Socially Embedded Learning of the Office-Conversant Mobile Robot *Jijo-2*

Hideki Asoh, Satoru Hayamizu, Isao Hara, Yoichi Motomura
Shotaro Akaho, and Toshihiro Matsui
jijo20etl.go.jp
Electrotechnical Laboratory
1-1-4 Umezono, Tsukuba, Ibaraki 305 Japan

Abstract

This paper explores a newly developing direction of machine learning called "socially embedded learning". In this research we have been building an office-conversant mobile robot which autonomously moves around in an office environment, actively gathers information through close interaction with this environment including sensing multi-modal data and making dialog with people in the office, and acquires knowledge about the environment with which it ultimately becomes conversant. Here our major concerns are in how the close interaction between the learning system and its social environment can help or accelerate the systems learning process, and what kinds of prepared mechanisms are necessary for the emergence of such interactions. The office-conversant robot is a platform on which we implement our ideas and test their feasibility in a real-world setting. An overview of the system is given and two examples of implemented ideas, i.e. dialog-based map acquisition and route acquisition by following, are described in detail.

1 Introduction - Socially Embedded Learning -

"Why can children and young animals learn complicated things so efficiently ?" "Why can't machine learning systems take off and overcome hard coded systems ?" In the last thirty or forty years the field of machine learning has developed many effective representation schemes and learning algorithms. Although these algorithms have achieved a wide variety of successes and brought about the systems which can recognize speech, human faces, medical diseases, etc., the learning capability of artificial systems is still far from that of humans. In many explanations which have been proposed to answer the questions, one of most likely is that the machine learning systems lack social relationships with the environment including teachers or other learning systems.

In the design of traditional learning systems, learning algorithms are implemented only in the learning systems. No special consideration is given to the teacher or the environment surrounding the system. The learning systems are fed with training data which is prepared by users, and learn simple functional relationships hidden in the data. In other words, only very narrow information channel is maintained between the systems and the teachers. The learning systems are rather isolated from the information-rich environment surrounding them.

On the other hand, human children apparently have very dense interaction with their surroundings, especially with their parents. Recent research in developmental psychology has been revealing that there is a huge number of innate mechanisms or tricks implemented not only by the learners (babies, children) but also by the teachers (parents, adults), which often function unconsciously and support the children's learning process by maintaining the close coupling between the learners and their environment.

For example, as for language development of newborn babies, Masataka[1992] analyzed the interaction between 3-and 4-month-old babies and their mothers, and reported that mothers unconsciously respond to babies' cooing (the most early stage of speech) by imitating these sounds. Babies associated this imitating response with a comfortable, safe feeling, which motivates them to imitate their mothers' sounds. This circle of imitating sound works as a very good training process for making vocal sound. Here the most important stimulus is the mother's appropriate response at the appropriate time.

We often observe that one of the most important information sources for learning of not only new-born babies but also elder children is the information given from surrounding adults "on the job". Giving appropriate instructions or advices at the appropriate time is the most powerful way to teach something. While teaching, adults estimate children's focus of attention and choose the most effective time for giving information.

This kind of closely coupled interaction with the environment significantly supports and accelerates children's learning. We call this aspect of the learning process "*Social Embeddedness of Learning.* Here the most important point is that the mechanisms are implemented in *both learners and teachers.* Not only learning systems but also the people in the system's environment actively participate in the learning processes and play important roles.

In the field of Artificial Intelligence, the importance of such close coupling with environment has been emphasized recently [Agre and Chapman, 1987; Brooks, 1991]. Terms such as "reactiveness", "situatedness'\ or "embeddedness in the environment" are used to express such understanding. Many trials to build situated intelligent systems have been undertaken and reported [Kaelbling, 1987]. Although the importance of such features for learning (one of the most important capabilities of intelligent systems) has been advocated recently [Kaelbling, 1993] it still has not investigated thoroughly.

Based on these considerations, we have launched a research project for fostering a learning system called "office-conversant robot *Jijo-2* [1]" which has ample interaction with its environment and plenty of assistance from nearby humans [Matsui *et al.*, 1995]. The office-conversant robot is a mobile robot which autonomously walks around in an office environment, actively gathers information through close interaction with this environment including sensing multi-modal data and making dialog with people in the office, and acquires knowledge[4] about the environment with which it ultimately becomes conversant.

Our research interest is in investigating and clarifying how the close interaction between learning systems and teachers can help and accelerate the learning processes. What kind of mechanisms are necessary to make the effect emerge ? In particular, our interest is not in the learning of lower level functions but in learning at higher levels, such as combinatorial symbolic structure in the environment, and how human teachers can assist the systems to learn.

In the following sections we give an overview of the office-conversant robot system, and present two examples of socially embedded learning, i.e. dialog-based map acquisition and route acquisition by following a person.

2 Office-Conversant Robot

We chose indoor autonomous mobile robots as a platform of our research for the following reasons:

- In order to communicate with humans naturally, it is desirable to have a physical body.
- The ability to move around helps data to be collected actively in a real world environment
- Compared with manipulators, the mechanisms of mobile robots are rather simple and easy to control.
- There are several commercial mobile robot bases available even to the novice of robotics.

The conventional application of machine learning in robotics is learning to control complex mechanisms. Learning of dexterous manipulation and smooth navigation are typical research issues. However, in our study, learning of controlling is not significant part. In other words, the office-conversant robot is a robot which survives to learn many things, not learns many things to survive. There are many interesting targets to be learned by office-conversant robots, such as

- the topological and geometric structure of environmental space (map),
- the relation between humans/objects/resources and locations,
- the relation between humans and objects or resources,
- the relation between humans, and
- the ontological structure of the office environment.

These targets are combinatorial symbolic structures in office environments. Learning combinatorial structures is one of the most important challenges for machine learning research.

As the robot becomes conversant with its office environment, it can function as an information server in the office. The existence of such a walking dictionary will facilitates smooth communication between members of the office and supports efficient group projects.

2.1 Hardware Architecture

Almost all hardware components are off-the-shelf (Figure 1). We use Nomad 200, a three-wheel mobile robot base provided by Nomadic Inc. (Stanford, U.S.A.). Our Nomad 200 is equipped with 16 ultra sonic sonar ring sensors, 16 infrared ring proximity sensors, touch sensors in the robot's bumper, odometrie sensors for measuring running distance and steering angle, and a compass. The on-board computer is an IBM/AT with Intel Pentium (180 MHz) controlled by Linux OS and connected to a LAN through radio Ethernet.

We added a small microphone, one CCD color camera, an image capture board, an analog radio transmitter for transmitting the speech signals to the host computer, and the Japanese speech synthesizer "Shaberimbo" (commercial product by NTT Data Co, Ltd.). The host computer is Sparc Station 20 with Super Sparc 60 MHz x4.

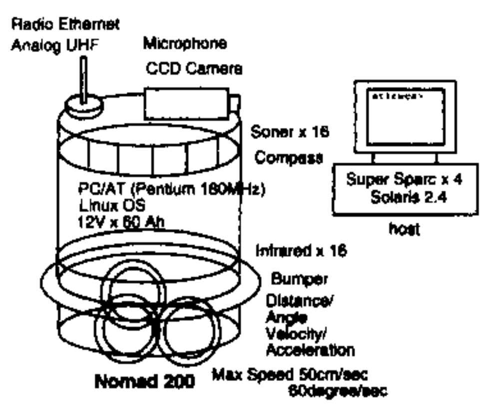

Figure 1: Hardware Architecture

[1]The name *Jijo-2* means "conversance" when pronounced in Japanese

2.2 Software Architecture

The control software is organized in a multi-agent based reactor-integrator architecture [Matsui *et a*/., 1997]. Ten agent modules are formed in a subsumptive architecture, and concurrently running modules communicate with each other in an event-driven manner using UNIX sockets (Figure 2). Some of these modules consist of sub-modules. In each module, several behavior classes are implemented and they are instantiated and invoked by request messages from other modules. With this modular architecture, we can rather easily realize concurrent event processing, which can cope with interrupts caused by exceptional event notification.

As is shown in Figure 2, all modules are divided into two groups. The lower group is devoted to reactive behaviors such as obstacle avoidance, and elemental motions (go straight along a corridor, turn right at a corner, etc.). Image capturing and simple early vision processing (visual tracking with correlation matching), interface with a speech synthesizer, and navigation-related local event (such as detecting an open space etc.) monitor are also included in the group. These modules are implemented on the on-board PC in C.

The upper group provides more deliberative goal-directed behaviors such as scheduling multiple goals, goal-directed route planning, and making simple goal-oriented dialog with humans. These modules are implemented with EusLisp [Matsui and Hara, 1995], which is a dialect of object-oriented Lisp on the host machine. The speech recognizing module is a Hidden Markov Model-based, speaker-independent continuous speech recognition system for Japanese sentences [Itou *et a*/., 1993]. The recognition rate in a previous controlled experiment was approximately 84.2 % of the spontaneous speech of 40 subjects (183 utterances). Because the conditions are not as favorable for our robot, we limited the variety of acceptable speech. Currently the robot's vocabulary is approximately 50 words. The dialog control is very simple using templates of dialog patterns. For example,

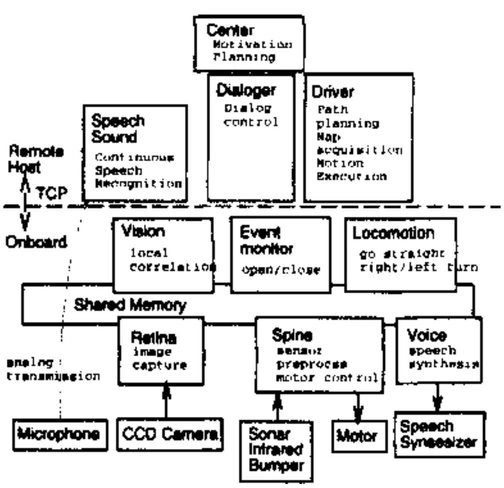

Figure 2: Multi-agent software architecture

a template for asking location is used to extract location names from the speech of humans. Some templates for answering to simple questions from humans are also prepared.

3 Socially Embedded Map Learning

As mentioned above, there are many targets to be learned by the office-conversant robot. The most fundamental one is map learning for efficient navigation in the environment. Although our chief aim is for a variety of knowledge about the office environment, we chose this as the first target and implemented two mechanisms for map learning using intensive human assistances.

As is well known, many representation schemes of maps have already been proposed. They can be roughly divided into two categories: "occupancy grids" [Buhmann, 1995; Elfes, 1992], and finite automat a-based topological maps [Kuipers and Byun, 1988; Mataric, 1992; Tani, 1995], which represent the space as a state transition graph. Each node of this graph usually corresponds to a specific location or landmark, and each edge corresponds to an elemental movement of the system.

We employed the latter one mainly because it does not require accurate global coordinates, and it is more natural for humans to designate the name of a specific place than to designate the x-y coordinate values of the place. To cope with the uncertainty of a real-world setting, we used a partially observable Markov model like that of Simmons and Koenig (1995) and Cassandra *et al* (1996). The map is a collection of conditional probabilities

P(StopLocation, Observations\Start Location, Action).

Each probability corresponds to an edge of the state transition graph of the partially observable Markov model. In our tentative system, *Action* is one of six elemental actions (go straight, turn right/left following free space, turn right/left following wall, and turn around) *Observations* are running distance and accumulated steering angle during execution of *Action.* They are measuied by odometric sensors. The distribution of these values are approximated by Gaussian. *Observation* also includes information acquired from dialog (e.g. location name) when the robot asks a question.

3.1 Dialog-based Map Acquisition

We implemented dialog-based map acquisition, where the system learns a map through making dialog with human trainers [Asoh *et al.,* 1996a, 1996b]. In the task of probabilistic map building, the robots' localization of itself was the most difficult part. Among its sensors, *Jijo-2* can use ultrasonic sonars, odometric sensors, and vision sensor (CCD camera) for locating itself in its environmental space. Although the vision sensor is the most powerful one, it requires a huge computational cost unless we prepare artificial visual landmarks. With sonars and odometric sensors, only very rough localization can be achieved. Accumulating the uncertainty of location seriously prevents the system from efficiently learning the map [Dean *et al.*, 1990].

Here is the first point where we introduced the assistance of humans to facilitate the learning process. Humans in the office can work as very high quality sensors which tell the identifiable name of the location to the robot. The robot always assesses its certainty about the location and if the certainty falls below some threshold, the robot asks a question to a nearby human.

The following is an example of a dialog between a human and *Jijo-2* (actual dialog was in Japanese):

R: Where am I ?

H: You are in front of Dr.Nakashima's Office.

R: Dr.Nakashima's Office ?

(Because the accuracy of speech recognition is not enough high, the robot always confirms human answer. We are planing to use certainty of the recognition to determine whether confirms or not.)

H: Yes

R: Where shall I go ?

H: Please go to Dr.Hara's Office.

R: Dr.Hara's Office ?

H: Yes

R: Sorry, I don't know how to go to Dr.Hara's Office.

H: Ok. Please go straight.

R: Go straight ?

H: Yes

R: I'll go straight.

(The robot goes straight till an end-of-action condition is satisfied. Normally the end-of-action condition is detecting a doorway or crossing.)

R: Where am I ?

H: You are in front of Dr.Matsui's Office.

(The robot goes straight but suddenly stops because of detecting a false end-of-action condition)

R: Where am I ?

H: You are on your way.

R: On my way ?

H: Yes.

R: Where shall I go ?

At the beginning of learning, the system knew nothing about the topological structure of the environmental space. It can only detect close space (corridor between long walls) and open space (doorway and crossing) using sonar. Repeating the execution of a command such as "go straight down the corridor" and asking the question "Where am I ?", it gradually acquired its probabilistic map. Even after the learning proceeded to some extent, while executing ordered navigation such as "Go to Dr. Matsui's office", the system always assessed its position and asked question, "Where am I ?", at appropriate times.

In this setting of map learning, the function of the information coming from dialog is not only to reduce the location uncertainty, but also to help the system to decide the time to learn and the time to add a new location node to the map. In the above dialog, by the human's answer, "You are on your way," the robot could understand that the end-of-action condition was satisfied incorrectly. In this situation the system did not use the data acquired in the previous movement,.

3.2 Route Acquisition by Following

The second point where we introduced human assistances was in teaching a long path to a goal. Although the above dialog-based method worked well, a user, who was also a teacher, had to command each elemental motion (go straight, turn right etc.) at the early stage of learning because the robot did not know anything about its environment. This was rather tedious. A solution is to give the robot more powerful dialog capability which could accept compound commands like "first go straight, then turn right, then....".

Here we introduced another assistance of humans. Instead of commanding a composed path through spoken dialog, a teacher simply commands the robot "Follow me.". The robot follows the teacher with visual tracking capability and learns a path to a goal during the guided tour (see Figure 3).

To acquire the path, that is, a sequence of elemental actions, the robot should segment the entire path to the goal and recognize each segment as an elemental motion.

Figure 3: Jijo-2 following a person

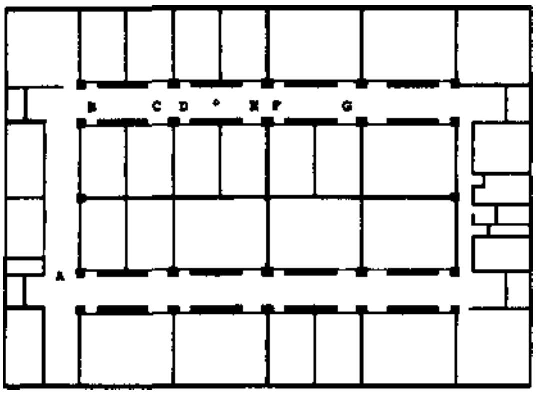

Figure 4: Geographic map of the ETL E-building

Here the self-motion recognition capability plays an important role. We employed the Dynamic Programming matching method to recognize each motion trajectory element [Asoh *et a*l., 1997].

3.3 Experimental Results

We evaluated the effectiveness of the methods in the real environment. The site of the experiment was a part of a floor of our laboratory building (Figure 4).

In an experiment the system executed 52 trial runs. Each run was a trip from their current place to an arbitrary selected place (denoted by the character A, B,... in Figure 4). The longest run was about 20 m. Ultimately the system succeeded in acquiring a topological map with 14 state nodes [Asoh *et* al., 1996a; Asoh *et* al., 1997]. Average moving speed was about 30 cm/sec.

4 Related Work

The book of Leslie Kaelbling [1993] which popularized the concept of "learning in embedded systems", is mainly concerned with learning state-action pairs using reinforcement learning. The interaction channels between learning systems and environments are sensing to recognize the current state and reward for the action.

Several researchers proposed a learning method in which a teacher observes learner's performance and provides appropriate advices [Clouse and Utgoff, 1992; Gordon and Subramanian, 1994; Maclin and Shavlik, 1996]. Doringo and Colombetti [1994] considered the interdependence between the environment, the learning agent, and the trainer, and applied reinforcement learning not only to learning agent but also to trainer program.

Recently in the field of robot learning, "learning from observation" and "learning by imitation" are being considered as promising learning schemes [Demiris 1994; Ikeuchi and Suehiro, 1994; Kawato *et* al., 1994; Kuniyoshi *et ai,* 1994]. These are also considered as an approach of utilizing wider communication channels between the learning systems and human teachers. Here the assisting information is mainly visual. Our second experiment of route acquisition by following can be viewed as a kind of learning by imitation.

One of the pioneering works of introducing dialog as a communication channel with learning systems was the well-known intelligent robot "SHAKEY" developed by Nilsson [1965]. A natural language understanding system "SHRDLU" by Winograd [1972] also had some learning capability. Recently, Mark Torrance [1994] has developed a mobile robot which learns a map through natural language dialog typed in from a keyboard.

The office-conversant robot presented in this paper, which was deeply inspired by these earlier projects, tries to integrate these ideas into a platform to evaluate their effectiveness in acquiring a higher level structural knowledge about the environment.

5 Conclusion and Future Work

We have described the architecture and functions of an office-conversant robot *Jijo-2,* a platform for the research toward socially embedded learning. The results have convinced us that introducing close interaction with the environment, especially human assistance into learning processes, is very effective in enlarging the scope and applicability of machine learning. We would like to continue our explorations in this direction and achieve learning of more functionalities.

Tentative status is far from satisfiable one. We should explore more variety of interaction between robot and its environment. The most urgent problems we currently face are in visual processing and making dialog. For visual processing, we plan to give the robot the capability of detecting humans in the environment and recognizing these humans using an active vision system. This capability is necessary to enlarge the communication channel between the robot and humans. In order to widen the content of dialog, we must implement semantic analysis of spoken sentences and reason with semantic representation. We also plan to utilize the framework of socially embedded learning for realizing these capabilities themselves. Scaling up the experiment is also important in evaluating the system's performance.

Acknowledgement

This work is supported by the Real-World Computing Program organized by the Ministry of International Trade and Industry of the Japanese Government.

References

[Agre and Chapman, 1987] P. Agre and D. Chapman. Pengi: An implementation of a theory of activity. In "Proceedings of the 1987 National Conference on Artificial Intelligence, 268-272, 1987.

[Asoh *ei* al., 1996a] H. Asoh, Y. Motomura, T. Matsui, S. Hayamizu, and I. Hara. Acquiring probabilistic map with dialog-based learning. In *Proceedings of ROBOLEARN-96,* 11-18, 1996.

[Asoh *et a*l., 1996b] If. Asoh, Y. Motomura, T. Matsui, S. Hayamizu, and I. Hara. Combining probabilistic map and dialog for robust life-long office navigation.

In *Proceedings of the 1996 1EEE/RSJ International Conference on Intelligent Robots and Systems,* 807-812, 1996.

[Asoh *et ai,* 1997] H. Asoh, I. Hara, and T. Matsui. Self-motion recognition of mobile robot and its application to map acquisition. Technical Report, Electrotechnical Laboratory, 1997.

[Bellman, 1957] R. Bellman. *Dynamic Programming,* Princeton University Press, 1957.

[Brooks, 1991] R. A. Brooks. Intelligence without reason. In *Proceedings of UCAI-91,* 1991.

[Buhmann, 1995] J. Buhmann. The mobile robot RHINO. *AI MAGAZINE,* Summer 1995, 31-38, 1995.

[Cassandra *et ai,* 1996] A. R. Cassandra, L. P. Kaelbling, and J. A. Kurien. Acting under uncertainty: discrete bayesian models for mobile-robot navigation. In *Proceedings of the 1996 IEEE/RSJ International Conference on Intelligent Robots and Systems,* 963-972, 1996.

[Clouseand Utgoff, 1992] J.A. Clouse and P.E. UtgofT. A teaching method for reinforcement learning. In *Proceedings of the Ninth Conference on Machine Learning,* 92-101, Morgan Kaufmann, 1992.

[Dean *et al,* 1990] T. Dean, K. Basye, R. Chekaluk, S. Hyun, M. Lejter, and M. Randazza. Coping with uncertainty in a control system for navigation and exploration. In *Proceedings of AAAI-90,* 1010-1015, 1990.

[Demiris, 1994] J. Demiris. Experiments towards learning by imitation. In *Proceedings of AAAI-94,* 1437, 1994.

[Doringo and Colombetti, 1994] M. Doringo and M. Colombetti. Robot shaping: developing autonomous agents through learning, *Artificial Intelligence,* Vol.71, 321-370, 1994.

[Elfes, 1992] A. Elfes. Multi-source spatial data fusion using Bayesian reasoning. In M.A.Arbib and R.C.Gonzales eds., *Data Fusion in Robotics and Machine Intelligence,* 137-163, Academic Press, 1992.

[Gordon and Subramanian, 1994] D. Gordon and D. Subramanian. A multistrategy learning scheme for agent knowledge acquisition. *Inform atica,* Vol.17, 331-364, 1994.

[Ikeuchi and Suehiro, 1994] K. Ikeuchi and T. Suehiro. Towards an assembly plan from observation, Part I: Task recognition with polyhedral objects", *IEEE Transactions on Robotics and Automation,* Vol.10, 1994.

[Itou *et al,* 1993] K. Itou, S. Hayamizu, K. Tanaka, and H. Tanaka. System design, data collection and evaluation of a speech dialogue system. *IE1CE Transactions on Information and Systems,* Vol.E76-D, 121-127, 1993.

[Kawato *et al,* 1994] M. Kawato, F. Gandolfo, H. Gomi, and Y. Wada. Teaching by showing in Kendama based on optimization principle, In Marinaro,M. and Morasso,P.G. (eds.) In *Proceedings of the 1994 International Conference on Artificial Neural Networks,* 601-606, 1994.

[Kaelbling, 1987] L. P. Kaelbling. An architecture for intelligent reactive systems. In *Reasoning About Actions and Plans,* M. Georgeoff and A. Lahsky (eds.), Morgan Kaufmann, 1987.

[Kaelbling, 1993] L. P. Kaelbling. *Learning in Embedded Systems,* The MIT Press, 1993.

[Kuniyoshi *et ai,* 1994] Y. Kuniyoshi, M. Inaba, and H. Inoue. Learning by watching: extracting reusable task knowledge from visual observation of human performance, *IEEE Transactions on Robotics and Automation,* Vol.10, 799-822, 1994.

[Kuipers and Byun, 1988] B.J. Kuipers and Y.-T. Byun. A robust, qualitative method for robot spatial reasoning. In *Proceedings of the 1988 National Conference on Artificial Intelligence,* 774-779, 1988.

[Maclin and Shavlik, 1996] R. Maclin and J.W. Shavlik. Creating advice-taking reinforcement learners. *Machine Learning,* Vol.22, 251-281, 1996.

[Masataka, 1992] N. Masataka. Early ontogeny of vocal behavior of Japanese infants in response to maternal speech, *Child Development,* Vol.63, 1177-1185, 1992.

[Mataric, 1992] M.J. Mataric. Integration of representation into goal-driven behavior-based robot. *IEEE 'Transactions on Robotics and Automation,* 8, 304-312, 1992.

[Matsui *et ai,* 1995] T. Matsui, S. Hayamizu, H. Asoh, I. Hara, and Y. Motomura. Office-Conversant mobile robot project. In *Proceedings of the 1995 Annual Conference of the Robotics Society of Japan,* 1995 (in Japanese).

[Matsui and Hara, 1995] T. Matsui and I. Hara. *EusLisp Reference Manual Ver.8.00,* ETL-TR-95-2, 1995.

[Matsui *et ai,* 1997] T. Matsui, H. Asoh, I. Hara. An event-driven architecture for controlling behaviors of the office conversant mobile robot, *Jijo-2,* In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation* (to appear), 1997.

[Nilsson, 1969] J. Nilsson. A mobile automaton: an application of artificial intelligence techniques. In *Proceedings of the First 1JCAI,* 509-520, 1969.

[Simmons and Koenig, 1995] R. Simmons and S. Koenig. Probabilistic robot navigation in partially observable environments. In *Proceedings of 1.JCA1-95,* 1080-1087, 1995.

[Tani, 1995] J. Tani. Self-organization of symbolic processes through interaction with the physical world. In *Proceedings of IJCAI-95,* 112-118, 1995.

[Torrance, 1994] M.C. Torrance. Natural Communication with Robots. *Master Thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science,* 1994.

[Winograd, 1972] T. Winograd. *Understanding Natural Language,* Academic Press, 1972.

LEARNING

Learning 6: Logic and ILP