Dennis and M. Fisher. *Verifiable Autonomous Systems*. Cambridge University Press, 2022. (In press).
12. Y. Falcone, S. Krstic, G. Reger, and D. Traytel. A Taxonomy for Classifying Runtime Verification Tools. *Int. J. Softw. Tools Technol. Transf.*, 23(2), 2021.
13. M. Farrell, M. Luckcuck, and M. Fisher. Robotics and

Integrated Formal Methods: Necessity Meets Opportunity. In *Proc. 14th IFM Conference*, volume 11023 of *Lecture Notes in Computer Science*. Springer, 2018.

14. M. Farrell, M. Luckcuck, L. Pullum, M. Fisher, A. Hessami, D. Gal, Z. Murahwi, and K. Wallace. Evolution of the ieee p7009 standard: Towards fail-safe design of autonomous systems. In *Proc. IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW)*, pages 401–406, 2021.
15. A. Ferrando, L. A. Dennis, R. C. Cardoso, M. Fisher, D. Ancona, and V. Mascardi. Toward a Holistic Approach to Verification and Validation of Autonomous Cognitive Systems. *ACM Transactions on Software Engineering and Methodology*, 30(4), 2021.
16. M. Fisher, A. Ferrando, and R. C. Cardoso. Increasing Confidence in Autonomous Systems. In *Proc. 5th ACM International Workshop on Verification and mOnitoring at Runtime EXecution (VORTEX)*. ACM, 2021.
17. S. Foster, F. Zeyda, and J. Woodcock. Unifying Heterogeneous State-Spaces with Lenses. In *Proc. 13th International Colloquium on Theoretical Aspects of Computing (ICTAC)*, volume 9965 of *Lecture Notes in Computer Science*, 2016.
18. M. Gleirscher, S. Foster, and J. Woodcock. New Opportunities for Integrated Formal Methods. *ACM Computing Surveys*, 52(6), 2020.
19. S. International. Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles. Technical Report J3016_202104, 2021.
20. J. D. Lee and K. A. See. Trust in Automation: Designing for Appropriate Reliance. *Human Factors*, 46, 2004.
21. M. Luckcuck, M. Farrell, L. A. Dennis, C. Dixon, and M. Fisher. Formal Specification and Verification of Autonomous Robotic Systems: A Survey. *ACM Computing Surveys*, 52(5), 2019.
22. A. Miyazawa, P. Ribeiro, W. Li, A. Cavalcanti, J. Timmis, and J. Woodcock. RoboChart: Modelling and Verification of the Functional Behaviour of Robotic Applications. *Software & Systems Modeling*, 18(5), 2109.
23. M. Salem, G. Lakatos, F. Amirabdollahian, and K. Dautenhahn. Would You Trust a (Faulty) Robot?: Effects of Error, Task Type and Personality on Human-Robot Cooperation and Trust. In *Proc. 10th HRI Conference*. ACM, 2015.
24. U. C. Türker, R. M. Hierons, M. R. Mousavi, and I. Y. Tyukin. Efficient State Synchronisation in Model-based Testing through Reinforcement Learning. In *Proc. 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2021.
25. M. P. Webster, D. G. Western, D. Araiza-Illan, C. Dixon, K. Eder, M. Fisher, and A. G. Pipe. A Corroborative Approach to Verification and Validation of Human-Robot Teams. *International Journal of Robotics Research*, 39(1), 2020.
26. A. F. T. Winfield, S. Booth, L. A. Dennis, T. Egawa, H. Hastie, N. Jacobs, R. Muttram, J. I. Olszewska, F. Rajabiyazdi, A. Theodorou, M. Underwood, R. H. Wortham, and E. Watson. IEEE P7001: A New Standard on Transparency. *Frontiers in Robotics and AI, section Ethics in Robotics and Artificial Intelligence*, 2021. To Appear.
27. J. Woodcock and A. Cavalcanti. A Tutorial Introduction to Designs in Unifying Theories of Programming. In *Proc. 4th International Conference on Integrated Formal Methods (IFM)*, volume 2999 of *Lecture Notes in Computer Science*. Springer, 2004.
28. J. M. Zhang, M. Harman, L. Ma, and Y. Liu. Machine Learning Testing: Survey, Landscapes and Horizons. *IEEE Transactions on Software Engineering*, 2020.

Mohammad Reza Mousavi is Professor of Software Engineering at King's College London. His main research area is in model-based testing, particularly applied to software product lines and cyber-physical systems. He has been leading several research initiatives and industrial collaboration projects on healthcare and automotive systems their validation, verification, and certification.

Ana Cavalcanti is Royal Academy of Engineering *Chair in Emerging Technologies* at the University of York. She works on Software Engineering for Robotics: modelling, validation, simulation, testing, and verification. She currently leads the RoboStar centre of excellence in this area, and is Chair of the board of the Formal Methods Europe Association.

Michael Fisher is Royal Academy of Engineering *Chair in Emerging Technologies* at the University of Manchester. His research concerns verification, responsibility, trustworthiness, and safety of autonomous robotic systems, and he co-chairs the IEEE Technical Committee on the *Verification of Autonomous Systems*: https://www.ieee-ras.org/verification-of-autonomous-systems.

Louise Dennis is a Senior Lecturer at the University of Manchester where she leads the Autonomy and Verification group. Her expertise is in rational agent programming languages and architectures for autonomous systems, particularly ethical machine reasoning, explainability and creating verifiable systems.

Rob Hierons is a Professor of Testing at The University of Sheffield. His research largely concerns the automated generation of efficient, systematic test suites on the basis of program code, models or specifications. He is joint Editor of the Journal of Software Testing, Verification, and Reliability (STVR) and is a member of the editorial board of The Computer Journal.

Bilal Kaddouh is is a Lecturer in Aerial Robotics at the University of Leeds. He primarily researches robotics and unmanned systems design, control and multi-robot mission management. He works on the development of UAS technologies for infrastructure inspection and maintenance, BVLOS operations, precision agriculture and atmospheric studies.

Effie Lai-Chong Law is a Professor of Computer Science, specialising in Human-Computer Interaction (HCI). Her long-term research focus is Usability and User Experience (UX) methodologies. Her recent research areas are: Multisensory emotion recognition; Conversational AI (chatbots); Mixed Reality. She has authored more than 200 peer-reviewed papers and played a leading role in a number of research projects on technology-enhanced Learning, health and welling.

Rob Richardson is a Professor of Robotics at the University of Leeds. His research interests cover a broad range of applied robotics including robotics for civil infrastructure inspection and repair, making smart robots, and robotics for 3D printing applications. He has key roles in many large-scale research projects including, Pipebots, Trustworthy Autonomous Systems and AMPI (advanced machinery and productivity institute). His robotic platforms operate in the air, on the ground, in the water and underground.

Jan Oliver Ringert is a lecturer in Software Engineering at King's College London. His main research area is model-based software engineering with a focus on applying formal analyses for verification and synthesis of reactive systems and software system evolution.

Ivan Tyukin, is a Professor of Mathematical Data Science and Modelling at King's College London. Prior to that he held positions at the University of Leicester, UK and RIKEN Brain Science Institute, Japan. His current research interests are in mathematical foundations of Artificial Intelligence (AI) and learning systems, mathematical modelling, adaptive systems, inverse problems with nonconvex and nonlinear parameterization, data analytics, and computer vision. Ivan Tyukin is a UKRI Turing AI Fellow pursuing a research programme to develop Adaptive, Robust, Resilient, Certifiable, and Trustworthy AI systems.

Jim Woodcock is a Professor of Software Engineering at the University of York, Professor and Distinguished Researcher at Aarhus University, and Professor at Southwest University, China. He is a Fellow of the UK Royal Academy of Engineering and editor in chief of the ACM journal Formal Aspects of Computing. He works on the theory and practice of formal methods for software engineering.