Verification of autonomous systems [TC spotlight]

Araiza-Illan, Dejanira; Fisher, Michael; Leahy, Kevin; Olszewska, Joanna Isabelle; Redfield, Signe

Published in:
IEEE Robotics and Automation Magazine

DOI:
10.1109/MRA.2022.3143966

Published: 22/03/2022

Document Version
Peer reviewed version

Araiza-Illan, D., Fisher, M., Leahy, K., Olszewska, J. I., & Redfield, S. (2022). Verification of autonomous systems [TC spotlight]. *IEEE Robotics and Automation Magazine*, *29*(1), 99-101. https://doi.org/10.1109/MRA.2022.3143966

Verification of Autonomous Systems

Dejanira Araiza-Illan[1]
Singapore Innovation Hub, Supply Chain Strategy, Innovation & Deployment, Johnson & Johnson, Singapore
Michael Fisher[2]
Department of Computer Science, University of Manchester, UK
Kevin Leahy[3]
MIT Lincoln Laboratory, Lexington MA 02421, USA
Joanna Isabelle Olszewska[4]
School of Computing & Engineering, University of West of Scotland, UK
Signe Redfield[5]
Naval Research Laboratory, Washington, DC, 20375, USA

I. Overview

The robotics and autonomous systems communities have seen a significant and rapid increase in both the development of robots and vehicles for commercial use and in using these systems across a wide range of novel applications. As these robots, vehicles, software and even embedded devices move towards much greater autonomy, techniques for verification, providing much higher confidence than usual, are becoming required. However, the analysis and evaluation processes used for traditional systems must be significantly enhanced to provide increased confidence in this next wave of *autonomous* systems. The need for well understood and effective verification techniques will become even vital, as we move to commercial applications that rely on complex Artificial Intelligence (AI) technologies, and the utilisation of these systems in safety-critical scenarios.

There are a growing number of research developments concerning the verification of complex systems that can all impact upon this problem. These are clearly of relevance for designing, constructing, and deploying autonomous systems, but also have importance to Psychology (e.g., social robotics), Philosophy (e.g., machine ethics), and Law (e.g., certification). Furthermore, constructing autonomous systems without strong behavioral guarantees can lead to serious outcomes and may consequently hold back the widespread adoption of these systems.

As the research is currently fragmented and often not well publicized, this Technical Committee$*$ aims to coalesce this activity, drive the research agenda forward, and instill the necessity for verification firmly within industry, government, standards and the public.

II. Background

Following a number of workshops and collaborations, the *Verification of Autonomous Systems* Technical Committee (TC-VAS) was established within the IEEE Robotics and Automation Society in 2019. It has now developed into a large community, with a mailing list of over 490 unique individuals[1], monthly webinars, workshops and input to standards activities.

$*$ www.ieee-ras.org/verification-of-autonomous-systems

[1] dejanira.araiza.i@gmail.com. ORCID: 0000-0003-4456-7554. Contributions to this article were made by Dejanira Araiza Illan in her personal capacity. The opinions expressed in this article are the author's own and do not reflect the view of Johnson & Johnson.
[2] michael.fisher@manchester.ac.uk. ORCID: 0000-0002-0875-3862
[3] kevin.leahy@ll.mit.edu.ORCID: 0000-0001-5894-7190. Kevin Leahy is currently an MIT Lincoln Laboratory employee. No Laboratory funding or resources were used to produce the result/findings reported in this publication.
[4] joanna.olszewska@ieee.org. ORCID: 0000-0001-5945-7505
[5] signe.redfield@nrl.navy.mil. ORCID: 0000-0002-5428-6054. Signe Redfield is currently an NRL employee. No Laboratory funding or resources were used to produce the result/findings reported in this publication.

[1] 493 unique people; 529 emails — checked on 2021/10/23

III. What is 'Verification'?

The term *verification* encompasses a range of techniques meant to assess, with varying degrees of strength, whether a particular system matches its required behaviour. Essentially, it consists in providing evidence that some system, **S**, matches its requirements, **R**. Since the borders between verification and validation can be blurred, this Technical Committee takes an inclusive approach to such classifications. The problem of defining requirements that correctly and completely describe the actual desired behavior is likewise included.

Verification techniques are often categorised into those that have a basis in mathematical logic/proof and those that rely on more empirical approaches. Within each of these classes, there remain many options. For example, within the former (termed *formal* verification), one might employ [8]:

- **Proof** — where a formal proof is carried out to establish that **R**, encoded as a logical formula, follows from the logical description of the behaviour of **S**;
- **Model Checking** — where **R** is exhaustively assessed against a representation of all possible execution paths of **S**;
- **Runtime Verification** — where **R** is assessed against the system **S** *as it is executing* [10].

All these above-mentioned options also have probabilistic versions, many of which have become popular in recent years [2], [7].

There are also a range of *empirical* verification approaches, such as:

- **Software Testing** — where **R** is checked on a subset of the possible executions of **S**;
- **Simulation-based Testing** — as above but where a *simulation* of the real environment is used for environmental modelling and interactions [1];
- **In-situ Testing** — where **R** is assessed against the actual working of the system in a real-world environment [3].

Again, there are many variations and options here ranging from the 'certain' (in the case of formal proof) to the stochastic. The Technical Committee aims to encompass work across all these areas as well as to link and support aspects such as transparency [11] and modularity [9], [6].

IV. What is 'Autonomy'?

Essentially, **autonomy** is the ability (and often requirement) of a system to make its own decisions and take its own actions. This Technical Committee takes an inclusive view on system autonomy, covering *full autonomy*, where decision-making and action is fully within the system's software (and so assessment of why decisions are made becomes crucial); *adaptive systems*, where decision-making and action are driven by (often continuous) interactions with the environment; and *automated systems*, where decision-making and action are pre-scripted, etc. Which form of decision-making is utilised will also have a strong impact on the effectiveness of any of the verification techniques which can be applied to.

V. Why is this important?

As the range of systems that are expected to act on their own expands, so the need for verification of these autonomous systems becomes more important. When there is a 'human in the loop', i.e. a human providing oversight and control of a system, the key decisions about the system can be delegated to that human, leaving the system analysis to address issues such as reliability. However, once there is a need for the system to make key decisions, much more evidence and confidence in this type of systems will be required. Developing the ability to establish and provide this evidence is then essential not only for Engineers, but for all stake-holders such as Regulators, the Public, and Governments, and so for this Technical Committee.

If the abilities of systems and the environments in which they are to work remain constrained, then realistic boundaries for system behaviours can be provided. However, once autonomous systems are deployed in hard to predict or unknown environments and we expect them to make key, and sometimes (safety, mission, security) critical decisions, then a much stronger analysis is required. In addition, what requirements might be assessed depend crucially on what is known of the system and its environment. Traditionally, it has been assumed that we can assess (before deployment) all potential issues/concerns and mitigate against these, which might be the case in highly controlled, closed environments. However, with autonomous systems increasingly being used in open, uncontrolled environments and with internal, software behaviour able to change in various ways, our ability to predict "everything that might go wrong" is severely limited. Furthermore, stochastic models of complex, unknown environments can never be complete and may have hard-to-predict errors. Therefore, this Technical Committee is concerned with the development of tools and techniques to verify autonomous systems even in such unconstrained and unstructured environments.

VI. Technical Committee Organisation

Leadership of the Technical Committee is provided by five people, as follows.

A. Technical Committee Co-Chairs:

- Dejanira Araiza-Illan
- Michael Fisher
- Signe Redfield

B. Technical Committee Junior Co-Chairs:

- Kevin Leahy
- Joanna Isabelle Olszewska

VII. Technical Committee Activities

The Technical Committee engages in a range of activities:

- Monthly hour-long webinars, each comprising 2×20 minute talks, with questions (and answers) and an approximate average number of attendees of 55;
- Annual meetings to discuss existing activities, collate feedback, and define new activities;
- Sponsorship of workshops, such as at IROS or ICRA; and
- Input into IEEE standards, such as IEEE 2817 [5] and IEEE P7009 [4].

As an example, the P7009 Standards Working Group, which studies the "Fail-safe Design of Autonomous Systems", asked this Technical Committee for input concerning autonomous systems verification. Thus, a subgroup among TC-VAS was created, met regularly, and subsequently provided the P7009 team with a summary document to be included in their standard.

VIII. Technical Committee Future Plans

The Technical Committee also has plans for a range of future activities:

- A **Roadmap** highlighting research challenges and developments required over the medium to long term;
- A range of **Educational Resources** on the "verification of autonomous systems" topic;
- A **Catalog and Repository of Tools** available for some aspect of the verification of autonomous systems;
- Continued sponsorship of workshops at CASE, IROS, or ICRA, but also a stand-alone event dedicated to the verification of autonomous systems.

IX. Closing Remarks

Individuals can become involved with this Technical Committee by visiting our home page at `https://www.ieee-ras.org/verification-of-autonomous-systems` and selecting the **Join Us** menu option. This will bring up a form that adds new members to our mailing list.

References

[1] H. Alghodhaifi and S. Lakshmanan. Autonomous Vehicle Evaluation: A Comprehensive Survey on Modeling and Simulation Approaches. *IEEE Access*, 9:151531–151566, 2021.

[2] C. Baier and J.-P. Katoen. *Principles of Model Checking (Representation and Mind Series)*. The MIT Press, 2008.

[3] A. Bertolino, P. Braione, G. D. Angelis, L. Gazzola, F. Kifetew, L. Mariani, M. Orrù, M. Pezzè, R. Pietrantuono, S. Russo, and P. Tonella. A Survey of Field-Based Testing Techniques. *ACM Computing Surveys*, 54(5), may 2021.

[4] 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. 32nd International Symposium on Software Reliability Engineering 2021 [Industry Track]*, 2021.

[5] IEEE. P2817: Guide for Verification of Autonomous Systems. `https://www.ieee-ras.org/industry-government/standards/active-projects/verification-of-autonomous-systems`.

[6] ISO. 22166: standard on Modular Robotics. `https://committee.iso.org/home/tc299`.

[7] A. Lavaei, S. Soudjani, A. Abate, and M. Zamani. Automated Verification and Synthesis of Stochastic Hybrid Systems: A Survey, 2021.

[8] 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):100:1–100:41, 2019.

[9] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng. ROS: An Open-source Robot Operating System. In *Proc. ICRA Workshop on Open Source Software*, 2009.

[10] C. Sánchez, G. Schneider, W. Ahrendt, E. Bartocci, D. Bianculli, C. Colombo, Y. Falcone, A. Francalanza, S. Krstic, J. M. Lourenço, D. Nickovic, G. J. Pace, J. Rufino, J. Signoles, D. Traytel, and A. Weiss. A Survey of Challenges for Runtime Verification from Advanced Application Domains (beyond software). *Formal Methods Syst. Des.*, 54(3):279–335, 2019.

[11] 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. `https://doi.org/10.3389/frobt.2021.665729`.