Autonomy Technology Road Maps

Although not all technology road maps are useful for engineering research insights, three recently published ones are worth a detailed read [1]–[3]. All three reports are extensive, on the order of 100 pages or more. Instead of focusing on hardware [robot and/or unmanned aerial vehicle (UAVs)], they provide a detailed perspective on the technology gaps in various application domains and the research challenges that must be addressed by various research communities, including the IEEE Control Systems Society. Anyone interested in control, robotics, and autonomy for applied and cyberphysical systems will find numerous interesting "research nuggets" within these road maps.

Reference [2] considers robotics for manufacturing, medical, health-care, service, space, and defense applications. While providing deep insights into the possible future impacts of robotic technology in the economy and improving the quality of life, the road map also highlights the need for improvements in several critical technologies, including robust perception, planning, and navigation; intuitive human-robot interaction; and safe robot behavior [2, p. 2]. Examples given include the need in manufacturing applications for safe and secure autonomous navigation in unstructured environments with obstacles, human-driven vehicles, and pedestrians. Improvements in learning and adaptation algorithms are also required to enable robots to operate in uncertain factory environments that can be reconfigured. The need for new planning and control algorithms for high degree-of-freedom robots (combining mobile bases, arms, and end effectors) that address greater uncertainty and wider tolerances is also emphasized. Since many of the applications will involve close interaction between humans and robots, such as with autonomous cars or robotic health care, there is a need to address the current significant gaps that exist in understanding and expressing intent between humans and robots and in improving trust, as enabled through verification and validation of autonomous systems.

Future research will require close interactions with the computer science community.

The Defense Science Board report [1] identifies several key autonomy-enabling technologies that must be improved (perceptual processing, planning, learning, human–robot interaction, natural language understanding, and multiagent coordination) [1, p. 8]. The report also includes several interesting observations, such as the fact that autonomous systems require an increased focus on software "throughout the entire design process, rather than treating it as an afterthought to the development of the hardware" [1, p. 10]. The report also stresses that "treating 'levels of autonomy' as a developmental road map has created a focus on machines, rather than on the human–machine system. This has led to designs that provide specific functions rather than overall resilient capability" [1, p. 4]. The report recommends that the U.S. Department of Defense abandon the widespread use of "levels of autonomy" and use instead a new autonomous systems reference framework that is proposed. In other words, developing a taxonomy can be useful, but dwelling on it too long can do more harm than good.

Faced with the distinct possibility that there could soon be a significant increase in the number of UAVs (remotely piloted and autonomous), a committee of the National Academies reviewed the technology, regulation, legal, and social barriers and issues to the increased use of autonomy in civil aviation systems and aircraft. The analysis and recommendations of the committee are described in [3]. Numerous barriers are listed, but four particularly challenging ones are highlighted: the certification process, decision making by adaptive or nondeterministic systems, trust in adaptive/nondeterministic systems, and verification and validation.

All three of these reports contain many other valuable insights and application domains, but this discussion shows that some common threads that have emerged—the need to improve 1) manned-unmanned teaming, 2) adaptation and learning, and 3) verification, validation, and trust. But what are the paths forward?

Some progress has been made on verification and validation, and these techniques are being used effectively by some large commercial aerospace

Digital Object Identifier 10.1109/MCS.2015.2449672
Date of publication: 16 September 2015

companies [2, p. 20]. But is the high cost of using formal methods to certify autonomous systems plausible for all industries that utilize robots and UAVs? Given the competitiveness of the commercial markets involved, it seems very unlikely, especially as these devices are deployed into more complex and less-well-structured environments. Thus, if autonomous service robots or cars and "increasingly autonomous" UAVs [3] are to be held to the same very high certification standards, there is a need to develop an appropriate certification technology that is much faster, cheaper, and more capable than current methods. The high cost and effort required to write, debug, and certify good software for autonomous systems poses a real barrier to implementation [4], and the control community must focus more attention on resolving this issue.

The situation is similar for adaptive control and learning algorithms. Although there is undoubtedly still room to improve the performance of these techniques, there is a significant gap in the ability to characterize the behavior of adaptive or nondeterministic systems. Gain and phase margins in frequency plots (such as Bode or Nyquist plots) are used extensively to certify nonadaptive controllers because these metrics provide clear guidelines (bandwidth, gain, and phase margins) in terms of the effect of time delay and of the energy retained in the system when excited at different frequencies. Practicing engineers and certifying agencies understand and agree upon these metrics and use them for both analysis and synthesis. Developing a set of techniques to characterize the transient and steady-state behaviors of adaptive or nondeterministic systems is much harder, and it is complicated even further by the fact that there are no quantifiable and agreeable metrics for analyzing the response of adaptive controllers to different conditions.

Future research will require close interactions with the computer science community. While the control community has a lot to offer in terms of the analysis of system stability and reachability, these skills need to be applied to systems with embedded software. The complexity of that analysis will be strongly related to the coding process, the language used, the operating system employed, and the specific hardware implementation. Although more work is required on the performance analysis of specific algorithms, it is likely that much of the future research will focus on developing techniques for analyzing an instantiation of an algorithm in a real system.

Anyone interested in control, robotics, and autonomy for applied and cyberphysical systems will find numerous interesting "research nuggets" within these road maps.

If formal methods do end up being too complex and stochastic simulations are all that can be achieved, then more work is also required to understand the confidence in the assessment tool. For example, if a Monte Carlo simulation is the assessment tool, then techniques are required to assess this assessor—how many simulations are needed, and how accurate are the stochastic models employed in simulating real-world conditions? What about unforeseen or highly unlikely, yet severely detrimental, events? Dealing with such high-impact but low-probability events also requires the careful design of the generating distribution of Monte Carlo simulation initial conditions and parameter selections, in particular, demonstrating that the average performance, although satisfactory, may not be sufficient. Furthermore, effort will be required to educate regulators on 1) the value of stochastic simulations (for example, what the results do, and do not, mean) and 2) the cost and difficulty of obtaining stronger results in the complex environments envisaged. This requires that effort be exerted to appropriately set the expectations of policy makers (and insurance agents) on what analyses are possible and how they should be interpreted.

Developing and characterizing the techniques to verify software and characterize adaptive systems are important steps toward ensuring much broader acceptance (and approved usage) of the advanced concepts; so the sooner they are developed the better. The alternative is that many researchers could find themselves in the undesirable position of having techniques that could improve the performance and resiliency of future autonomous systems, but not having the capability to verify them in a way that leads to the approval to deploy them.

ACKNOWLEDGMENTS

I would like to warmly thank Prof. Eric Feron and Prof. Girish Chowdhary for their valuable comments and discussions on this topic.

REFERENCES

[1] R. Murphy and J. Shields, "Defense science board-task force report–the role of autonomy in DOD systems," Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, Washington, D.C., Tech. Rep CD 1172, July 2012.

[2] (2013). A roadmap for U.S. robotics: From internet to robotics. [Online]. Available: http://roboticsvo.us/sites/default/files/2013%20Robotics%20Roadmap-rs.pdf

[3] Autonomy Research for Civil Aviation: Toward a New Era of Flight, By Committee on Autonomy Research for Civil Aviation, Aeronautics and Space Engineering Board, Division on Engineering and Physical Sciences, National Research Council.

[4] C. Philippe. (2015, May 4). Verification, Validation, and certification challenges for control systems. [Online]. Available: http://ieeecss.org/sites/ieeecss.org/files/documents/IoCT-Part4-10VVCofAerospaceControl-LR.pdf

Jonathan P. How