



This space is reserved for the Procedia header, do not use it



A Software Architecture for Generally Intelligent Humanoid Robotics

Ben Goertzel¹, David Hanson², and Gino Yu³

¹ OpenCog Foundation, Hong Kong, China
ben@goertzel.org

² Hanson Robotics, Hong Kong, China

³ School of Design, Hong Kong Polytechnic University

Abstract

This paper summarizes the authors' thinking regarding the design of a software framework for interfacing between general-intelligence-oriented software systems and complex mobile robots, including humanoid robots. The framework describes incorporates perception synthesis, action orchestration, and high level control, and is designed to effectively leverage existing relevant software frameworks such as ROS and Blender. An initial case study motivating this work is the use of the OpenCog AGI (Artificial General Intelligence) software framework to help control humanoid robots created by Hanson Robotics.

Keywords: Consciousness, Information Integration, Global Workspace, Nonlinear Dynamics

1 Introduction

One approach to achieving intelligent behavior in complex mobile robots such as humanoids, is to integrate a general-purpose cognitive architecture with an appropriate conglomeration of robotics-specific software. This is by no means the only possible approach, e.g. purely subsumption-oriented [1] or biomorphic [4] approaches are also interesting and potentially effective. However, the integrative cognitive architecture approach is also worthy of exploration.

In this brief position paper we present some ideas regarding the design and engineering of a software system for connecting cognitive software with robotics software to achieve an integrated cognitive architecture for intelligent robotics. The initial motivation for this system is a current project involving the use of the OpenCog Artificial General Intelligence (AGI) software system ¹ [2, 3] to control humanoid robots created by Hanson Robotics. However, the goal is to create a general-purpose architecture, not restricted to this particular application.

¹ <http://opencog.org>

2 High Level Requirements

The creation of a requirements specification for a software framework connecting cognitive architectures to robots, is itself a complex pursuit, and it is hoped that discussion of this conference paper will facilitate collection of ideas in this regard from the community, which can then be used to formulate a detailed requirements list. What is presented here is merely an outline of high level requirements, intended to evoke the basic intentions underlying the software framework:

- Perception: Intake of perceptual data from a variety of different robotic sensors, each with their own drivers and preprocessors; synthesis of data from multiple sensors into an overall perceptual world-view, or (depending on the nature of the data) a collection of overlapping partial world-views; export of processed, synthesized versions of perceptual data to cognitive systems
- Action: Intake of high-level or medium-level action plans from cognitive systems, and transformation of these into detailed action plans suitable for enaction by robotic systems; intake of detailed action plans from external systems (e.g. telerobotics controllers); delivery of detailed action plans to specific robotic systems, for driving robot actions; synthesis of multiple high or medium level action plans, intended for simultaneous or overlapping action, into a single detailed action plan suitable for driving action of a specific robotic system
- Control / Coordination: Determination of which actions a robot should take, given its perceptions at a given time and also conditional on the understanding of the current goals and context, where two important cases are: the understanding of goals and context is provided by reference to an external cognitive system, and the understanding of goals and context is provided by relatively simple ruled supplied directly within the robot control/coordination framework (hence allowing for rapid response in the case of relatively simple “reflex” behaviors or other behaviors that are desired to be preprogrammed in a specific context); and, receipt of robot control signals from remote controls or other user interfaces (noting that in some cases, control signals from UIs may be treated as absolute commands; in other cases they may be treated as special perceptual signals, to be considered alongside other perceptions in formulating appropriate action plans)

3 Reference Tools

Among the numerous tools needed to fulfill the above broad requirements are: A cognitive architecture, capable of ingesting perceptual data, forming a model of the current context of a robot therefrom, and suggesting appropriate actions; a robot simulator; a software framework for interacting with robots; and, optionally, software for generating robot movements, e.g. a 3D graphic art program and/or motion capture program.

To the extent these various tools are already able to work together, development will be easier. For example, one option would be to make use of the Blender infrastructure, e.g. MORSE, integrated with Blender, for robot simulation; or Gazebo for more accurate simulation of certain aspects of robot movement such as dynamic stabilization; ROS, integrated with Blender / MORSE / Gazebo, for robot control; BGE, Blender Game Engine, for scripting interactions involving simulated robots; Blender for art asset creation (including perhaps additional plugins like MakeHuman).

the cognitive architecture, as stated above our current work involves OpenCog but the framework sketched here doesn't depend on OpenCog specifically.

4 Architecture Sketch

At a very high level, we suggest to fulfill the requirements posited above via creation of the following software components: Action Orchestrator (AO), Action Creator (AC) , Perception Synthesizer (PS) , Robot Controller (RC) , Action Data Source (ADS), an interface that might wrap up more than one such data source, and Robot Control User Interface (RCUI).

There are plausible deployment scenarios in which the AC, ADS or RCUI might be implemented on a variety of platforms (e.g. Windows, Mac; or for the RCUI, Android or iOS). On the other hand, we can assume for the foreseeable future that the RC, PS and AO will be run on Linux. Hence one deployment possibility is to deploy the RC, PS and AO as ROS nodes; but implement communication with the other components via some more platform-independent method like ZeroMQ.

Figure 1 depicts a high-level overview of the architecture.

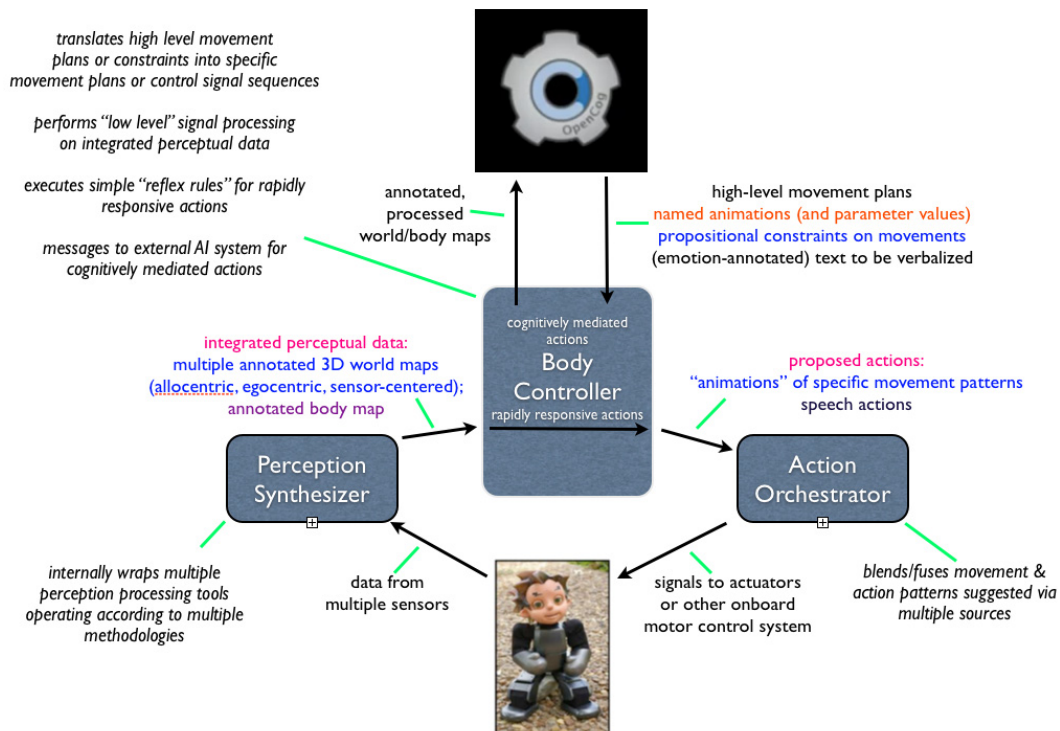


Figure 1: High-Level Overview of Proposed Software Architecture

4.1 Action Orchestrator

The AO receives signals from the RC, indicating what set of high level actions to carry out at a given point in time. It then directs the robot how to carry out the actions. In general the

signals received by the AO would be lists of action descriptions. Each action description would contain an action name and optionally some (discrete or quantitative) action parameters. The actions usable by the AO should be specified in a configuration file, not hard coded, so it is easy to add new actions or change the parameter-sets of existing ones.

The AO could be connected to a physical robot, or else to a simulated robot in the chosen simulation platform. This would be determined in a configuration file. Ideally the commands sent from the AO to the physical or simulated robot would be identical.

In an initial implementation of the AO, each action type would simply come with a hard-coded mapping into a set of detailed motor commands for a particular robot. A slightly more advanced implementation would handle blending, i.e. would have a way of dealing with multiple simultaneous action requests by blending the corresponding lists of motor commands.

Ultimately, a sophisticated AO would contain an internal hierarchical representation of actions, representing each action as a hierarchically structured set of trajectories of body parts. This would allow flexible blending of actions that started at different times and involved overlapping parts of the body. But this represents a significant AI undertaking, specification of which goes beyond the scope of this document. (Some early thoughts in this regard are given in OpenCog documentation, e.g. in Engineering General Intelligence.)

4.2 Action Creator

The AC translates movement scripts created externally, into robot action sequences to be enacted in a robot simulator or physical robot.

As one example, the AC would translate animations created in an art program (e.g. Blender, Maya, etc.) into robot action sequences. If the bone structure of the animated character and the robot are identical or very similar, this is relatively straightforward (though still not trivial as it requires knowledge of the parameters of the robot motors, etc.). If the bone structures are different this requires some fancier mathematical translation. The action sequences created in the AC are then exported to the AO, which can invoke them as needed.

4.3 Perception Synthesizer

The PS receives data from the robot, e.g. sound from a microphone, 3D visual data from a Kinect, visual data from a webcam, etc. It then processes this data into an appropriate form for further utilization. For example, the PS might overlay webcam visual data with Kinect visual data. Or the PS might be configured to pass along to the RC only those parts of the visual data stream that have changed since the last message passed along to the RC. Basic signal processing could also be done in the PS. In essence the PS cleans up and filters data for passing along to the RC (and potentially on to other sources as well, e.g. a data store or OpenCogs Atomspace).

The PS could receive information from the physical robot, or the simulated robot. Ideally the data received in the two cases would be identical in format. (Of course the particulars of e.g. visual data obtained from a simulation would not be the same as from actual visual sensors, though.)

In a future version the AO and PS could work in collaboration, but initially they will be separately operating pieces of software.

Assuming one utilizes ROS as an infrastructure, the AO, PS and RC could be three separate ROS nodes, in frequent communication. The AO and PS would then communicate with other ROS nodes representing particular aspects of the robot or associated software.

4.3.1 Multiple Spatial Maps

One subtlety regarding the PS is hinted by the fact that the human brain possesses multiple spatial maps: an allocentric (i.e. top-down, third person) map, an egocentric face-centered map, two egocentric eye-centered maps, one for each eye, and a map of the body to which it corresponds (a somatosensory map) separate from all these visual world-maps.

To emulate this sort of approach, the output from the PS would be an object consisting of multiple components, each component corresponding to one of: an allocentric map (in the form of something like an OctoMap, a common ROS component) , an egocentric face-centered map (assuming we're dealing with a robot that has a face or some analogue, i.e. that is facing in some direction) , a camera-centered map, i.e one for each of the robot's cameras (these are like the brain's eye-centered maps) , a body-map, which is really a data structure giving relevant information about the parts of the robot's body (this may be fairly robot-specific I would think, as different robot bodies will provide different sorts of information)

The information sent from the PerceptionSynthesizer to the RobotController would then be the "diff" between the last map sent and the current map, for each of these component maps.

The key point of this approach would be that: sometimes a robot will have a clear idea of where certain things observed are relative to its own location, but NOT of where it (or these things observed) are relative to the allocentric map; whereas, sometimes a robot will have a clear idea of where certain things are relative to its allocentric map of the room (e.g. if something against a wall across the room shifts to the right, it can notice this in terms of a movement of the object relative to the wall), even if it's unclear of its own location (in this case it is clear about an event's position on the allocentric map but not the egocentric map).

Of course, the robot would also contain, on its allocentric map, knowledge about where it believes it is located on the allocentric map. However, this knowledge could be represented probabilistically, i.e. it might not be sure exactly where it is.

4.4 Action Data Source

An Action Data Source comprises an external piece of software that produces descriptions of actions to be sent to the Action Creator.

One example would be a Telerobotics Data Source, containing a mechanism for observing the movement of a human face, and packaging these movements as a stream to pass to the PS for perceptual understanding; and in a tele-operation scenario, to the AC as well.

4.5 Robot Control UI

In many cases it will be useful to have a human being directly control a robot, via manipulating a User Interface that sends messages to the RC.

E.g., one possibility is to make such a UI in the form of an Android phone app. This avoids the need to use special hardware, but makes the UI mobile, so that a person can hold the phone in their hand like a remote control and manipulate the robot. It also leaves the opportunity open for advanced manipulation of control parameters as well as simple controls.

4.6 Robot Controller

The RC contains the logic for choosing actions based on perceptions and based on user controls. Initially, for a standalone application, this logic could be fairly simple. For instance, instances of non user control driven action logic might include: "Move the head to face someone who is

detected to be talking”, ”Move the head to look at the nearest human face, or at a human face that has recently become mobile”, ”Imitate the movements of the person being tracked by the ADS”, ”Imitate the movements of the person being tracked by the ADS; or if this person is looking at a certain person, then look at that person.”

The best way to formalize this action logic is not clear. One option is to leverage OpenCog, and use the OpenCog AtomSpace to contain symbolic abstractions of perceptual data, and action commands; using Implication relations in OpenCogs AtomSpace to encode (potentially probabilistically weighted) action logic rules. OpenCog provides a highly general and flexible formalism here, including probability weighting of actions; and also provides a framework in which learning of action logic rules can naturally take place (although the initial action logic rules will be hard coded).

The RC must be able to receive control signals from external sources as well – from humans via the RCUI, or from other software programs. The right messaging protocol must be chosen carefully here; ZeroMQ is one candidate currently under consideration.

5 Conclusion

A rough sketch of a software architecture for interfacing between cognitive architectures and robotics software systems has been outlined. This represents current work in progress, on which feedback is solicited from others working in the area. Feedback will be incorporated in the ongoing design and engineering work, hopefully resulting in a system with broad utility beyond the initial application of flexibly interfacing OpenCog with Hanson Robotics robots.

References

- [1] R. A. Brooks. *Flesh and Machines*, Pantheon Books, New York. NY, 2002.
- [2] Ben Goertzel, Cassio Pennachin, and Nil Geisweiller. *Engineering General Intelligence, Part 1: A Path to Advanced AGI via Embodied Learning and Cognitive Synergy*. Springer: Atlantis Thinking Machines, 2013.
- [3] Ben Goertzel, Cassio Pennachin, and Nil Geisweiller. *Engineering General Intelligence, Part 2: The CogPrime Architecture for Integrative, Embodied AGI*. Springer: Atlantis Thinking Machines, 2013.
- [4] Brosl Hasslacher and Mark W. Tilden. Living machines. *Robotics and Autonomous Systems*, page 143169, 1995.