March 26, 2006

Proposal to Center for Advanced Robotics Steering Group

From Stan Franklin and the “Conscious” Software Research Group (CSRG)

LIDA-AV: Cognitive Robotics for Autonomous Vehicles

Introduction to Cognitive Robotics

Traditional robotics has been mostly concerned with designing robots that know how to perform certain tasks. Cognitive robotics adds the dimension of deciding what tasks are to be performed and when. In addition, a cognitive robotics agent should serve to connect and integrate the various subsystems of the robot, thereby providing a complete control structure, a “mind.” This cognitive agent control structure would encompass the robot’s sensory system, its perception, various memory systems, attention, action selection, and its effectors. The cognitive agent control structure would be responsible for ALL of the various tasks assigned to the robot.

A successful Autonomous Vehicle (AV) will need to operate in a very dynamic, confusing environment. Given a scenario where an AV's navigational system makes the suggestion to move forward towards its next waypoint, while simultaneously the obstacle avoidance system tells it to turn right because of an object in its path, and a human controller tells it to turn left because of another priority objective: a LIDA-like cognitive system to orchestrate and manage this decision making process, seems a necessity.

Let’s refer to a LIDA-based cognitive robotics agent control structure for an autonomous vehicle as LIDA-AV. In this case LIDA-AV would integrate, and be responsible for all the various subsystems, robotic control (effectors), vision (sensing and perceiving), localizaton, navigation, obstacle handling and, of course, cognitive robotics (action selection). Each of these various subsystems would be integrated into the LIDA architecture in a way appropriate to its role in the overall operation of the AV.

This LIDA architecture is best viewed as being implemented by a continual and frequent iteration of its cognitive cycle–sample the environment, process the information, select an action, and act on the environment. These cognitive cycles are conveniently divided into nine steps: 1) sensing, 2) perceiving, 3) cueing episodic memory, 4) selective attention, 5) broadcast, 6) recruitment of internal resources, 7) instantiation of action schemes, 8) action selection, 9) action taken. Higher-level cognitive processes require multiple cognitive cycles for their operation.

Robot control is associated with control of effectors, with how to actually carry out selected actions. Routine tasks corresponds to sensory-motor automatisms in biological systems, and has been the major emphasis of traditional robotics. It comprises step 9 of the LIDA cognitive cycle, the action taken step.

Vision in the AV is associated with steps 1 and 2 of the cognitive cycle, sensing and perceiving. Implemented through perceptual associative memory (the slipnet) in the LIDA architecture, its tasks include the identification of individuals (say an individual aircraft), the recognition of objects (for example, a person or a container), and understanding of situations (for example, another AV crossing its path).

Localization, clearly needed under the assumption of traditional AI planning, might still be useful in the context of the continual, frequent action selection of the LIDA cognitive cycle. If so, it may require GPS as an additional sense, and utilize perception, perhaps over several cycles to recognize landmarks. If needed, the environmental mapping can be built into perceptual associative memory (the slipnet).

Navigation algorithms are to be implemented in the LIDA architecture as action schemes (step 7), and navigation results from the continual cyclic action selection (step 8). Navigation, and the entire LIDA-AV operation must be in real-time.

Obstacle handling (recognition and avoidance) would be implemented as multi-cyclic action selection in LIDA-AV.

Subsystem to be addressed

Our primary subsystem to be addressed will be **cognitive robotics** though, as described in the preceding section, we expect to interact closely with each of the other research groups and their subsystems. We envision each of the other groups contributing their knowledge and research results to the design and implementation of LIDA-AV, which will provide the overall control structure for the autonomous vehicles. More detail on this control structure and its relationship to the other subsystem research groups will be found in the *Open questions to be addressed* section below. LIDA-AV should be able to be efficiently developed within, and utilizing, the envisioned overall blackboard system architecture

Summary of known techniques

The LIDA architecture consists of the IDA architecture together with several forms of learning: perceptual learning, episodic learning, and procedural learning (Franklin 2005). IDA is derived from a working software agent designed for, and funded by, the US Navy for personnel work (McCauley and Franklin 2002).

The IDA architecture includes modules for perception (Zhang, et al. 1998), various types of memory (Anwar and Franklin. 2003, Franklin et al 2005), "consciousness" (Bogner, Ramamurthy and Franklin. 2000), action selection (Negatu and Franklin. 2002), constraint satisfaction (Kelemen, Liang, and Franklin. 2002), deliberation (Franklin 2000), and volition (Franklin 2000). The mechanisms of these modules are derived from several different "new AI" sources (Hofstadter and Mitchell. 1994, Jackson 1987, Kanerva 1988, Drescher 1991, Maes 1989). All these constitute known techniques to be employed in the proposed research.

Open questions to be addressed

1. *Tuning internal LIDA parameters:* The LIDA architecture employs a number of internal parameters that need to be tuned correctly for efficient operation. The IDA internal parameters have been tuned for the personnel work domain of the Navy. Some will need to be retuned for the domain of this proposal. All of the added LIDA parameters must be tuned.
2. *Relating vision to perceptual associative memory (the slipnet):* Here we must produce primitive feature detectors to interface LIDA's perceptual module with what's given it by the vision subsystem research group. This effort will be facilitated by our as yet unpublished joint research with David Friedlander of SET Corporation (Washington, DC).
3. *The need for localization and a built-in environmental map:* If we were to employ traditional AI planning for action selection, localization and an environmental map would clearly be required. They may well be still desirable as part of a LIDA-AV control structure. If so, localization will require an additional GPS sensor that must be interfaced with the LIDA perception module, as well as action schemes implementing landmark recognition algorithms. In addition, we must solve the problem of integrating the environmental map into the LIDA architecture. All this will require cooperation with the localization subsystem subgroup.
4. *Navigation:* Algorithms for navigation must be implemented in the LIDA architecture as action schemes. These algorithms may well include perception for localization, and added complexity of dealing with trailers. This implementation of the algorithms will be accomplished in collaboration with the navigation subsystem research group using their results.
5. *Obstacle handling:* Obstacles must first be recognized and then avoided, or otherwise handled. Recognition is to be accomplished by LIDA's perceptual module, with help from the vision subsystem research group.

The handling of obstacles will require algorithms implemented as action schemes, together with behavior streams that operate over multiple cycles. These must be designed and implemented in collaboration with the obstacle handling subsystem research group.

6. *Sensory-motor automatisms (SMA):* Traditional robotics has emphasized designing robots that know how to do things. The ability to perform tasks is critical for any robot. When at the end of each cognitive cycle an action (instantiated action scheme) is selected by the LIDA-AV control agent, that action scheme must call an appropriate SMA to actually perform the action. The algorithms for these SMA's will be mostly provided by other subsystem research groups. However, the mapping between instantiated action schemes and SMA's will be an open problem for our group. We are hopeful that some of these SMA's, and their mapping, can be learned, as is often the case in humans.
7. *Communication with humans (dispatchers) and with other autonomous AV's:* The LIDA-AV control agent may have to receive and understand commands from a human dispatcher. It may also have to coordinate activities with other autonomous AV's. Thus, communication may become a necessity, and thus an open problem. Communications with humans in English is a problem largely solved in the IDA Navy personnel agent. There are inter-agent communication languages (for example FIPA) that can be used. The coordination problem will be more difficult than the communication problem, but should be solvable in a straightforward way within the LIDA architecture.
8. *Integrating and testing:* The various pieces of the entire LIDA-AV system must be integrated and tested.

Though not all of these open questions can be fully addressed during the eighteen-month time frame of this proposal, some initial headway can be expected in each of them. Do note that the role of the cognitive robotics subsystem research group, as envisioned here, can be expected to be substantial from the beginning of the project, and to increase more than linearly as the project progresses.

Resources needed

The proposed research by the cognitive robotics subsystem research group will be conducted by the "Conscious" Software Research Group <http://csrg.cs.memphis.edu/csrg/index.html>, led by Stan Franklin, as part of its ongoing development of the LIDA model. Existing CSRG team members, in particular Franklin, Colman, Friedlander, and D'Mello <http://csrg.cs.memphis.edu/csrg/html/people.html>, will act as part-time, no-cost consultants to the LIDA-AV effort.

New team members will be added with academic full-time responsibility to the LIDA-AV research and development. One new graduate research assistant (RA) will be hired, along with an undergraduate RA. The latter, a senior computer science major who intends to go on to graduate work, will provide continuity to the project when the graduate RA is graduated. It will be important to bring these RA's on board as soon as possible, since the LIDA architecture is quite complex and bringing them up to speed will take a little time.

These new team members will use existing workstations for their development work. However, one additional high-end workstation will be needed to house and experiment with the LIDA-AV agent. The following budget has line items for the two RA's and this workstation.

Budget

Salaries and Wages			
A.	Principal Investigator	Stan Franklin	
	No funds requested		
B.	Graduate Research Assistant		
	1. Acad year		9,000
	2. Summer		7,000
C.	Undergraduate assistant @ $5.15		7,004
		Subtotal I	**23,004**
Fringe Benefits			
	@ 7.65% of IC		536
		Subtotal II	**536**
Travel			
		Subtotal III	**-**
Operating Expenses			
	Workstation		3,000
	Tuition/fees for graduate assistant		6,717
		Subtotal IV	**9,717**
Equipment			
		Subtotal V	**-**
		Total Direct Costs	**33,257**
Facilities and Administration Costs			
		Subtotal VI	**-**
		Total Project Costs	**33,257**

Budget justification

The graduate RA salary is computed at $1,000, which is less than the $1,125 the Computer Science Department pays its master's level graduate assistants. We hope that the guaranteed summer employment, plus the chance for interesting work leading to a master's thesis, will be enough to enable us to recruit a suitable RA. The undergraduate RA salary is computed at minimum wage. Here we're relying on interesting work and credit for independent study to allow us to recruit an appropriate computer science undergraduate student.

If necessary, the purchase of the high-end workstation to accommodate the full LIDA-AV agent can be postponed until the second year. The system server or an existing workstation would likely be sufficient for developing the initial integrated versions.

Timeline for R&D activities

A timeline for the cognitive robotics subsystem research group is much more difficult to generate than for other subsystem research groups for two reasons. First, the objectives of this research are open-ended with performance criteria not yet specified. Second, the research of this subsystem research group is so dependent on that of all the other subsystem research groups, that our time expectations cannot be estimated without seeing theirs. Thus we will not include a timeline.

Some qualitative predictions of how our research time will be spent can be made. The open questions outlined above can also be viewed as tasks. During the early stages of this research, much attention will be devoted to task 1, the tuning of parameters, without which nothing else can be accomplished. However, this tuning can only be accomplished along with progress on several of the other tasks (2, 4, 6). Thus, much emphasis must be placed on visual perception, on algorithms for navigation, and on the mapping between selected actions and SMA's. The latter will require close coordination with each of the other subsystem research groups. Research and development on all of these can, even must, occur in parallel. The same is true of the other listed tasks as well.

By August 2007 we would expect to have completed the tuning of parameters, to have a running perceptual module related to vision, and to have a running integrated LIDA-AV agent, with some stubs for environmental inputs and actions, with which to experiment. Much progress will have been made on the other tasks, with the exception of communication, which will be postponed until later in the project. Communication is essentially a solved problem, with the AV domain being much simpler than the previous Navy personnel domain. It can be implemented quickly when needed.

References

Anwar, A., and S. Franklin. 2003. Sparse Distributed Memory for "Conscious"

Software Agents. *Cognitive System s Research* 4:339-354.
Bogner, M., U. Ramamurthy, and S. Franklin. 2000. Consciousness" and Conceptual Learning in a Socially Situated Agent. In Human Cognition and Social Agent Technology, ed. K. Dautenhahn. Amsterdam: John Benjamins.
Drescher, G. L. 1991. *Made-Up Minds: A Constructivist Approach to Artificial Intelligence*. Cambridge, MA: MIT Press.
Franklin, S. 2000. Deliberation and Voluntary Action in 'Conscious' Software Agents. *Neural Network World* 10:505-521.
Franklin, S. 2005. Cognitive Robots: Perceptual associative memory and learning. In Proceedings of the 14th Annual International Workshop on Robot and Human Interactive Communication (RO-MAN 2005).
Franklin, S., B. J. Baars, U. Ramamurthy, and M. Ventura. 2005. The Role of Consciousness in Memory. *Brains, Minds and Media* 1:1-38, pdf.
Hofstadter, D. R., and M. Mitchell. 1994. The Copycat Project: A model of mental fluidity and analogy-making. In *Advances in connectionist and neural com putation theory, Vol. 2: logical connections*, ed. K. J. Holyoak, and J. A. Barnden. Norwood N.J.: Ablex.
Jackson, J. V. 1987. Idea for a Mind. *Siggart* Newsletter, 181:23-26.
Kanerva, P. 1988. *Sparse Distributed Mem ory*. Cambridge MA: The MIT Press.
Kelemen, A., Y. Liang, and S. Franklin. 2002. A Comparative Study of Different Machine Learning Approaches for Decision Making. In *Recent Advances in Sim ulation, Com putational Methods and Soft Com puting*, ed. E. Mastorakis. Piraeus, Greece: WSEAS Press.
Maes, P. 1989. How to do the right thing. *Connection Science* 1:291-323.
McCauley, L., and S. Franklin. 2002. A Large-Scale Multi-Agent System for Navy Personnel Distribution. *Connection Science* 14:371-385.
Negatu, A., and S. Franklin. 2002. An action selection mechanism for 'conscious' software agents. Cognitive Science Quarterly 2:363-386.
Zhang, Z., S. Franklin, B. Olde, Y. Wan, and A. Graesser. 1998. Natural Language Sensing for Autonomous Agents. In *Proceedings of IEEE International Joint Sym posia on Intellgence System s 98*.