Verifying Autonomous Robots: Challenges and Reflections

Clare Dixon
Department of Computer Science, The University of Manchester, UK
https://www.research.manchester.ac.uk/portal/clare.dixon.html
clare.dixon@manchester.ac.uk

Abstract

Autonomous robots such as robot assistants, healthcare robots, industrial robots, autonomous vehicles etc. are being developed to carry out a range of tasks in different environments. The robots need to be able to act autonomously, choosing between a range of activities. They may be operating close to or in collaboration with humans, or in environments hazardous to humans where the robot is hard to reach if it malfunctions. We need to ensure that such robots are reliable, safe and trustworthy. In this talk I will discuss experiences from several projects in developing and applying verification techniques to autonomous robotic systems. In particular we consider: a robot assistant in a domestic house, a robot co-worker for a cooperative manufacturing task, multiple robot systems and robots operating in hazardous environments.

2012 ACM Subject Classification Computer systems organization → Dependable and fault-tolerant systems and networks; Software and its engineering → Software verification and validation; Theory of computation → Logic

Keywords and phrases Verification, Autonomous Robots

Digital Object Identifier 10.4230/LIPIcs.TIME.2020.1

Category Invited Talk

Funding *Clare Dixon*: This work was funded by the Engineering and Physical Sciences Research Council (EPSRC) under the grants Trustworthy Robot Systems (EP/K006193/1) and Science of Sensor Systems Software (S4 EP/N007565/1) and by the UK Industrial Strategy Challenge Fund (ISCF), delivered by UKRI and managed by EPSRC under the grants Future AI and Robotics Hub for Space (FAIR-SPACE EP/R026092/1) and Robotics and Artificial Intelligence for Nuclear (RAIN EP/R026084/1).

Acknowledgements The work discussed in this document was carried out collaboratively with researchers on the following funded research projects: Trustworthy Robot Systems[1]; Science of Sensor Systems Software[2]; Future AI and Robotics Hub for Space[3]; and Robotics and Artificial Intelligence for Nuclear[4].

1 Formal Verification of Autonomous Robots

Autonomous robots are being developed for many purposes across society. These may be autonomous cars or pods, home robot assistants, warehouse robots, museum guides, delivery robots, companion robots, agricultural robots etc. Whilst these robots have the potential to be of great use to society, improving our lives, we need to make sure that they are reliable, safe, robust and trustworthy. We discuss work from several projects about experiences verifying autonomous robots.

[1] www.robosafe.org
[2] www.dcs.gla.ac.uk/research/S4/
[3] www.fairspacehub.org
[4] rainhub.org.uk

Formal verification is a mathematical analysis of all behaviours using logics and tools such as theorem provers or model checkers. Formal verification is often applied to an abstraction of the real system to obtain a discrete and finite system that is not too large. There has been recent interest in applying formal verification to autonomous robot systems, see for example [14] for a overview.

Here we focus on temporal verification using automatic tools and techniques such as model checking and deduction (see for example [8]) that do not require user interaction. Model checking [4, 12, 3] is a fully automatic, algorithmic technique for verifying the temporal properties of systems. Input to the model checker is a model of the system and a property to be checked on that model specified in temporal logic. For temporal deduction both the system (S) and the property (P) are specified in logic and a calculus is applied to check that P is a logical consequence of S.

2 Domestic Robot Assistants

Robots can be used in healthcare or elderly care environments enabling people who need assistance to continue living in their own homes. Such robots can help carrying things, provide reminders to drink water or take medicine, inform the user when something in the house needs attention such as the doorbell is ringing or the bath is overflowing, or provide alerts to care givers if the person is not responding.

We considered a personal robot assistant located in the robot house at the University of Hertfordshire [16]. The house has sensors that provide information to the robot about the kettle or the television being on, the fridge door being left open, someone sitting on the sofa etc. The robot is controlled by a set of "if-then" rules (with priorities) where the "if" part checks whether a condition on the sensor data or internal Boolean flags is satisfied and the "then" part contains actions for the robot to execute and flags to be set. We modelled the decision making rules using a model checker, checking properties relating to the expected execution of the rules. We considered two different approaches: one modelling the system using Brahms [17] a human agent modelling language translated into the model checker Spin [20, 19]; and the other via direct modelling in the NuSMV model checker [5, 9]. Experiments with real robots were also carried out considering whether mistakes by the robot affected people's trust in them [15].

3 Collaborative Manufacture

A second use case for robot assistants is in collaborative manufacture. We considered a robot co-worker with a scenario of a handover task for collaboratively constructing a table [21]. We considered this scenario using three types of verification: formal verification using model checking, simulation based testing [1] and user experiments with the robot.

We modelled the scenario using probabilistic timed automata and carried out verification using the PRISM model checker [11]. Simulation based testing involved developing a simulation of the system under consideration and generating and executing tests systematically that cover interesting parts of the state space. Also experiments with the robot and users were carried out. Issues found using one form of verification were used to improve the models and therefore the verification for the others. Some properties were more amenable to verification with some of these methods rather than others.

4 Swarm Robots and Wireless Sensor Networks

A robot swarm is a collection of simple, often identical, robots working together to carry out some task. Each robot has a small set of behaviours, is typically able to interact with other nearby robots and with its environment usually without a central controller. Robot swarms are often thought to be fault tolerant in that it may be possible to design a swarm so that the failure of some of the robots will not endanger the success of the overall mission. Wireless sensor networks are similar being a collection of simple, often identical, sensors with no centralised control.

We have modelled and verified properties of swarm robots relating to coherence [6] and energy optimisation for foraging robots [13] and relating to synchronisation properties for wireless sensor networks [10, 18] enabling us to detect cases and parameter settings where the required property does not (or is unlikely to) hold.

5 Robots in Dangerous Environments

In certain environments it may be preferable or we may need to use robots to carry out tasks because they are dangerous or hard to access for example underwater, space or nuclear clear up. Ensuring such robots are reliable and robust is particularly important as we may not be able to access them to reset or repair them if they go wrong. Here we advocate using different types of verification for different components [7, 2].

6 Conclusions

We have discussed experiences in verifying autonomous robots from robot assistants in the home and for collaborative manufacture, to swarm robots and robots in hazardous environments. We believe that the decision making aspects of the robot should be separated from other components to allow verification and explainability of the decisions made. We advocate using a range of verification techniques including formal verification, simulation based testing and end user experiments to improve the reliability of such systems. Different components or different types of property may require different types of verification. We believe better verification will lead to improvements not only in their safety and reliability but may also be used as evidence of this to regulators and improve trust in them by the public. Many challenges remain including how to verify systems that learn, designing systems in a modular way amenable to verification, modelling the environment, and how to provide evidence to certify autonomous robotics.

References

1 D. Araiza-Illan, D. Western, A. G. Pipe, and K. Eder. Coverage-driven verification - an approach to verify code for robots that directly interact with humans. In N. Piterman, editor, *Hardware and Software: Verification and Testing - 11th International Haifa Verification Conference, HVC 2015, Haifa, Israel, November 17-19, 2015, Proceedings*, volume 9434 of *Lecture Notes in Computer Science*, pages 69–84. Springer, 2015.

2 R. C. Cardoso, M. Farrell, M. Luckcuck, A. Ferrando, and M. Fisher. Heterogeneous verification of an autonomous curiosity rover. In *Proceedings of the NASA Formal Methods Symposium*. Springer, 2020.

3 A. Cimatti, E. Clarke, E. Giunchiglia, F. Giunchiglia, M. Pistore, M. Roveri, R. Sebastiani, and A. Tacchella. NuSMV Version 2: An OpenSource Tool for Symbolic Model Checking. In

Proc. International Conference on Computer-Aided Verification (CAV 2002), volume 2404 of *LNCS*, Copenhagen, Denmark, July 2002. Springer.

4 E. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, 2000.

5 C. Dixon, M. Webster, J. Saunders, M. Fisher, and K. Dautenhahn. "The Fridge Door is Open"-Temporal Verification of a Robotic Assistant's Behaviours. In M. Mistry, A. Leonardis, M. Witkowski, and C. Melhuish, editors, *Advances in Autonomous Robotics Systems*, volume 8717 of *Lecture Notes in Computer Science*, pages 97–108. Springer, 2014.

6 C. Dixon, A.F.T. Winfield, M. Fisher, and C. Zeng. Towards Temporal Verification of Swarm Robotic Systems. *Robotics and Autonomous Systems*, 2012.

7 M. Farrell, R. C. Cardoso, L. A. Dennis, C. Dixon, M. Fisher, G. Kourtis, A. Lisitsa, M. Luckcuck, and M. Webster. Modular verification of autonomous space robotics. In *Assurance of Autonomy for Robotic Space Missions Workshop*, 2019.

8 M. Fisher. *An Introduction to Practical Formal Methods Using Temporal Logic*. Wiley, 2011.

9 P. Gainer, C. Dixon, K. Dautenhahn, M. Fisher, U. Hustadt, J. Saunders, and M. Webster. Cruton: Automatic verification of a robotic assistant's behaviours. In L. Petrucci, C. Seceleanu, and A. Cavalcanti, editors, *Critical Systems: Formal Methods and Automated Verification, Proceedings of FMICS-AVoCS*, volume 10471 of *Lecture Notes in Computer Science*, pages 119–133. Springer, 2017. doi:10.1007/978-3-319-67113-0_8.

10 P. Gainer, S. Linker, C. Dixon, U. Hustadt, and M. Fisher. Multi-Scale Verification of Distributed Synchronisation. *Formal Methods in System Design*, 2020.

11 A. Hinton, M. Kwiatkowska, G. Norman, and D. Parker. PRISM: A Tool for Automatic Verification of Probabilistic Systems. In *Proc. 12th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 3920 of *LNCS*, pages 441–444. Springer, 2006.

12 G. J. Holzmann. *The SPIN Model Checker*. Addison-Wesley, 2003.

13 S. Konur, C. Dixon, and M. Fisher. Analysing Robot Swarm Behaviour via Probabilistic Model Checking. *Robotics and Autonomous Systems*, 60(2):199–213, 2012.

14 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.

15 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 *International Conference on Human-Robot Interaction (HRI)*, pages 141–148, Portland, Oregon, USA, 2015. ACM/IEEE.

16 J. Saunders, N. Burke, K. L. Koay, and K. Dautenhahn. A User Friendly Robot Architecture for Re-ablement and Co-learning in A Sensorised Home. In *European AAATE (Associated for the Advancement of Assistive Technology in Europe) Conference*, Vilamoura, Portugal, 2013.

17 M. Sierhuis, W. J. Clancey, and R. J. J. van Hoof. Brahms an agent-oriented language for work practice simulation and multi-agent systems development. In R. H. Bordini, M. Dastani, J. Dix, and A. El Fallah-Seghrouchni, editors, *Multi-Agent Programming, Languages, Tools and Applications*, pages 73–117. Springer, 2009.

18 M. Webster, M. Breza, C. Dixon, M. Fisher, and J. McCann. Exploring the effects of environmental conditions and design choices on IOT systems using formal methods. *Journal of Computational Science*, 2020.

19 M. Webster, C. Dixon, M. Fisher, M. Salem, J. Saunders, K. Koay, and K. Dautenhahn. Formal verification of an autonomous personal robotic assistant. In *Proceedings of Workshop on Formal Verification in Human Machine Systems (FVHMS)*. AAAI, 2014.

20 M. Webster, C. Dixon, M. Fisher, M. Salem, J. Saunders, K. L. Koay, K. Dautenhahn, and J. Saez-Pons. Toward Reliable Autonomous Robotic Assistants Through Formal Verification: A Case Study. *IEEE Transactions on Human-Machine Systems*, 46(2):186–196, April 2016.

21 M. Webster, D. Western, D. Araiza-Illan, C. Dixon, K. Eder, M. Fisher, and A. Pipe. A corroborative approach to verification and validation of human–robot teams. *International Journal of Robotics Research*, 39(1):73–99, 2020.