

A Multi-Agent Control Architecture for a Rescue Robot

Adam Haber

ARC Centre of Excellence in Autonomous Systems
School of Computer Science and Engineering, UNSW
Sydney, NSW, 2052, Australia.
a.haber@student.unsw.edu.au

Despite many years of research and progress in the field of artificial intelligence, there is still no universally accepted definition of the word *intelligence*. Instead, researchers have identified a multitude of tasks, skills, and behaviours that *require* intelligence, such as perception, learning, planning, reasoning, and many others. Much A.I research is focused on improving each of these separate skills within an isolated problem domain, leaving aside the problem of integrating the various components. Although the initial thrust of A.I in the 1950s was towards this kind of integrated system, in recent times the problem of integration has become conspicuous by its absence in the field, but is essential to improve our design of complete intelligent systems, and consequently our understanding of our own brains.

A subfield of A.I research known as cognitive architecture has developed, which attempts to address this problem explicitly; research in this area seeks to identify the components and associated interactions necessary for an intelligent system. Typically, work in this field is conducted by describing and justifying a set of principles or axioms, and then describing an architecture which is governed by these axioms, and occasionally, implementing an intelligent system based on the architecture. Prominent examples in the literature are ACT-R (Anderson 2007), Soar (Laird, Newell, and Rosenbloom 1987), ICARUS (Langley and Choi 2006) and RCS (Albus 2005). Despite the relatively large number of proposed architectures in the field, implementations are rare, thus it has proven difficult to evaluate the success of the various axioms, and methodologies proposed (Langley, Laird, and Rogers 2009).

A natural step towards addressing this difficulty is implementing robotic systems which are controlled by systems incorporating cognitive architectures; simultaneously advancing the development of autonomous systems (Long and Hanford 2010). Such systems represent a great potential to improve the state of robotics, and present new insight into embodied cognition (Cox 2007). My research aims to work towards the grounding of cognitive architectures in testable implementations, by proposing and implementing a cognitive architecture designed specifically to control a mobile robot. The remainder of this extended abstract begins by outlining the principles which form the backbone of the archi-

ture, the testing environment in which the implementation will be embedded, and then describes the work completed so far. Finally we will address the body of work still to be completed and plans for future research.

Architectural principles

Design of our architecture is guided by three principles which were derived by myself and my supervisor, namely: multiplicity, heterogeneity, and adaptability.

Multiplicity. One of the few points of consensus within cognitive architecture research is that architectures must be composed of modular, independent components. This is a consequence of the multi-faceted nature of information processing, and confirmed the modular structure of the human brain. Most architectures are composed of less than ten components, each of which has a single (very complex) function. Many authors have argued convincingly that a much finer level of modularity must exist for sufficient resourcefulness (Minsky 1988). Accordingly, we introduce the first and most central principle of our architecture: that this modularity exists at several scales, both the large scale modularity common to existing cognitive architectures, and also, a much finer-grained modularity, with each module contain a number of smaller entities which we will refer to as agents (we use this term without reference to agent-orienting software engineering). This type of design often arises naturally in robotics - usually a multiplicity of methods can provide a more robust solution than a monolithic algorithm (Sammut 2010).

Heterogeneity. Many existing architectures commit to a single form for knowledge representation or reasoning, such as production rules in ACT-R and Soar, or conceptual inference in ICARUS, but a system with heterogenous components can prove more resourceful by taking advantage of the different strong points of different systems.

Adaptability. Though the principle of multiplicity ensures our architecture is provided with a collection of agents with which to solve problems, it does not specify any organisation, in particular, which agents should be applied or active at any given time; although certainly this question is dependent on the environment. We thus arrive at the third principle of our architecture: that the agents which are active (or take precedence), and thus the composition of the architecture, should respond to changes in the environment

- the activation and deactivation of agents within that architecture should be *learned*.

Completed Work

A challenging test domain for the proposed system is provided by robotic urban search and rescue (USAR). This task requires exploration within an unstructured environment, where the system must pursue a multi-objective search and be faced with unpredictable obstacles. Widely accepted standards and metrics for USAR have been developed for the international RoboCup Rescue competition, and we perform our experiments in a test area modelled on its competition arena. Initial experiments were and are being conducted in a high fidelity 3D physics-based simulation written (by myself) for this purpose based on the open source jMonkeyEngine3 game engine. No unrealistic information about the environment is given to the robot, all data comes from a suite of simulated noisy sensors. Later experiments will be performed using a real robot.

The complete control architecture has three large modules, each of which is composed of a collection of smaller agents as described above, and a controller which adaptively moderates the agents interaction. The functions of the modules are passive perceptual system, an active perceptual, and a goal-action. The task of the passive perception system is to search each frame of all available sensor data and determine whether any victims are in view. In keeping with the multiplicity principle, this is achieved by three separate agents, each of which act on sensor data independently; they are warm-skin detector, head-shape detector, and movement detector. Warm-skin detection operates on data from thermal and colour cameras to detect regions which are both skin coloured, and within the temperature of human skin. Head-shape detection uses information from a 3D depth camera to search for objects which have an approximately head-like shape (Xia, Chen, and Aggarwal 2010), while movement detection uses a frame-differencing algorithm acting on the colour camera data (Bradski 2000).

The controller for this module is provided by a system that learns which agent is preferable for various perceptual settings and activates the agents accordingly. The verdicts from the various agents are then combined, weighted by their learned confidence values, to produce the systems best guess at whether there is a victim(s) in shot and its location(s). This system has been implemented using a Naive Bayes Classifier which receives ground-truth feedback from the simulation (for the physical robot, a teacher must provide the ground truth) to incrementally learn the agent activation hypothesis, and can successfully identify victims under all the conditions simulated (including in shadow, and partially obscured). Both the simulation, and the implementation were entirely my own work except where otherwise cited.

Current and future work

The main thrust of my research, and the area with which I am currently occupied, is the implementation of higher cognitive functions; in particular, the last two modules de-

scribed above. Active perception takes advantage of the capability of a robot both to move and to influence its environment; this module is designed to improve the detection of victims by improving the information available to the robot. In this module, the agents will be look-around-by-moving-robot, look-around-by-moving-head, etc. The controller for this module requires more sophisticated learning and planning algorithms, since only one agent can operate at a time in contrast to passive perception where all agents can operate in parallel. A learning system based on Inductive Logic Programming (ILP) to provide adaptability is a current area of my research.

Future work consists of implementation of the goal-action module, which will produce the long-term behaviour for the robot. This module may prove most challenging and innovative, since it is designed to allow the robot to pursue multiple goals concurrently, while switching (productively) between their pursuit; a difficult problem in A.I robotics. The constituent agents of the module will each pursue their own goal, i.e observe-victims, do-not-get-stuck, etc. The aim of the controller for this module is to allow these agents to compete, and dictate which goal should be pursued at which moment, such that the result is productive, and the decisions made (such as whether to risk harm to the robot to reach a victim) are appropriate. Testing and evaluation of the architectures performance both in simulation and in reality will be according to NIST USAR metrics, although the environments will be modified to produce more cognitively challenging situations.

References

- Albus, J. 2005. Rcs: A cognitive architecture for intelligent multi-agent systems. *Annual Reviews in Control*.
- Anderson, J. R. 2007. *How can the human mind occur in the physical universe?* New York: Oxford University Press.
- Bradski, G. 2000. The opencv library. *Dr. Dobb's Journal of Software Tools*.
- Cox, M. 2007. Perpetual self-aware cognitive agents. *AI magazine* 28(1):32.
- Laird, J. E.; Newell, A.; and Rosenbloom, P. 1987. Soar: An architecture for general intelligence. *Artificial intelligence*.
- Langley, P., and Choi, D. 2006. A unified cognitive architecture for physical agents. *Proceedings of the National Conference on Artificial Intelligence* 21(2):1469.
- Langley, P.; Laird, J. E.; and Rogers, S. 2009. Cognitive architectures: Research issues and challenges. *Cognitive Systems Research* 10(2):141–160.
- Long, L., and Hanford, S. 2010. Cognitive robotics using vision and mapping systems with soar. *SPIE Defence and Security Conference*.
- Minsky, M. 1988. *The Society of Mind*. New York: Simon and Schuster.
- Sammur, C. 2010. Robot soccer. *WIREs Cognitive Science* 1:824–833.
- Xia, L.; Chen, C.; and Aggarwal, J. K. 2010. Human detection using depth information by kinect. *Asian Conference on Computer Vision*.