Autonomy beyond Anomalies and Goals: A Strategic Perspective

Don Perlis PERLIS@CS.UMD.EDU
Department of Computer Science, University of Maryland, College Park, MD 20742 USA

Michael T. Cox MCOX@CS.UMD.EDU
Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 USA

Abstract

In recent years there has been strong interest in both reasoning about goal-identification and selection and metacognitive handling of anomalous situations. These two concerns are usually framed in terms of making agents more autonomous and flexible in dynamic and complex domains. Here we wish to argue that there is a natural unifying perspective that includes both concerns and that may point the way to a yet more powerful kind of autonomy.

1. Introduction

An agent often has routine activities in which it is forming and/or following plans in pursuit of existing goals. And there are also situations in which it has to stop and ask itself: what do I do now? One major example of the latter is that of anomaly-handling: something seems out of the ordinary, contrary to expectation, and might indicate the need to do a form of error-correction. This has been the focus of much recent work, for instance Meta-AQUA (Cox & Ram, 1999), the Metacognitive Loop (Anderson & Perlis, 2005), and other similar efforts. Another example is goal-driven autonomy, in which an agent may autonomously alter or add to its goals if circumstances so warrant (Aha, Cox, & Munoz-Avila, 2013; Aha, Klenk, Muñoz-Avila, Ram, & Shapiro, 2010).

We wish to call attention to a level of processing at which an agent considers quite generally what to do: select from among several existing goals, form a plan to achieve an existing goal, continue with a current plan-in-action, alter such a plan, identify a new goal, abandon a plan or a goal, adopt new subgoals in response to unexpected events, explore opportunities for possible goals or other benefits, do a reality-check of beliefs and expectations, and so on. This could perhaps be called the executive level of processing (borrowing that phrase from cognitive psychology), although the terminology already is in use in various cognitive architectures and so might not be the best choice. Instead, let us call this *the strategic level*.

In what follows we briefly sketch some ideas related to the idea of such a processing level.

2. A Sketch of an Approach

We postulate a *metacognitive monitoring activity (MMAC)* that runs in parallel with an agent's normal routine activity of planning-acting in pursuit of already-identified goals. MMAC will be aware of such routine activities that are underway, and also of their aims and expectations, and of how (at least some) events are actually unfolding (which may or may not be as expected). As MMAC processes this real-time information, it also asks itself over and over: What should I do now? What choices are there? Is there anything that would be better to do than what I am doing? MMAC would normally run in the background, unless something pops into prominence in virtue of a certain salience or threshold that is reached. What can govern such an event?

One way to envision this is in terms of the A-distance (Kifer, Ben-David, & Gehrke, 2004), which assesses alterations in time-series data that exceed a given threshold. This is a crucial kind of hedging-factor. For any given set of expectations will almost certainly fail to be fully identical to observed events. Tiny variations are the norm, and one cannot possibly attend to all of them (nor would it make sense to do so if it were possible). Yet how can such thresholds be determined, when context means everything? In some contexts, a small variation in color or noise-level may be insignificant, and in others may flag major problems or opportunities.

We think that learning is a promising approach here: an agent can learn, for a given context in which it may be operating (or planning to operate in), which are the important things to attend to. This can be partly at the explicit symbolic level (e.g., a teacher can tell the agent some items to watch for and some to ignore) and partly subsymbolic (experience can provide ranges of "normalcy" that the agent trains into its routines.[1]

Now, we doubt A-distance alone will be enough to cover all the cases that we envision for MMAC. For instance, another agent might simply tell our agent that something is important. That is unlikely to cross an A-distance threshold, since conversations may go on all the time, with words flowing rapidly back and forth. It would presumably require a rather high-level reasoning process to understand language sufficiently well to distinguish in a general principled way between, say, "such matters are unimportant" and "don't ever assume that such matters are unimportant." Thus we hypothesize that the strategic level of MMAC is a complex organizer of cognition that integrates activities and seeks to improve their lot.

3. Conclusion

In this paper, we suggest that an underlying research issue exists of considerable potential for enhanced autonomy: *how to design an agent with an effective and general-purpose "what do I do now" capacity*. The capacity bears on many cognitive processes and seems crucial for high-level reasoning in complex ever-changing environments. Researchers have at times studied aspects of what we describe under the MMAC banner in terms of metacognition. Other researchers have examined some of these issues in terms of goal reasoning and goal-driven autonomy. It may be the case that we are all speaking of the same process.

Acknowledgements

This material is based upon work supported by ONR Grants # N00014-12-1-0430 and # N00014-12-1-0172 and by ARO Grant # W911NF-12-1-0471.

[1] A-distance was developed largely for the latter situation, with continuous real-valued data; but recent work has shown that it also is effective for discrete symbolic data (Cox, Oates, Paisner, & Perlis, 2013).

References

Aha, D. W., Cox, M. T., & Munoz-Avila, H. (Eds.) (2013). *Goal reasoning: Papers from the ACS workshop* (Tech. Rep. No. CS-TR-5029). College Park, MD: University of Maryland, Department of Computer Science.

Aha, D. W., Klenk, M., Muñoz-Avila, H., Ram, A., & Shapiro, D. (Eds.) (2010). Goal-Directed Autonomy: Papers from the AAAI Workshop. Menlo Park, CA: AAAI Press.

Anderson, M., & Perlis, D. (2005). Logic, self-awareness and self-improvement. *Journal of Logic and Computation 15*, 21–40.

Cox, M. T., Oates, T., Paisner, M., & Perlis, D. (2013). Detecting change in diverse symbolic worlds. In L. Correia, L. P. Reis, L. M. Gomes, H. Guerra, & P. Cardoso (Eds.), *Advances in Artificial Intelligence, 16th Portuguese Conference on Artificial Intelligence* (pp. 179-190). University of the Azores, Portugal: CMATI.

Cox, M. T., & Ram, A. (1999). Introspective multistrategy learning: On the construction of learning strategies. *Artificial Intelligence, 112*, 1-55.

Kifer, D., Ben-David, S., & Gehrke, J. (2004). Detecting change in data streams. *Proceedings of the Thirtieth Very Large Databases Conference* (pp. 180-191).