A Methodology for Developing Competent Agents Without Sensor and Actuator Profusion*

Henry H. Hexmoor

Department of Computer Science
SUNY at Buffalo, Buffalo NY 14260
hexmoor@cs.buffalo.edu

March 30, 1994

1 Introduction

Making a cup of coffee; frying and serving an egg; preparing buttered toast; playing catch; inserting and playing a video-tape; vacuum/mopping the floor or mowing the lawn; digging a hole/trench; (un) locking a door with a key; opening, passing through, and closing a door; feeding someone using a fork, knife, spoon, cup, etc.; retrieving a screwdriver from the toolbox in the garage; folding clothing; and moving large furniture objects (boxes, chairs, furniture) are all tasks that all performed by humans. An agent that performs these tasks requires human-like physiology. How about find a door and get out of an office space and fetch a coffee pot. These were tasks that a Lego robot was expected to perform at the 1993 AAAI robot building event. Obviously, those Lego robots are in no position to do any of the tasks in the first list. Nevertheless, the Lego robots were considerably competent at variants of tasks in the second list just as humans are competent in tasks in the first list.

We have developed an architecture for intelligent autonomous agents that endows them with physical abilities that are "natural" for the agents, i.e., correspond to their physiology. We call our architecture GLAIR (Grounded Layered Architecture with Integrated Reasoning), [HN92, HLS92, HCBS93, HLCS93, HLS93, LHS93, HS93].

GLAIR models encoding and abstraction of sensory information into an agent's concurrent unconscious and conscious processes. In these concurrent processes, the sensory information is interpreted in different ways and tied to motor actions in various ways.

We believe, competence and versatility of agents is derived more from profusion of pathways between sensing and acting and interpretations of sensory data than in profusion of sensors and actuators. The economy in sensors and actuators is traded off with diversity of concurrent ways of using them.

2 GLAIR

Figure 1 schematically presents GLAIR. GLAIR architecture is characterized by a three-level organization into a Knowledge Level (KL), a Perceptuo-Motor Level (PML), and a Sensori-Actuator Level (SAL). GLAIR is a general multi-level architecture for autonomous cognitive agents with integrated sensory and motor capabilities. GLAIR offers an "unconscious" layer for modeling tasks that exhibit a close affinity between sensing and acting, i.e., behavior based AI modules, and a "conscious" layer for modeling tasks that exhibit delays between sensing and acting. GLAIR provides learning mechanisms that allow for autonomous agents to learn emergent behaviors and add it to their repertoire of behaviors. GLAIR motivates the concept of *embodiment*. Concept learning provides an important motivation for embodiment. Win-

*This work was supported in part by Equipment Grant No. EDUD-US-932022 from SUN Microsystems Computer Corporation. This paper was presented in the AAAI spring symposium workshop "Toward Physical Interaction and Manipulation", March 23, 1994, Stanford, California.

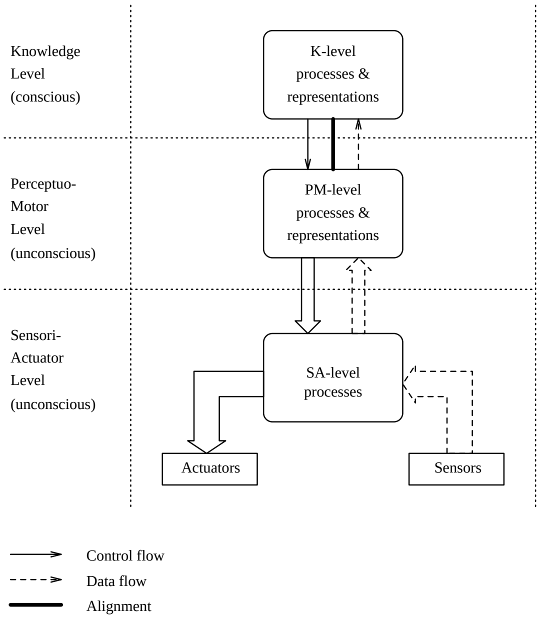

Figure 1: Schematic representation of the GLAIR architecture. Width of control and data paths suggests the amount of information passing through (bandwidth). Sensors include both world-sensors and propriosensors.

ston's Arch program [Win75] is an early example of a system that learns concepts through examples. This program relies heavily on feature analysis. However, the feature concepts and the concepts of objects learned lack embodiment, as is typical for traditional symbolic AI work. Most of this work never gets implemented in actual agents that can physically interact with their environment (other than via a keyboard and monitor, which is a rather uninteresting case), i.e. the symbols never inhabit a body. It is this kind of approach that Brooks criticizes in papers like [Bro90]. According to Brooks, symbolic representations should be matched to the agent's sensors and actuators, and this grounding provides the constraints on symbols necessary for them to be truly useful. It is this matching of representations to the agent's sensors and actuators that we call embodiment.

The most general motivation behind our work is the desire to be able to "program" a robotic autonomous agent by telling it *what* to do at a high level and have it "understand", rather than telling it *how* to do something in terms of primitive operations, with little or no "understanding".[1] For instance, we want to tell it to go find a red pen, pick it up, and bring it to us, and not have to program it at a low level to do these things.[2] One might say that we want to communicate with the robot at the *speech act* level. To do this, the agent needs a set of general-purpose (relative to its own body) perceptual and motor capabilities along with an "understanding" of these capabilities. The agent also needs a set of concepts which are similar enough to ours to enable easy communication (as [Win75] already mentioned). The best way to accomplish this, in our opinion, is to endow the agent with embodied concepts, grounded in perception and action.

A major objective for GLAIR is learning emergent behaviors. Like Agre's improvised actions [AC87] and Brooks's subsumption [Bro85] we believe complex behaviors emerge from interaction of the agent with its environment without planning. However, previous work in this area hard-coded primitive actions and did not attempt to learn the improvised behavior.

3 Principles of Physical Interaction and Manipulation

Our first *principle of physical interaction and manipulation* is that there is a mapping/correspondence between physiology and a set of tasks. Of Course, given a set of tasks, we can reverse engineer and find a physiology that suits the tasks. We are interested in agents that learn manipulation and embodied concepts on their own. This is in contrast to traditional robotics where concepts are defined that are suitable for humans and human physiology. These concepts are then encoded in agents and they are expected to be as competent in tasks as their human counterparts. We hypothesize that human natural language comprehension mechanism translates representations to embodied representa-

[1] There was a similar motivation behind the move from imperative to logic-based or declarative programming languages, but there are no further similarities.

[2] Retrieving "canned" parameterized routines is still a low-level programming style that we want to avoid.

tions and vice versa. With this we believe we need to model multi-level representations within an agent. Our work focuses on embodied representations of knowledge which are "natural" for the agent and leaves the abstraction which make communication easier as a separate problem.

It is awesome how humans use their five senses and limbs to accomplish an enormous range of interactions with physical objects. Our point here is that it is not the number of sense types or limbs that is crucial but rather how they are used to create an intricate web of interactions. Clearly, many connections between sensing and acting require deep analyses and deliberation. However, there are others that are immediate, rapid, and tied to bodily functions, i.e., reflexes. In fact there is a continuum of level of cognition between sensing and acting. Our second *principle of physical interaction and manipulation* is the need for discovering and modeling various pathways between sensing and acting. Our architecture addresses this continuum by providing concurrent conscious and unconscious processes.

In our implementation of GLAIR-based agents, each actuator, each sensor, each reflex, each behavior at the unconscious level, and the conscious reasoning is made into a process in a multiprocessing environment. With this scheme we promote concurrent execution of sensor data processing and actuator control. All sensors and actuators are active at all times. We also make all reflexes and behaviors compete concurrently. Reflexes are given priority over behaviors but certain behaviors can be programmed to override reflexes. Similarly, consciously reasoned actions can be made to override unconscious behaviors and reflexes.

In addition to concurrent sensing and acting, for this principle we also suggest that control and coordination of actions take place at multiple levels within an agent. This is in contrast to central action arbitration. In GLAIR, this action arbitration takes place independently at each of the three levels.

For many real-time tasks, the close coupling of sensing to acting, i.e., immediacy of sensing to acting is a key element in success of physical interactions. Our third *principle of physical interaction and manipulation* is that we need to model this close coupling between sensing and acting. One form of this close coupling is reflexes. A more general form of this type of coupling is antecedent driven actions. The agent experiences a condition and in response produces an action. The agent may use the consequences of actions in deciding to perform actions. This is based the operant model of behavior. Another form of coupling is learned patterns of interaction that are unconscious for the agent. Agents that live in real worlds need to generate timely actions on a routine basis and often without time to reason about actions. This can be considered to be habitual. The agent decides to perform an action solely based on its predecessor action.

An agent's monitoring to accomplish a task requires starting and stopping actuator functions. We believe that this is at too fine a level of control and we ought to start and stop behaviors instead. We ought to treat interactions between sensing and acting in behaviors as closed boxes, learned and encoded by the agent separately. At the level of behaviors, an agent will not have complete control over its actuators. Actuators will asynchronously receive commands from various behaviors and compute a resultant outcome for themselves. Our fourth *principle of physical interaction and manipulation* is that we need to develop feedback control loops, one per behavior and all concurrently active.[3]. With this principle we are promoting equality among behaviors. As in behavior-based AI [Bro90], perceptions with actions are paired up in modules. However, unlike Brooks's subsumption, no behavior inhibits another. Signals in each behavior are sensation. Outputs will be the commands behavior sends to the actuator. The comparator will compare the difference between the reference and the actual signal. In the decision part of the loop, we use the discrepancies detected by the comparator to make decisions about actuator commands. No reference is made to other behaviors. In our architecture, we provide a language for building a variant of these feedback control behaviors. All coordination between behaviors is achieved at the KL. Since actuators have processes independent of behavior processes, each feedback control loop sends messages to actuators asynchronously. Behaviors are not in direct control

[3] A feedback control loop consists of a reference signal, a comparator of sensed signal against the reference signal, and a decision mechanism that generates and sends commands to the actuators

of the actuators. The result of an actuator's performance depends on the combination of commands received from behaviors.

We have developed agents that play air-battle games and learn to fly and out-maneuver other agents [HCBS93]. We are currently developing physical and simulated mobile robot agents using our architecture. Instead of describing our previous work, in the remainder of this paper we sketch broad outline of an agent performing two tasks of inserting a video tape and locking a door with a key. Our point with this very simple example is firstly to point out the principles in GLAIR and secondly that very few sensors and actuators can be used to accomplish large number of tasks.

4 An Example

Let's have a robotic agent with one cylindrically shaped limb (arm), two fingers at the tip, and force/tactile sensing at the tip of a finger. We will use three actuators (i.e., motors), one for moving the arm axially, second one to open/close fingers, and the third one to turn the fingers at the wrist. External sensations are force at wrist: low-resistance, high-resistance; tactile: key-imprint, video-imprint, no-imprint (when fingers are closed). Internal sensations are axially-moving-out (i.e., reaching), axially-moving-in (i.e., retracting), turning-clockwise, turning-counterclockwise, fingers-close, fingers-open. Actuators can be commanded separately. Move-arm-forward, move-arm-backward, stop-arm, turn-clockwise, turn-counterclockwise, stop-wrist, open fingers, close-fingers, stop-fingers are our primitive actuator commands.

We are making many simplifying assumptions. For instance when placing the video tape in the VCR or key in the key hole, if the arm is not perfectly positioned, we will be "wedging" instead of "inserting". However, these assumptions were in light of the physiology. Since we did not provide for the arm to move any other way but axially, we couldn't expect corrective actions. We also assume that actuator speeds are slow and the sensation time is negligible.

From the primitive sensations, we can infer more abstract sensations like slippage. For our robotic agent, we define an object slipping if while being held its imprint changes substantially during motion. With this we have an abstract sensation of slippage. We define arm-stop as an abstract actuator command to stop all three motors.

Figure 2 shows an example of the association between primitive sensations and more abstract ones in the KL. Primitive sensations are filtered through lower levels of GLAIR and when appropriate concepts are identified, if there are corresponding terms in other levels, they are updated.

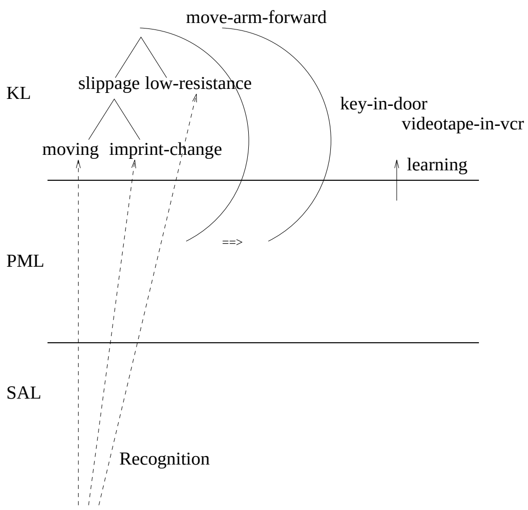

Figure 2: An example of GLAIR knowledge base. This knowledge is in transit between KL and PML.

Now for mapping sensing to actuators, we write a bunch of behavior feedback loops in PML. Antecedent parts of these behaviors are nondecomposable at the PML. While (and (not slip) low-resistance) move-arm-forward, When slip arm-stop, When (and (or video-imprint key-imprint) high-resistance (not slip)) open-fingers, When (and fingers-open high-resistance) move-arm-backward, When (and fingers-open moving-moving-in) close-fingers, When (and low-resistance fingers-closed no-imprint) move-arm-forward, When (and high-resistance no-imprint) arm-stop, when (and key-imprint high-resistance (not slip)) turn-clockwise, when (and key-imprint high-resistance slip) arm-stop.

Since these behavior feedback loops exist in the PML and not in KL, it is as if the agent has somehow acquired the skills involved in these tasks but can't yet explain how it uses them. The agent can learn either every step of the feedback loop or just when to perform each task at KL. Learning here means migrating knowledge form PML to KL as shown in Figure 2. In this example, the agent can observe the frequency of performing inserting a video tape and locking a door with a key. Pattern of actions that occurs frequently in these tasks will be noticed by the agent and made into a concept.

5 Conclusion

We have outlined several key principles we find useful for physical interaction and manipulation. Matching physiology to tasks, a model for the continuum of level of cognition involved between sensing and acting, many sensing and acting behaviors need to be coupled closely, and that behavioral feedback control is needed are our proposed principles.

References

[AC87] Philip E. Agre and David Chapman. Pengi: An implementation of a theory of activity. In *Proceedings of AAAI-87, Seattle, Wa.*, pages 268–272, July 1987.

[Bro85] Rodney Brooks. A robust layered control system for a mobile robot. Technical Report 864, MIT AI Labs, MIT, 1985.

[Bro90] Rodney A. Brooks. Elephants don't play chess. *Robotics and Autonomous Systems*, 6:3–15, 1990.

[HCBS93] Henry Hexmoor, Guido Caicedo, Frank Bidwell, and Stuart Shapiro. Air battle simulation: An agent with conscious and unconscious layers. In *University of Buffalo Graduate Conference in Computer Science-93 (TR93-14)*. Dept. of Computer Science, SUNY at Buffalo, New York, 1993.

[HLCS93] Henry Hexmoor, Johan Lammens, Guido Caicedo, and Stuart C. Shapiro. Behavior Based AI, Cognitive Processes, and Emergent Behaviors in Autonomous Agents. In *Proceedings of AI in Engineering*, 1993.

[HLS92] Henry Hexmoor, Joe Lammens, and Stuart Shapiro. An autonomous agent architecture for integrating perception and acting with grounded, embodied symbolic reasoning. Technical Report CS-92-21, Dept. of Computer Science, SUNY at Buffalo, New York, 1992.

[HLS93] Henry Hexmoor, Johan Lammens, and Stuart C. Shapiro. Embodiment in GLAIR: A Grounded Layered Architecture with Integrated Reasoning. In *Florida AI Research Symposium*, pages 325–329, 1993. also available as SUNYAB CS Dept. TR93-10.

[HN92] Henry Hexmoor and Donald Nute. Methods for deciding *what to do next* and learning. Technical Report AI-1992-01, AI Programs, The University of Georgia, Athens, Georgia, 1992. Also available from SUNY at Buffalo, CS Department TR-92-23.

[HS93] Henry Hexmoor and Stuart C. Shapiro. Examining the expert reveals expertise. In *3rd International Workshop on Human and Machine Cognition: Expertise in Context*, 1993. to appear.

[LHS93] Johan Lammens, Henry Hexmoor, and Stuart C. Shapiro. Of elephants and men. In *Proceedings of NATO-ASI*. Springer Verlag, 1993. to appear.

[Win75] Patrick Henry Winston. Learning structural descriptions from examples. In Ronald J. Brachman and Hector J. Levesque, editors, *Readings in Knowledge Representation*, pages 141–168. Morgan Kaufmann, San Mateo CA, 1985 (orig. 1975).