Automatic Language Acquisition by an Autonomous Robot

(Invited Paper)

Stephen Levinson, Weiyu Zhu, Danfeng Li, Kevin Squire,
Ruei-sung Lin, Matthew Kleffner, Matthew McClain, and Johnny Lee
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign, Urbana, IL 61801
Email: {*sel, weiyuzhu, dli1, k-squire, rlin1, kleffner, mrmcclai, jjlee5*}@ifp.uiuc.edu

***Abstract*— There is no such thing as a disembodied mind. We posit that cognitive development can only occur through interaction with the physical world. To this end, we are developing a robotic platform for the purpose of studying cognition. We suggest that the central component of cognition is a memory which is primarily associative, one where learning occurs as the correlation of events from diverse inputs. We also posit that human-like cognition requires a well-integrated sensory-motor system, to provide these diverse inputs. As implemented in our robot, this system includes binaural hearing, stereo vision, tactile sense, and basic proprioceptive control. On top of these abilities, we are implementing and studying various models of processing, learning and decision making. Our goal is to produce a robot that will learn to carry out simple tasks in response to natural language requests. The robot's understanding of language will be learned concurrently with its other cognitive abilities. We have already developed a robust system and conducted a number of experiments on the way to this goal, some details of which appear in this paper. This is a first progress report of what we believe will be a long term project with significant implications.**

I. INTRODUCTION

Cognitive development has been studied in various environments–on the playground by the psychologist, under the microscope by the neuro-scientist, and in the armchair by the philosopher. Our study occurs in a robotics lab, where we attempt to embody cognitive models in steel and silicon.

How did we choose this particular habitat? First and foremost, we are scientists and engineers, which immediately suggests forming constructive theories and building things to test them. The particular question we are examining is one of the most fascinating questions that has been asked in the last century: Can machines think?

Alan Turing raised this very question back in 1950 [1]. He introduced the idea of a machine engaging in "pure thought" and communicating to the world via teletype writer. As an answer to The Question, he suggested that when the machine's discourse (via teletype) was indistinguishable from a human's, we could say that the machine was thinking. He goes on at the end of the paper to suggest that initially, machines should perhaps learn to compete with men at some purely intellectual task, such as chess, but then, he suddenly presents an alternative approach for creating machine intelligence:

> "It can also be maintained that it is best to provide the machine with the best sense organs that money can buy, and then teach it to understand and speak English. This process could follow the normal teaching of a child. Things would be pointed out and named, etc." [1]

The artificial intelligence community has largely followed the former proposal. We believe the latter holds more promise.

Our research is based on a few fundamental principles. First, we believe that a mind cannot be disembodied–it must interact with the real world. Second, we posit that memory is primarily associative, and that learning is based on the correlation of information from diverse inputs. Third, in humans and higher animals, these two assertions are fulfilled by a complex sensory-motor system. We posit that such a system is necessary for human-like cognition. Finally, we suggest that these ideas provide the basis for a mind which can learn a semantic representation of reality, upon which higher cognition and all linguistic structure is established.

Our experiments are based on these concepts. We are developing a robotic platform with basic sensory-motor capabilities, including binaural hearing, stereo vision, tactile sense, and basic proprioceptive control. On top of this system, we are implementing various processing and learning models, and studying how they contribute to semantic understanding.

We would like to acknowledge work by others we feel is related to our own. We have found work by Rodney Brooks et. al. on Cog [2] and other projects quite inspiring. We have also followed with interest work by John Weng et. al. [3]–[5], whose ideas and methods are very similar to our own.

The rest of this paper is organized as follows. In the next section, we discuss the cognitive cycle, its instantiation, and give an example of how it might work in our robot. Following that, we describe some of the experiments we have conducted, including how they have furthered our goal of creating a robot able to comprehend language. Finally, we describe our plans for the next stage of development, and offer some conclusions.

II. COGNITIVE CYCLE

We have used the basic cognitive cycle depicted in Figure 1 to guide the design of our robotic system. This simple flow diagram divides cognition into four distinct components: sensory input, long term memory, a central decision maker, and motor output. The divisions suggested by this model are simple, yet we feel they provide the necessary framework for embodied learning. We describe these in more detail below, with discussion on 1) how they relate to human cognition, and 2) implementation of functional equivalents.

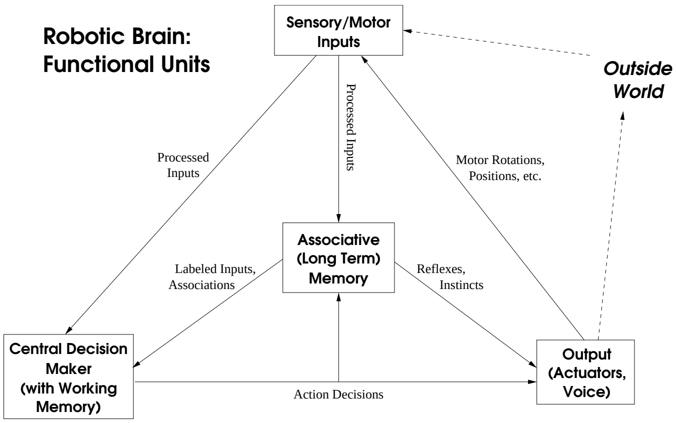

Fig. 1. Cognitive Cycle

A. Cognitive Cycle Components

1) Sensory Input: The necessary start of the cognitive cycle is the gathering of information from the environment through the sensory inputs. In humans, these inputs include the five senses–tactile (touch), gustatory (taste), olfactory (smell), auditory (hearing), and visual (sight). We also perceive information about ourselves, through proprioception (sense of body position and movement) and interoception (internal sensory perception of such things as hunger and body temperature). From these we draw all of our experience, and while we can learn and adapt without one or more of them, these senses are a prerequisite to our cognitive abilities.

For embodied cognition, we would like to implement as many of these senses in our robot as possible. It is relatively easy to find functional equivalents for eyes (cameras) and ears (microphones). Touch sensors, while not nearly as versatile as skin, allow for limited input of tactile sensations. Some proprioceptic sensing can be implemented using feedback sensors on actuators. Olfactory and gustatory sensors are more difficult to include, and we chose to ignore these senses for now. However, research has progressed in the development of artificial skin [6], noses [7], and tongues [8]. Sometime in the not too distant future, researchers will be able to use these organs to allow a robot to perceive an even richer set of sensory inputs. For now, we have chosen to focus on the senses of sight, sound and touch, with minimal simulation of the others (e.g., proprioception) as needed. Some of the particular processing we perform on the sensory inputs is discussed in Section III-A.

2) Long Term Memory:

> Memory is the most important function of the brain; without it life would be a blank. Our knowledge is all based on memory. Every thought, every action, our very conception of personal identity, is based on memory.... Without memory all experience would be useless. (Edridge-Green, 1900) [9]

This quote illustrates why we placed memory at the center of the cognitive cycle in Figure 1. Long-term memory is where almost all learned information is acquired and stored. There are many facets and theories about the learning of long term memories, some of which we have worked into our models. In this area, we have investigated object recognition, speech recognition, vision-based map learning and learning of behavior. Sections III-B and III-C give more details about some of the learning research we have conducted.

One aspect of learning fundamental to our work is the idea that learning and recall occur mostly as the correlation of data from various inputs. Under this assumption, we have started developing a new long-term memory model using a cascade of multi-modal hidden Markov models (HMMs). Some details of this proposal appear in Section IV-A.

3) Central Decision Maker: The central decision maker in the diagram is the goal seeker. It contains a working (scratch) memory, and uses information from the current inputs and long term memory to make decisions and specify actions. It is the glue which ties the other components together.

The workings of the decision making process can be gleaned from the behaviors we want the robot to express. This is discussed more in Section II-C and in future work at the end of this paper. While our present model is simple, we believe that as we progress it will grow in complexity and richness to more accurately reflect corresponding processes in humans.

4) Motor Output: The last stage of the cognitive cycle involves actuating motion in response to the activity of the other components. In particular, we can identify two classes of motion actuation in humans that we need to model:

1) movement in the environment
2) other articulated body movement (e.g., movement of arms and head, speech)

Humans act and react by controlling their skeletal-muscular system. In a robot, this is replaced by a steel body and a system of motorized actuators. To do the most human-like cognitive studies, we would like to work with something as anthropomorphic as possible.

B. Implementation

The previous section gives some minimum specifications and goals we feel are necessary to start building an intelligent robot. Here we will describe our particular implementation.

We chose to work with Arrick Robotics' Trilobot [10]. The robot's anthropomorphic capabilities are rich enough to suit our purposes. In particular, the robot can move freely via wheels, can move its head, and use its arm to manipulate common objects, allowing relatively complex behaviors. For sensory input, the robot has a number of touch and other sensors available. We have also added cameras and microphones for stereo vision and hearing capabilities. An on-board computer collects input from the cameras, microphones, and sensors, and sends control commands to the robot. The computer can also handle limited processing of the data, but a wireless transmitter is available to transmit the data to other workstations, where most processing occurs. We have developed a robust distributed computing system to manage the actual data transmission and processing, with various modules for sensory data processing, learning, decision making, and control. We currently have two such robots, whom we have named *Illy* and *Norbert*, and are continuing to develop and refine them.

C. Desired Behavior

With infants and small children for motivation, we can envision the type of behavior we want our robots to express. We want the robot to learn about its environment. When exploring autonomously, the robot should respond to sensory inputs that it finds interesting–bright colors, loud noises, or other signals that contrast with the background of its environment. Upon discovering something of interest, the robot may try to move closer or look at it from a different position. If it is an object, the robot may try to run into it, pick it up and take it home, or ask a nearby human to name it. The inputs it responds to and the behavioral outputs it expresses may be hardwired (corresponding to innate or instinctual behaviors in humans), or learned. Initially, to ease implementation and testing, much of the behavior should be hardwired.

Another way the robot should learn about its environment is through interaction with a benign teacher. In particular, we wish to explore the robot's acquisition of semantic concepts related to the its environment. The concepts that we expect the robot to be able to learn include such things colors, directions (left, right, forward, etc.), actions, and types of objects. For example, for learning the color red, the robot might be presented with a coke can, a red ball, and a red patch on the wall. These would be named by color while the robot was looking at them. After learning the concept of "red", as well as the concept of "can" and "get", the robot should understand and respond correctly to a request to "get the red can" from an environment containing red, blue, and green cans. The key point here is that the learning of semantic concepts occurs through the correlation of words with various sensory inputs. We feel that this idea forms the basis for the semantic understanding of language.

We have conducted much of our research with these scenarios in mind. The next section describes some of our experiments and how they relate to semantic understanding.

III. Experiments

A. Sensory System

Before semantic understanding can occur, we must consider how information is processed before it reaches the brain. It is well known, for example, that the human ear acts as a spectral filter (see e.g., [11]). As well, a large amount of feature extraction occurs in the retina before the signal even leaves the eye (see e.g., [12]). In our robot, also, we need to process this data before using it for learning or decision making.

For visual inputs, we use mostly standard image processing and computer vision techniques. See [13]–[15] for details. For auditory inputs, in addition to standard techniques, we have developed some processing techniques specifically for anthropomorphic behavior. These include binaural sound source localization and sound understanding.

1) Binaural Sound Source Localization: We have proposed hypothesis-driven approach for binaural sound source localization. The basic idea is to treat binaural processing as a pattern recognition problem instead of solely a signal processing problem. At a given moment in time, a set of features is extracted from the sensory inputs, and a hypothesis is made about a certain pattern of attributes contained in these features. This hypothesis is then checked by the knowledge built in or learned by the model, and subsequently accepted or rejected.

For hypothesis testing, we adopted a Bayes rule based hierarchical decision making framework. Three localization cues–inter-aural time differences (ITDs), inter-aural intensity differences (IIDs) and spectrum–are extracted from the binaural inputs. The location probability is first calculated using the ITDs and the *a priori* information. This probability serves as the *prior* for the second step, where the location probability is refined using the IIDs. The refined probability in turn serves as the *prior* for the last step, and the spectral cues are combined to make the final decision. The reasons for such an approach comes from both psychophysical evidence and engineering consideration [16]. It has be shown in human listener experiments that ITDs are the most robust localization cues, followed by IIDs, and then spectral cues.

Applying this model, 3D localization can be realized only using binaural inputs. We have implemented a robust ITD estimator on the robot, which can successfully guide the robot toward interesting sound sources in many real life environments, including computer labs, seminar rooms, or even the atrium of an office building. The details of the algorithm can be found in [17] and [18].

2) Sound Understanding: Sound understanding refers to the ability to identify a sound using some of its unique characteristics. Neuro-science studies suggest that neurons in the auditory cortex exhibit a variety of complex responses to sound stimuli, and that the pattern of the responses for a given sound might provide its identity.

We have proposed a generic sound understanding model according to this observation, which serves as the "auditory cortex" for the robot. Following the principle of autonomous learning, no labelled training data is required to build the model. At the beginning of learning, the robot listens to various sounds from the environment. A Gaussian mixture model (GMM) is trained after "enough" audio samples have been heard. This process can be thought of as the early development of the "auditory cortex", where each individual "neuron" (individual Gaussian in the GMM) takes in shape.

Once the model is built, a histogram is used to represent a sound, where each bin in the histogram corresponds to an individual Gaussian in the model, and the number of bins corresponds to the number of Gaussians in the GMM. For each input feature vector, the value for each bin is the normalized observation probability of this Gaussian. With a sequence of feature vectors, the value for each bin is the accumulated observation probability. Only one histogram will be built per sound, regardless of the length of the sound. Histogram intersection is applied to measure the similarity between sounds,

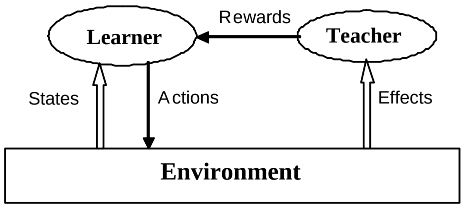

Fig. 2. The architecture of an intelligent action learning system

based on the assumption that larger intersection indicates higher similarity. With this simple sound representation and similarity measurement mechanism, the robot can learn the meaning of and associate different sounds with non-auditory inputs, such as vision. For details see [19].

B. Reinforcement Learning and Navigation

Just as a child must learn to move and interact with the world, our robot needs to learn to move around and interact with its environment. In fact, the ability to interact with the world purposively plays an important role in natural language acquisition. For example, in order to learn the semantic meaning of such complex behaviors as “go over there” or “bring me the can”, it is essential that the robot can execute these behaviors in the real world.

Our aim in these studies is to develop the robot’s ability to acquire general-purpose skills through interaction with the world. Both studies below are based on reinforcement learning. Figure 2 shows the architecture of the reinforcement learning systems we have used. At each time instant, the robot (learner) receives encoded descriptions of the state of the environment from peripheral sensory inputs, and selects actions to perform. The evaluation of the effect of an action is given by a teacher (which may be a person or some measured criteria), and fed back to the agent in the form of positive or negative rewards. According to the reward received, the agent then adjusts its action policy accordingly with the aim to maximize the expected future reward that the robot may receive. The output of the learning is the association between observations and actions, which fits into the role of the long-term memory in the cognitive cycle depicted in Figure 1. Below we describe two experiments which use this methodology.

1) Visual Navigation and Map Learning: In these experiments, we have the robot learn, via reinforcement learning, visual navigation as well as a conceptual map of a maze.

Memory-based algorithms store all previous instances in a large memory. When trying to solve a problem, they simply retrieve from memory instances similar to the current situation, and make a decision based on these retrieved instances. For a memory-based reinforcement learning algorithm, each instance contains information about the current state, the action chosen, and the expected cumulative reward after performing the action. By using this information, the Q-value function can be estimated, and the optimal action policy can then be derived.

We applied this learning algorithm to solve a vision-based robot learning problem. The problem is defined as follows: starting from a randomly chosen position in a maze, the objective of the robot is to learn to move to a goal via the shortest path. Initially, the robot does not have the skill to navigate inside of the maze, so it will act clumsily and might bump into walls from time to time. When the robot bumps into a wall it will receive a negative reward, and when it reaches the goal it will receive a high positive reward. All navigation is done by vision, so the robot must simultaneously learn, via reinforcement, vision-based navigation skills and the cognitive map of the maze.

We decompose the problem using a hierarchical model. At the bottom of the hierarchy, the goal is to learn the navigation skills based on high-dimensional feature vectors extracted from the visual input. We use locally weighted learning [20] to estimate the Q-value function for this subproblem. At the top of the hierarchy, the goal is to learn the maze. It turns out that maze-learning has to be modeled as a partially observable Markov decision process. To solve this subproblem, we use McCallum’s nearest sequence memory [21].

Our experimental result shows that after a number of trials, the robot can acquire both the skill to navigate inside a maze as well as the conceptual map of the maze.

2) PQ-Learning for Navigation: Reinforcement learning methodology has recently received increasing attention for robot skill learning with little or no *a priori* knowledge of the environment [22]–[24]. However, efficient reinforcement learning with a real robot in the real world is still an unsolved problem, especially when the learning state space is large. To address this challenge, we have proposed a novel learning algorithm, called PQ-Learning, which effectively solves the slow convergence problem inherent in the traditional reinforcement learning methodology, by introducing a special value propagation technique in learning (see [25] for details). We have applied this algorithm on *Illy* for goal-oriented navigation learning. The robot was able to successfully learn the optimal action policy to achieve its goal in a few tens of learning episodes, much faster than traditional Q-learning.

C. Language Learning

Here we describe some higher level experiments we have conducted specifically with reference to language understanding and acquisition.

1) Speech and Action Concept Learning: One of our earlier studies looked at the learning of associations between tactile and speech input, in a similar manner to [3]. In this experiment, conducted on an older robot named *Alan*, a benevolent teacher would push on a touch sensor on the robot while speaking a movement command. For example, the teacher might push on a sensor on the back of the robot and say “forward”. A touch on the rear sensor would “push” the robot forward (its wheel would straighten and its motor would start running). After a training period, the robot could, on a speaker dependent basis, be controlled by voice. The most important

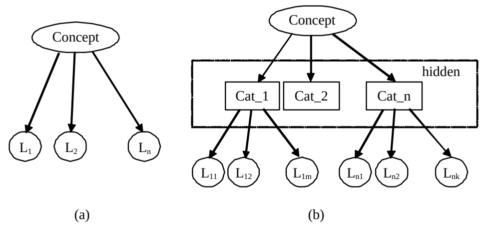

Fig. 3. (a) Supervised learning; (b) Scene Concept Learning

aspect of this research is the idea of associating speech with a built in understanding of direction. See [26] for more details.

2) Scene Concept Learning: The research on automatic scene concept learning is aimed at vision-based language concept understanding by our robot, which is a preliminary step toward the goal of language acquisition. The basic idea and scenario of this study is similar to supervised learning, in that the agent receives labeled scenes of interest during training, and learns the semantic meaning of each label. However, the two learning paradigms are in fact quite different. First, as shown in Figure 3, the task of scene concept learning includes a hidden layer of concept categories such as color, shape, and texture. For each learning example, what the robot receives as the label resides on the leaf nodes only, where the nature of the label, the category type, is unknown to the agent in advance. Second, for scene concept learning, each scene example may receive multiple labels from different concept categories instead of only one as in the supervised learning case. The existence of hidden layers and multi-labeling imposes great challenges on the proposed concept learning tasks.

In this study, we have proposed a feature JPDF (Joint Probability Density Function) based concept representation model for general-purpose visual scene understanding by the robot. Based on this model, we have developed a learning system for small-sized visual scene concept understanding by a robot [27]. The set of the scene concepts consists of 6 color concepts, 3 shape concepts, and 13 object concepts drawn from 15 natural objects. Each object consists of a number of shapes and colors from different view angels. After 40 training episodes, the robot achieved 94% scene concept retrieval accuracy in 40 independent testing examples.

3) Speech Imitation: One topic we have not touched on yet is speech production, which in humans is an essential part of language learning. Children learn speech by mimicking those around them, so it makes sense to use speech imitation as a vehicle for learning to speak in the robot. One of us has developed a robust method for speech imitation, involving extracting phonetic and phonemic features from the sound stream which give an internal representation correlating to the vocal tract shape, while taking into account the resolution of the human ear. The features that are extracted can be reused for speech synthesis or used with features from other modalities for recognition and learning. See [28] for more details.

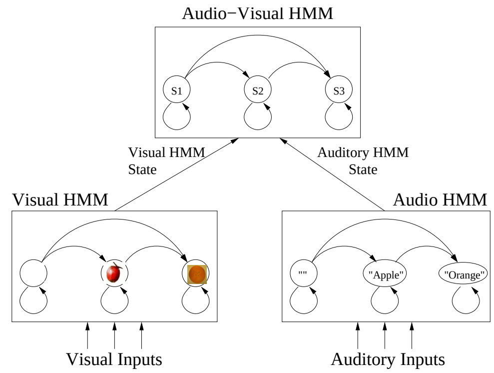

Fig. 4. Cascade of Hidden Markov Models

IV. FUTURE WORK

In additional to the ideas already mentioned, we are currently developing and testing new models related to memory and decision making. Highlights of some of those ideas appear below.

A. Associative Memory Using a Cascade of Hidden Markov Models

The task and goal of our robot's memory is to associate related sensory input from multiple senses. For this scenario, we have a novel proposal for learning from multiple sources. Hidden Markov models (HMMs) have proven quite robust at learning speech models. For our method, we propose having a separate group of HMMs for each modality (i.e., a set of HMMs for auditory inputs, a set of HMMs for visual inputs, etc.), and then using the state outputs of these HMMs as the inputs to the next layer of HMMs. Our approach is demonstrated in Figure 4. While others have suggested and developed hierarchical HMMs (HHMMs) [29], we believe our particular approach to be novel, and moreover capable of modeling some of the deep structure of language.

One important aspect of learning in the robot is that it should be incremental and adaptive. We have had some outside success using the incremental training techniques described in [30], and plan to use these training methods in the HMMs running on our robots. More importantly, we believe this training can occur simultaneously for all models.

B. Central Decision Maker

Through various studies, Gopnik et. al. [31] and others have shown that children learn a great deal through mimicking others. We are planning to implement an imitation mode for the robot in which the robot could learn to perform certain tasks through observation of humans or other robots. Ideally, as the robot matures, we would also like to include the ability for it to use or test some of its basic associations as they develop. A rich environment with a wide variety of stimuli provide the potential for a very complex set of behaviors.

V. CONCLUSION

Our ultimate goal is nothing less than construction and explanation of a mechanical "mind". While the study of mind has an intrinsic theoretical and philosophical component, the matter cannot be resolved by a thought experiment. Some constructive approach, however crude, is required. We consider our project to be a humble but serious beginning to a long-range research program which has significant technological and social implications.

We have proposed three fundamental hypotheses upon which we believe a constructive cognitive theory should rest. First, manipulation of our mental model of reality is primarily accomplished by storing, fetching and comparing memorized associations. Second, this mental model depends critically on a fully integrated sensory-motor periphery. Third, the dominant structure of language is semantics. We have proposed to test these hypotheses through the vehicle of an autonomous intelligent robot, trained in a reinforcement paradigm.

On the basis of these hypotheses, our robot has already acquired a great number of important abilities and behaviors. It can locate and localize sound sources very robustly, and learn how to characterize and understand those sounds. It can learn about its environment visually through reinforcement learning, with a teacher or exploring on its own. It has also begun to learn concepts by recognizing the correlation among speech, objects, tactile inputs, and directions. Underlying all of these behaviors is a robust communications framework allowing the various system components to interact and run concurrently.

We are now at a critical juncture in experiments at which simple behaviors are transformed into complex ones. We believe this complexity will arise from the interaction of numerous simpler components. Although our ultimate goal is still far off, we have made some progress defining the function and interaction of these components, and obtained very encouraging results.

Our work is quite challenging and ambitious, and perhaps controversial. Yet we feel that our experiments are technically feasible and potentially of great practical value if successful. Most importantly, however, in our best scientific and technical judgment, when a mechanical mind is eventually constructed, it will much more closely resemble the ideas expressed herein than the mainstream ideas being pursued so vigorously at the present.

References

[1] A. Turing, "Computing machinery and intelligence," *Mind*, pp. 433–460, 1950.

[2] R. Brooks, C. Breazeal, M. Marjanovic, B. Scassellati, and M. Williamson, "The Cog project: Building a humanoid robot," in *Computation for Metaphors, Analogy and Agents*, ser. Springer Lecture Notes in Artificial Intelligence, C. Nehaniv, Ed. Springer-Verlag, 1998.

[3] Y. Zhang and J. Weng, "Grounded auditory development by a developmental robot," in *Proc. INNS/IEEE International Joint Conference of Neural Networks 2001 (IJCNN 2001)*, Washington DC, July 14-19 2001, pp. 1059–1064.

[4] J. D. Han, S. W. Zeng, K. Y. Tham, M. Badgero, and J. Weng, "Dav: A humanoid robot platform for autonomous mental development," in *Proc. 2nd International Conference on Development and Learning*. Cambridge, MA: IEEE Computer Society Press, June 12-15 2002.

[5] J. Weng, "A theory for mentally developing robots," in *Proc. 2nd International Conference on Development and Learning*. Cambridge, MA: IEEE Computer Society Press, June 12-15 2002.

[6] M. Hakozaki, H. Oasa, and H. Shinoda, "Telemetric robot skin," in *Proceedings of the 1999 IEEE International Conference on Robotics and Automation*, Detroit, Michigan, May 1999.

[7] A. Loutfi, S.Coradeschi, T. Duckett, and P.Wide, "Odor source identification by grounding linguistic descriptions in an artificial nose," in *Proc. SPIE Conference on Sensor Fusion: Architectures, Algorithms and Applications V*, vol. 4385, Orlando, Florida, 2001, pp. 273–282.

[8] S. Savoy et. al., "Solution-based analysis of multiple analytes by a sensor array: toward the development of an electronic tongue," in *SPIE Conference on Chemical Microsensors and Applications*, vol. 3539, Boston, MA, November 1998.

[9] F. W. Edridge-Green, *Memory and its cultivation*. D. Appleton and Co., 1900, sited in [32].

[10] See: http://www.robotics.com/.

[11] L. Rabiner and B.-H. Juang, *Fundamentals of Speech Recognition*. Prentice Hall, Inc., 1993.

[12] M. J. Tovee, *An introduction to the visual system*. Cambridge University Press, 1996.

[13] W. Zhu and S. E. Levinson, "Edge orientation-based multiview object recognition," in *Proc. IEEE Int'l Conf. on Pattern Recognition*, vol. 1, Barcelona, Spain, 2000, pp. 936–939.

[14] W. Zhu, S. Wang, R. S. Lin, and S. E. Levinson, "Tracking of object with SVM regression," in *Proc. IEEE Int'l Conf. on Computer Vision & Pattern Recognition*, vol. 2, Hawaii, 2001, pp. 240–245.

[15] R. S. Lin, "Learning vision-based robot navigation," Master's thesis, University of Illinois at Urbana-Champaign, 2002.

[16] E. A. Macpherson and J. C. Middlebrooks, "Listener weighting of cues for lateral angle: The duplex theory of sound localization revisited," *J. Acoust. Soc. Am.*, vol. 111, no. 5, pp. 2219–2236, May 2002.

[17] D. Li and S. E. Levinson, "A bayes-rule based hierarchical system for binaural sound source localization," in *Proc. IEEE Internat. Conf. Acoust. Speech Signal Process*, Hong Kong, Apr. 2003.

[18] ——, "A robust linear phase unwrapping method for dual-channel sound source localization," in *International Conference on Robotics and Autiomation*, Washington D. C., May 2002.

[19] D. Li, "Computational models for binaural sound source localization and sound understanding," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2003.

[20] C. G. Atkeson, A. W. Moore, and S. Schaal, "Locally weighted learning," *Artificial Intelligence Review*, vol. 11, pp. 11–73, 1997.

[21] A. K. McCallum, "Instance-based state identification for reinforcement learning," in *Advanced in Neural Information Processing Systems*, vol. 7, 1995.

[22] H. Asada, S. Noda, S. Tawaratsumida, and K. Hosoda, "Purposive behavior acquisition for a real robot by vision-based reinforcement learning," *Machine Learning*, vol. 23, pp. 279–303, 1996.

[23] V. Gullapalli, J. Franklin, and H. Benbrahim, "Acquiring robot skills via reinforcement learning," *IEEE Control Syst. Mag.*, vol. 14, pp. 13–24, 1994.

[24] S. Thrun, "An approach to learning mobile robot navigation," *Robotics and Autonomous Systems*, pp. 301–319, 1995.

[25] W. Zhu and S. E. Levinson, "PQ-learning: an efficient robot learning method for intelligent behavior acquisition," in *Proc. 7th International Conference on Intelligent Autonomous Systems*, vol. 1, Marina del Rey, CA, USA, March 2002, pp. 404–411.

[26] Q. Liu, "Interactive and incremental learning via a multisensory mobile robot," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2001.

[27] W. Zhu and S. E. Levinson, "JPDF-based visual concept learning by an autonomous agent," to appear in Proc. of International Conference on Vision Interface 2003.

[28] M. Kleffner, "A method of automatic speech imitation via warped linear prediction," Master's thesis, University of Illinois at Urbana-Champaign, 2003.

[29] S. Fine, Y. Singer, and N. Tishby, "The hierarchical hidden Markov model: Analysis and applications," *Machine Learning*, vol. 32, no. 1, pp. 41–62, 1998.

[30] V. Krishnamurthy and J. B. Moore, "On-line estimation of hidden Markov model parameters based on the Kullback-Leiber information measure," *IEEE Trans. Signal Processing*, vol. 41, no. 8, pp. 2557–2573, August 1993.

[31] A. Gopnik, A. N. Meltzoff, and P. K. Kuhl, *The Scientist in the Crib*. HarperCollins, 1999.

[32] M. H. Ashcraft, *Human Memory and Cognition*. Harper Collins, 1989.