An outlook for self-awareness in computing systems

Peter Lewis, Marco Platzner, and Xin Yao

Self-awareness theories from psychology are inspiring new methodologies for the engineering of complex computing systems.

As computing systems continue to advance, they are increasingly comprised of large numbers of different types of subsystems, each with their own local perspective and goals and connected in changing network topologies. As a result, having humans understand and manage these systems is becoming increasingly infeasible. Future computing systems, from robots to personal music devices to web services, should be able to achieve advanced levels of autonomous behaviour to adapt themselves at runtime and learn behaviours appropriate to changing conditions. Nevertheless, users engaging with different parts of the system still expect high performance, reliability, security and other qualities.

Such systems will be faced with the challenge of managing trade-offs between these conflicting goals at runtime, both at the global and at the local level, in response to changing conditions. For a system to adapt itself effectively, it is important that it has the ability to be self-aware. Self-awareness is concerned with the availability, collection and representation of knowledge about something, by that something. A self-aware node has knowledge of itself, permitting reasoning and intelligent decision making to support effective autonomous adaptive behaviour. While some work in self-aware computing exists, there is no general methodology for engineering self-aware systems or for validating and benchmarking their behaviour. In the EPiCS (engineering proprioception in computing systems) project,[1] we aim to address this gap.

The study of self-awareness emerged as a field within psychology in the 1960s. Morin[2] defines self-awareness as "the capacity to become the object of one's own attention." As a prerequisite for this, the organism must have the ability to monitor or observe itself. Our work in the EPiCS project is concerned with taking inspiration from self-awareness theory to design methodologies for engineering, validating and benchmarking self-aware systems. In a recent paper,[3] we reviewed key concepts from self-awareness literature that are relevant to our work of realizing self-awareness in computing systems.
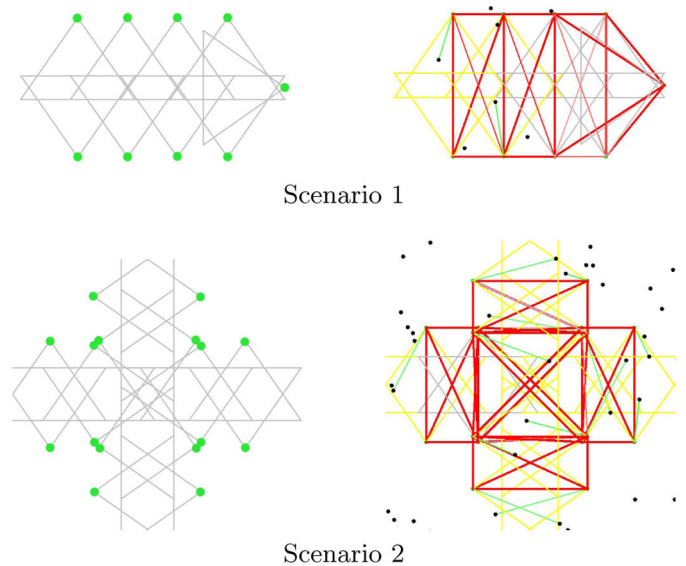
Figure 1. Illustrations of distributed smart camera network scenarios, showing the initial state of the system where no neighbourhood information is present (left) and the learnt neighbourhood structure after some time (right). Each camera is represented by a green circle, with its field of view indicated by the associated triangle. The cameras learn which other cameras have adjacent fields of view over time, through interactions, as the system operates, tracking objects (black dots). Red lines indicate links in the vision neighbourhood graph.

The first of these concepts is the distinction between public and private self-awareness classes,[4] which are concerned with knowledge of phenomena external and internal to the individual, respectively. A second concept is the theory of different levels of self-awareness,[5] ranging from basic awareness of environmental stimuli, by means of an awareness of interactions and time, to an awareness of one's own thoughts. Advanced organisms also engage in meta-self-awareness,[6] an awareness that they themselves are self-aware. A further concept is that self-awareness can be an emergent property of collective systems, even when there is no single component with a global awareness of the whole system.[7] This is a key observation that can contribute to the design of self-aware systems, because it implies

that we do not need a self-aware system to possess a global omniscient node.

There are several clusters of research in computer science and engineering that have explicitly used the term self-awareness.[3] However, there is no common framework for describing or benchmarking the self-awareness properties of these systems or the benefits that self-awareness brings. One cluster focuses on self-awareness as a part of metacognition, and Cox[8] suggests that this is tantamount to the algorithm selection problem. Here, the task is to choose the most efficient algorithm from a set of possibilities. We generalize this idea to decisions with multiple objectives, representative of the conflicting goals of adaptive nodes in dynamic, heterogeneous environments.

At a more fundamental level, some studies focus on how to engineer systems that explicitly consider knowledge about themselves. Examples here include Agarwal et al.,[9] who argue that, through self-awareness, the need to consider availability of resources and constraints at design-time can be avoided or reduced. However, this point is discussed only in a specific application context; our aim is to generalize the idea. Indeed, we have demonstrated this effect in recent work on distributed smart cameras.[10] Here, nodes are cameras that need to track objects moving through their fields of view, while making intelligent decisions at runtime about whom to communicate with and how to exchange tracking responsibilities. Using a market mechanism and a pheromone-inspired learning approach, the outcome is an efficient balance of the trade-off between tracking performance and communication overhead. Unlike previous approaches to this handover problem, the cameras did not require any a priori knowledge of their environment or the camera neighbourhood structure; this is all learnt online (see Figure 1).

Techniques such as this provide options from which meta-self-aware nodes can select. As part of the methodological work undertaken in EPiCS, we are studying how such learning techniques can be combined and chosen at runtime, and how to compare these decision processes. In addition to smart cameras, we are also studying this approach in two quite different case studies: heterogeneous computer clusters and interactive mobile media devices.

Looking forward, it is clear that there is still much to understand about how to incorporate self-awareness properties into computing systems. How can nodes learn and adapt to changing conditions at runtime, while still considering trade-offs between both system goals and the overheads associated with learning about themselves? How should we characterize system-level behaviour that emerges from the interactions of self-aware nodes that have only local information? What can we say about expectations or guarantees of this behaviour? If we are to improve the adaptability of systems through self-awareness, how should we measure this adaptability, and how should we describe assumptions about scenarios and dynamics that we have considered? Is it possible to generalize these types of claims? At the conceptual level, our current work is concerned with establishing methodologies for engineering, validating and benchmarking self-aware systems. However, realizing self-awareness in computing systems will require contributions from many disciplines, including psychology, philosophy, economics, complexity science, artificial and computational intelligence and electronic and software engineering.

Author Information

Peter Lewis and Xin Yao
Cercia
University of Birmingham
Birmingham, United Kingdom

Peter Lewis is a post-doctoral research fellow at the Centre of Excellence for Research in Computational Intelligence and Applications (Cercia) in the School of Computer Science at the University of Birmingham. His research is concerned with investigating algorithms and techniques for achieving self-awareness and self-expression in decentralized computational systems. Particular focuses include economics-inspired computational techniques and online learning algorithms, such as those using evolutionary computation and other nature-inspired techniques.

Xin Yao is a professor of computer science in the School of Computer Science at the University of Birmingham and the director of Cercia. He is a Fellow of the IEEE, a distinguished lecturer of the IEEE Computational Intelligence Society and a distinguished visiting professor at the Nature Inspired Computation and Applications Laboratory (NICAL) of University of Science and Technology of China, in Hefei. His research interests include evolutionary computation (evolutionary optimization, learning and design), neural network ensembles and multiple classifiers (especially on the diversity issue), meta-heuristic algorithms, data mining, computational complexity of evolutionary algorithms and various real-world applications.

Marco Platzner
Department of Computer Science
University of Paderborn
Paderborn, Germany

Marco Platzner is a professor of computer engineering at the University of Paderborn. Previously, he held research positions at the Computer Engineering and Networks Lab at ETH Zurich, Switzerland, the Computer Systems Lab at Stanford University, USA, the GMD – German National Research Center for Information Technology (now Fraunhofer IAIS) in Sankt Augustin, Germany, and the Graz University of Technology, Austria. Platzner holds diploma and PhD degrees in Telematics (Graz University of Technology, 1991 and 1996), and a 'Habilitation' degree for the area of hardware-software co-design (ETH Zurich, 2002). His research interests include reconfigurable computing, hardware-software co-design and parallel architectures. He is a senior member of the IEEE, a member of the ACM, serves on the programme committees of several international conferences (e.g., FPL, FPT, RAW, ERSA, DATE) and acts regularly as reviewer for scientific journals (e.g., IEEE Transactions on Computers, IEEE Transactions on VLSI, ACM Transactions on Embedded Computing Systems, ACM Transactions on Reconfigurable Technology and Systems). He is an associate editor of the International Journal of Reconfigurable Computing, the EURASIP Journal on Embedded Systems and the Journal of Electrical and Computer Engineering. He also serves as a reviewer for several research funding agencies (European Commission, Dutch Technology Foundation, French Agency for Research Funding, Belgian Research Funding and Innovation Stimulation Agency, Swiss National Science Foundation, German Research Foundation and the German Ministry for Economics and Technology).

References

1. http://www.epics-project.eu Engineering Proprioception in Computing Systems (EPiCS) project. Accessed 13 March 2012.
2. A. Morin, *Levels of consciousness and self-awareness: A comparison and integration of various neurocognitive views*, **Consciousness and Cognition 15** (2), pp. 358–371, 2006.
3. P. R. Lewis, A. Chandra, S. Parsons, E. Robinson, K. Glette, R. Bahsoon, J. Torresen, and X. Yao, *A survey of self-awareness and its application in computing systems*, **Proc. 5th IEEE Int'l Conf. Self-Adaptive and Self-Organizing Syst. Workshops (SASOW 2011)**, pp. 102–107, IEEE Computer Soc. Press, 2011.
4. C. Goukens, S. Dewitte, and L. Warlop, *Me, myself, and my choices: The influence of private self-awareness on preference-behavior consistency*. Tech. Rep., Katholieke Universiteit Leuven, 2007.
5. U. Neisser, *The roots of self-knowledge: perceiving self, it, and thou*, **Ann. N.Y. Acad. Sci. 818**, pp. 19–33, 1997.
6. A. Morin and J. Everett, *Conscience de soi et langage interieur: Quelques speculations. [Self-awareness and inner speech: Some speculations]*, **Philosophiques XVII** (2), pp. 169–188, 1990.
7. M. Mitchell, *Self-awareness and control in decentralized systems*, **Working Papers AAAI 2005 Spring Symp. Metacognition in Computation**, pp. 80–85, AAAI Press, 2005.
8. M. Cox, *Metacognition in computation: A selected research review*, **Artificial Intell. 169** (2), pp. 104–141, 2005.
9. A. Agarwal, J. Miller, J. Eastep, D. Wentziaff, and H. Kasture, *Self-aware computing*. Tech. Rep. AFRL-RI-RS-TR-2009-161, MIT, 2009.
10. L. Esterle, P. R. Lewis, M. Bogdanski, B. Rinner, and X. Yao, *A socio-economic approach to online vision graph generation and handover in distributed smart camera networks*, **Proc. 5th ACM/IEEE Int'l Conf. Distributed Smart Cameras (ICDSC 2011)**, pp. 1–6, IEEE Press, 2011.