

The Robot Baby and Massive Metacognition: Future Vision

Preeti Bhargava¹, Michael T. Cox¹, Tim Oates², Uran Oh¹, Matthew Paisner¹, Don Perlis¹, Jared Shamwell¹

¹University of Maryland, College Park ²University of Maryland, Baltimore County

Email: {prbharga,mcox,uranoh,mpaisner,perlis,shamwell}@cs.umd.edu, {oates@cs.umbc.edu}

Abstract— Metacognition is a powerful tool that can play machine learning and commonsense reasoning off one another synergistically. Based on our previous work in metacognition, we present the vision of a long-term project called “Robot in a Room” in this paper. Similar to newborn babies trying to learn about themselves, surroundings and things that they are able to do, this new project involves a mobile robot that does not know about its environment and its own capabilities i.e. it does not start with a robust self-model. However, it has a primary goal: to learn about itself and explore its environment (the room it lives in).

Index Terms—Autonomous learning agent, developmental robotics, metacognition

I. INTRODUCTION

Autonomous and adaptive agents have long been objects of interest to the AI community; and have scored several successes. However, in many common situations, a robot does not have detailed information about the functioning of its sensors or how the environment may react to its actions. It is in these scenarios that developmental robotics may prove useful.

We envision our robot, Jack (Fig. 1), eventually becoming aware of its own manipulators and sensors. That is, as he uses them in an initial “random flailing mode”, he will come to notice regularities (such as certain visual flows coinciding with certain wheel-motor impulses). This in turn will provide the initial formation of abstract concepts such as “arm” and perhaps even “self” (similar to developing meaningful representations of objects as described in [1]). Jack will also store observations in a memory schema designed to associate actions with input from its visual sensors, as discussed in Olsson et al [2]. The key difference in our work is that we introduce the Metacognitive Loop (MCL) [3] that will serve as an umbrella for guiding all processes in the robot. Jack will use MCL autonomously to guide the creation of subgoals and plans in the pursuit of its primary goal: to learn. MCL will also coordinate various types of learning that the robot will undertake (supervised and unsupervised); internal (learning how to improve its own processes) and external (learning about the nature of the world). MCL's primary advantage in these areas is adaptability – it will allow Jack to adjust to changes in the room and its contents as well as limitations in its own abilities.

II. CURRENT WORK

As the first step of this project [4], Jack uses the Growing Neural Gas (GNG) algorithm [5] to find useful patterns in a wide range of data types, including visual, tactile and even

internal state data. GNG is an unsupervised learning algorithm that is intended to give an agent that has little or no apriori knowledge of its environment, the ability to create a model of states in its environment autonomously. It will allow Jack to access powerful tools such as planning or logical reasoning that require a description of discrete states and actions to perform optimally.




Fig. 1. The robot that is used for the project, Jack, has the visual ability to see its environment and detect objects. Its mobility allows Jack to navigate the room and interact with the world with its arm and gripper.

III. METACOGNITIVE LOOP (MCL)

We conceived MCL as a generalized and reason-driven anomaly-handling capability. It is largely domain independent, involves only modest amounts of background knowledge and computation, and can be implemented in any automated system. The basic procedure consists of the “NAG” cycle: (i) Note expectations in order to detect any anomaly that might arise, (ii) Assess it in terms of available responses, and (iii) Guide any chosen response into place (and monitor the progress of that response for further anomalies). Fig. 2 depicts the host system and the generalized MCL module. The MCL module implements three ontologies, each implemented as a Bayesian network: *indications* ontology to represent anomalies, *failures* ontology for assessment, and *responses* ontology for repairs. The core nodes in the indications and responses ontologies represent domain-general concepts of anomalies and responses while the fringe nodes at the bottom represent specific anomalies and responses that might be suggested to the system under supervision. Links between nodes express relationships between the concepts they represent; for example, the node representing an anomalous sensor reading might be linked to that representing a faulty sensor failure. During operation, the system can adjust its expectations based on its experience.




Fig. 2 Host and the generalized metacognition module (MCL)

When an expectation is violated, the host will report it to MCL by activating indications fringe nodes corresponding to the type of violation. MCL will then use Bayesian inference to analyze the problem abstractly, producing an output that is articulated through the response fringe nodes in the form of an action that the host can carry out.

IV. FUTURE DIRECTIONS WITH MCL

A. GNG

Jack has been using the GNG algorithm combined with knowledge of its own rotation to address a simple problem: "How many red things are there in the room?"[4]. The GNG algorithm requires tuning several parameters and we intend to use MCL to help Jack learn to perform such adjustments.

B. A Distance

A-distance [6] is a metric for detecting shifts in data streams combined with classification margins to detect domain shifts. Jack can use MCL along with A-distance to detect changing or anomalous data, and perhaps even to recognize when the underlying paradigm of its environment has shifted.

C. Learning and Planning

As Jack learns about itself and its environment, it should be able to distinguish between interesting concepts worth learning and uninteresting ones that are unlikely to lead anywhere. It should be able to set its own goals, generate its own plans and correct its methodology in the face of failure. MCL, along with reinforcement learning, predicated on Jack's underlying desire for knowledge, can be used to guide this process.

D. Self-reflection and self-other distinction

Self-reflection in a robot is the ability of the robot to be aware of its own sensors, manipulators, computational capabilities and internal processes. Several approaches have been developed to accomplish this, including neural networks [7], and multi-agent architectures [8]. There is also the problem of self-other distinction. Nagai and colleagues [9] propose a computational model for early development of the mirror neuron system (MNS). Their robot observes motion of self and others and is able to make the distinction between them with increasing spatiotemporal resolution.

E. Alfred (Natural Language Processing)

Alfred 2.0 [10] is a universal interfacing agent that accepts a small set of English sentences and translates them into commands appropriate for different domains provided it is supplied with the requisite grammatical information about both the domain language and the English sentences that will be

used as input. It is also designed to be a general, common sense reasoner and hence is integrated with MCL. In time, we plan to use Alfred as an intermediary between Jack and its human controllers, allowing us to give him commands in English.

V. DISCUSSION

Our greatest obstacle is simple: Jack can execute commands such as "turn left and lower arm-position," and observe the changed optical flow that results. However, at present, its reasoner has no concept of an arm; it must be told to move it programmatically. Our hope is that the robot's learning modules (especially GNG) can eventually fill that gap. The fundamental problem is that Jack needs symbols with which to reason. GNG can provide these, in the form of patterns detected in Jack's optical flow, command-traces, and other varieties of monitoring. In particular, we anticipate that it will come to associate the execution of certain commands with certain patterns of visual flow ("when I'm facing a view that looks like this and turn left, I tend to see a view that looks like that"); then with certain stable regularities in the room ("when I'm facing the brown rectangular thing and turn left, I tend to see a black rectangular thing covered with white lines") and perhaps even with self/non-self-distinctions ("when I execute move-arm, my arm always moves, while books and chairs are more mercurial in their responses").

This will be a major test of our "commonsense core" hypothesis: that there is a modest set of general-purpose and computationally tractable anomaly-handling strategies that can allow a system to autonomously adjust to and learn from unanticipated circumstances so as to become more skilled, more knowledgeable, and more flexible.

REFERENCES

- [1] P.R. Cohen, T. Oates, C.R. Beal, and N. Adams, "Contentful Mental States for Robot Baby", Proceedings of AAAI'02, no. 18, pages 126-131
- [2] L. Olsson, C. L. Nehaniv, and D. Polani, "From unknown sensors and actuators to actions grounded in sensorimotor perceptions," *Connection Science*, vol. 18, no. 2, pp. 121-144, 2006.
- [3] M. Schmill, M. Anderson, S. Fults, D. Josyula, T. Oates, D. Perlis, H. Shahri, S. Wilson, and D. Wright. (2011). "The metacognitive loop and reasoning about anomalies." In M. T. Cox & A. Raja (Eds.), *Metareasoning: Thinking about thinking* (pp. 183-198). Cambridge, MA: MIT Press.
- [4] J. Shamwell, T. Oates, P. Bhargava, M. T. Cox, U. Oh, M. Paisner and D. Perlis, "The Robot Baby and Massive Metacognition: Early Steps via Growing Neural Gas", 2012 *IEEE International Conference on Development and Learning (ICDL)*, in press.
- [5] B. Fritzsche, "A Growing Neural Gas Networks Learns Topologies", *Advances in Neural Information Processing Systems*, 7, Brussels, 625-632
- [6] M. Dredze, T. Oates, and C. Piatko. "We're Not in Kansas Anymore: Detecting Domain Changes in Streams", *Empirical Methods in Natural Language Processing (EMNLP)*, pp. 585-595, 2010.
- [7] J.C. Zagal and H. Lipson. "Self-Reflection in Evolutionary Robotics: Resilient Adaptation with a Minimum of Physical Exploration". In *Proceedings of the Genetic and Evolutionary Computation Conference, Late Breaking Papers, GECCO (2009)*
- [8] K. Kawamura, D. C. Noelle, K.A. Hambuchen, T.E. Rogers and E. Turkay, "A Multi-Agent Approach to Self- Reflection for Cognitive Robotics", *International Conference on Advanced Robotics*, Coimbra, Portugal, June 30- July 3 2003
- [9] Y. Nagai, Y. Kawai and M. Asada, "Emergence of Mirror Neuron System: Immature vision leads to self-other correspondence", 2011 *IEEE International Conference on Development and Learning (ICDL)*
- [10] D. P. Josyula. "A Unified Theory of Acting and Agency for a Universal Interfacing Agent", *PhD Thesis*, University of Maryland, 2005.