A System for Epigenetic Concept Development through Autonomous Associative Learning

Laura M. Grabowski
grabow38@msu.edu

Matthew D. Luciw
luciwmat@msu.edu

Juyang Weng
weng@cse.msu.edu

Dept. of Computer Science and Engineering
Michigan State University
East Lansing, MI 48824 USA

***Abstract*— In early development, an autonomous agent must learn to understand its sensors. This is thought to be done via exploratory movements. These result in low-level understanding by exploring sensorimotor associations. This paper is concerned with the further development of internal sensors that can sense and convey a concept. A concept is an abstract and compact representation of information from multiple sources. This paper introduces a system for the development of concepts, with the key features of associative learning, internal attention pathways, and exploratory movements. It is shown, within a mobile agent in a simulation environment, to develop an internal sensor that measures the semi-concrete concept of distance traveled. Distance is shown to be understood correctly when the agent is moved (either passively or actively) in different environments and at a set of movement speeds. This system is general, and can be used to develop other concepts for other types of agents.**

***Index Terms*— concept learning, developmental mental architectures, internal attention, sensorimotor pathways, emergent cognition**

I. Introduction

A. Motivation

The notion of "concept" raises controversies in psychology, philosophy, and cognitive science. These disputes extend to the foundations of theories of concepts. Laurence and Margolis 1999 [1] highlight some of these fundamental debates: Are concepts mental representations or abstract entities? Do concepts embody mental theories, or are they collections of features? Are concepts objects, or behavioral abilities? Differences of opinion on these kinds of issues can be quite distinct between disciplines, and relate to discussions of the nature of cognition itself. Vernon et al. 2007 discuss the issue of cognition in the context of artificial systems, voicing the perspective from which the present work arises: "Cognition... can be viewed as the process by which the system achieves robust adaptive, anticipatory, autonomous behavior, entailing embodied perception and action." [2, p. 151]. This type of emergent systems viewpoint allows for cognition that is not necessarily in the self-reflective domain, greatly widening the pool of cognitive systems.

This work is supported in part by an IGERT Fellowship (NSF DGE 0114378).

The current work presents an emergent, bottom-up approach to concept development, in which the concept emerges through association of perceptual experiences from different modalities (e.g., visual, auditory, somatosensory). The representation of the concept encapsulates the different perceptual information that constitute the concept, and its output represents all the constituent properties. From the standpoint of information theory, a concept can be thought of as a package that allows information to be transferred in a compressed, coupled form, instead of in many individual pieces. A concept is different from the communication, or utilization, of the concept. The concept is the underlying representation. For example, the act of verbally identifying an object requires such an abstract and multimodal representation. A spoken word can represent a package of information over multiple modalities, such as visual (recognize the object despite possible variations), somatosensory (speak the object's name), audition, and touch. However, the same concept can be communicated in a different way (e.g., a written word).

From the computational neuroscience perspective, how might concepts be developed and stored? We guess that concepts are both formed and stored at associative areas. An associative area has multisensory input, and thus the output of neurons can represent multiple modalities. From the different input modalities, the sensory inputs might first be pre-processed to be more abstract, as they are in the human brain, where, information to each of the three association cortices comes from nearly every modality, usually from the higher-order sensory cortices (Kupfermann 1991 [3]). This convergence of information allows for what Menzel and Giurfa 2001 [4] call "horizontal processing," analogous to associative learning. This horizontal processing is the basis for the capabilities in honeybees that would be considered cognitive by Vernon et al.'s definition [2]. Of course, honeybees do not possess a cortex, but there are two dense associative areas in their brains, called the mushroom bodies. This could be the neural area that contributes, via associative learning, to the bees' concepts, as demonstrated, for example, in visual categorization (Menzel 2006 [5]).

There is an ongoing debate about whether the human child is born with basic knowledge about the physical world in core domains (see e.g., Wellman and Gelman 1992 [6]).

As neuroscience advances, we expect that these issues will reduce to phenotypical phenomena of the genetic regulations in epigenetic development, which are likely more low level than what can be called "knowledge about the physical world."

B. Types of Concepts and Concept Development

Concept development is studied extensively in developmental psychology (Carey 2000 [7]; Opfer and Siegler 2004 [8]). A concept can be thought of as a mental representation of tangible objects in the brain-external world (e.g., a spoon) and intangible ideas and feelings in the brain-internal world (e.g., colors, emotions). A concept differs from a skill; a skill is an ability to do a task (e.g., tracking a moving toy, perform path integration during navigation). From this perspective, concepts can be grouped into three categories: (1) concrete concepts, those that are tangible (e.g., a dog); (2) semi-concrete concepts, those that relate to something which can be demonstrated, but is not tangible (e.g., blue, distance); (3) abstract concepts, those relating to mental activities (e.g., fear). In the Self-Aware and Self-Effecting (SASE) mental architecture proposed in Weng and Zeng 2005 [9], the activities of brain-internal space are treated as internal context, in parallel with the brain-external contexts. Therefore, our proposed model treats abstract and semi-concrete concepts in the same way. In the work presented here, we focus on a computational scheme of autonomous associative learning through exploratory movements, which enables the development of concrete concepts and semi-concrete concepts. Previous work has shown how concrete concepts can be developed; this paper focuses on the development of semi-concrete concepts. The success of concept development is shown by performing skills that utilize the concepts. The sophistication of such acquired skills is an important indication of different stages of development and, at the current stage of research, we expect that the experimentally demonstrable sophistication corresponds to early development.

Here, we present a computational model for concept development through autonomous associative learning using biologically inspired mental architectures (*Type-2*, as defined in Weng and Zeng 2005 [9]), and the building block of the architecture, the sensorimotor pathway model called Multilayer In-Place Learning Network (MILN; Weng and Luciw 2006 [10]). To our knowledge, this is the first work in the literature to deal with the important issue of concept development computationally by a biologically inspired developmental architecture (Type-2) and a developmental neural network (MILN). Due to the great challenges of such a study, we will demonstrate experimental results with an agent in an unknown but simplified simulation environment. However, the principles presented in the paper are applicable to more sophisticated settings.

C. Fundamental Mappings

Exploration of sensorimotor space and development of sensorimotor pathways in artificial agents have been investigated by many researchers. Pierce and Kuipers 1997 [11] presented a sequence of learning methods enabling a robot to learn a cognitive map of its sensorimotor apparatus and environment, without domain-dependent knowledge of its environment or sensorimotor apparatus. Olsson etal. 2004 [12] investigated how an agent with neither innate knowledge of its sensory data modalities nor a model of its own physical configuration can discover structure in its sensory devices from raw sensor data. Philipona et al. 2003 [13] presented an algorithm that discovers information about the structure of the world by analyzing the law linking motor outputs and sensory inputs, regardless of the agent's body.

The problem of grounding (Harnad 1990 [14]) is well known in artificial intelligence and cognitive science: how can the semantics of an agent's internal state be intrinsic (self-generated) and not imposed from outside (hand-programmed)? An agent has direct access only to its own internal state, never to the external world or its properties (Choe and Smith 2006 [15]). Natural agents can make the inference from internal state to the external properties that state represents, but the mechanisms remain unknown. Choe and Smith 2006 [15] suggested that motor primitives play a key role in solving this problem by associating external stimulus properties with internal sensory states. Sensory invariance was used to enforce the mapping between sensory state and environmental property, grounding internal state representations. These fundamental mappings are crucial, enabling an agent to perform a number of simple tasks. The development of more complex behaviors requires concepts to serve as the foundation for new capabilities.

D. Path Integration

Animals, including humans, employ a variety of navigational strategies. Successful navigation utilizes several different concepts, such as distance, direction, time, and landmarks. Path integration, also called dead reckoning, is a fundamental and ubiquitous strategy, operating in diverse species, both invertebrate and vertebrate, including humans (Loomis et al. 1999 [16]. Path integration is the continual updating of position relative to a location, based on velocity, temporal, and acceleration information (Roche et al. 2005 [17], functioning automatically and constantly whenever the animal moves in continuous space (Etienne and Jeffery 2004 [18]. Path integration allows the animal to return to the starting location by the shortest, most direct route, without retracing the outbound path, and functions even in terrain where landmark cues are absent or unreliable. Precisely what an animal's nervous system is adding or integrating and how the computations are performed are unknown.

The animal requires travel direction and distance information for the path ingtegration computations. Distance measurements are given mostly by self-motion cues, derived from different sensory sources (visual, vestibular, proprioceptive). Directional information can be obtained by reference to a fixed landmark of known orientation, or by measuring rotations since the last known heading, with or without the use of external references (Wehner, Michel, and Antonsen 1996 [19]). In this paper, we study the development of a distance concept. Visual, vestibular, and somatosensory information contribute to the distance concept in humans; the underlying representation may be that of a velocity profile. Humans may be attempting to recreate an experienced velocity profile in order to recreate distance traveled (Bremmer and Lappe 1999 [20]).

II. Architecture for Autonomous Associative Learning

Unlike the traditional Markov Decision Process (MDP) (Kaelbling 1996 [21]), a developmental architecture is observation-driven: the agent's mental states are automatically generated from experience, instead of hand-selected by a designer. In the latter case, the meaning of each state is easy to understand, since that meaning was assigned. But in the observation-driven case, the meaning of each state is not initially known by either the agent or the programmer. One way to learn this meaning is through associating experiences in one modality with experiences in other modalities (see e.g. Pfeifer and Scheier 1997 [22]). This is associative learning; however, autonomous associative learning requires additional selective capability. It must (a) autonomously select two subparts of recent experience to attend (autonomous attention) as conditional context and dependent context, and (b) establish the association (mapping) between them.

The architecture of the agent in this paper is classified as a *Type-2* architecture, an Observation-driven selective Markov Decision Process, as defined in Weng and Zeng 2005 [9]. It is observation driven as explained above. It is selective via use of attention. A methodological novelty in this paper is the selective aspect of the architecture. Two innate attention pathways are used to quickly—in few samples— learn two meaningful associations : sensorimotor associations from conditional sensation to dependent action, and the inverse. After this two-way mapping is refined, if only one type of information is observed, the other is automatically implied or "filled in". This is possible since neurons in an association area that represents multiple information modalities will fire even when some information is missing; the missing context is automatically implied by the neuron's firing. Generally, all properties that constitute a concept are implied by the firing of a neuron representing them; however, not all properties have to be observed for that neuron to fire.

We use exploratory movements (locomotion) to develop a spatial, body-centered concept of forward distance traveled. The mathematical space of coupled possible sensations and actions can be called the sensorimotor manifold. Exploratory or randomized movements allow the agent to actively explore this manifold. The distance concept develops from exploratory movements, and contains (1) knowledge of how the sensed environment should have changed when the agent locomotes in a certain way, and (2) how the agent should have locomoted during a sensed change of the sensed environment. Both of these expectations are learned from experience.

A. Environment and the Agent Design

The agent is placed in a sparse environment (Fig. 1). The only features are the agent, the home location, and a set of b *beacons*. A beacon is a relatively stable feature in the environment that is somehow salient, e.g., a distinctive visual feature that might serve as a landmark. The beacons are placed randomly in each environment, with some constraints to prevent beacons from being too far from or too near the home location. The agent itself has a front and a back, and an orientation, a certain direction that it "faces". A developmental agent is equipped with a set of sensors, which it can use to experience the world, and effectors, which it can use to change the world. Some sensors and effectors are internal (affecting the brain only); others are external. This particular agent, modeled in part on Vickerstaff and Di Paolo 2005 [23], is equipped with three external sensors S_{e1}, S_{e2}, S_{e3} and one external effector E_{e1}.

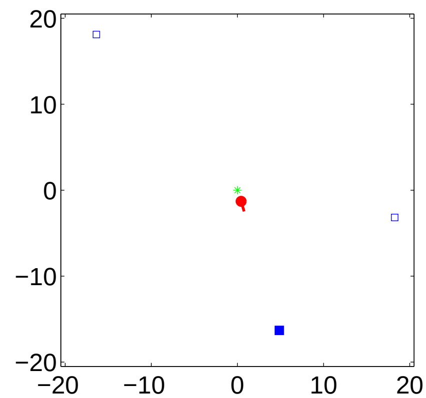

Fig. 1. Sample environment with agent. The agent is shown with the filled red circle, with its orientation indicated by the needle. Beacons are shown by the blue squares; the filled square is the currently attended beacon. The agent's home location is shown by the green star.

The first two sensors are sensitive to the phenomenon of placement of beacons in the environment, sensing the vector of angles θ from the beacons to the agent's forward facing direction and the distances $\mathbf{d}$ from the beacons to the agent, respectively. This agent does not have internal access to this absolute world knowledge. The first sensor $S_{e1}(t)$ represents the angles from the b beacons to the agent's orientation by a vector with $2b$ components. The index t stands for time, which is discrete, incrementing from 0. The following activation functions give the actual sensed values in this vector for the $\lceil i/2 \rceil$-th beacon (where $1 \leq i \leq 2b$):

$$x_{e1i}(t) = \frac{cos\left(\theta_{\lceil i/2 \rceil} - \frac{\pi}{2}\right)}{2} + 0.5 \qquad (1)$$

$$x_{e1(i+1)}(t) = \frac{cos\left(\theta_{\lceil i/2 \rceil} + \frac{\pi}{2}\right)}{2} + 0.5 \qquad (2)$$

These two sensor components can be considered as placed directly on the right and left sides of the agent: the response of the right sensor will be 1 for a beacon directly to the agent's right, and 0 for a beacon directly to the agent's left.

The second sensor $S_{e2}(t)$ represents the b distances from the beacons to the agent. Let $\mathbf{d} = (d_1, ..., d_b)$ where d_j is the absolute distance to the j-th beacon. Then

$$x_{e2j}(t) = \begin{cases} 0 & \text{if } d_j > r \\ 1 - \dfrac{d_j}{r} & otherwise \end{cases} \quad (3)$$

where r is the range of the sensor. The sensor activation increases as the beacon distance decreases.

For movement, the agent is equipped with two wheels. It can move within the environment by rotating either or both of these wheels. The rotations of the wheels are controlled by the effector $E_{e1}(t) = (a_{e11}(t), a_{e12}(t))$ where each component is constrained to be between 0 and 1.

This action is fed back to the third external sensor

$$S_{e3}(t) = (x_{e31}(t), x_{e32}(t)) = (a_{e11}(t-1), a_{e12}(t-1)) \quad (4)$$

that detects what actions were done in the previous time step.

The agent's *observation* $\mathbf{x}(t)$ contains all its sensations at time t, concatenated into a vector. The *outcome* $\mathbf{p}(t) \in \mathcal{P}$ is the next sensations, e.g., at $t+1$. That is, the next outcome will be the sensed result of the current observation (which includes the sensed last actions). A critical aspect of a developmental architecture is that, as the agent undergoes more experience, it is better able to predict the next outcome from all previous observations. One way to do so is requiring that at time t, the agent must estimate the probability distribution of $P(\mathbf{p}(t)|H_t = h)$ where $H_t = \{\mathbf{x}(t), \mathbf{x}(t-1), ..., \mathbf{x}(0), \mathbf{p}(t-1), ..., \mathbf{p}(0)\}$ is the history, the set of all observations and outcomes starting at $t = 0$.

Practically, as t increases, a mental architecture that uses the entire experience history to predict the next outcome is not feasible. Instead, reduce the temporal length of H_t and call this the *last context* l_t. The last context is a set parameterized by the number k of last observations and outcomes to store: $l_t = \{\mathbf{x}(t), \mathbf{x}(t-1), ..., \mathbf{x}(t-k), \mathbf{p}(t-1), ..., \mathbf{p}(t-k)\}$. Concatenate all vectors in this set into a vector $\mathbf{l}(t) \in \mathcal{L}$. An architecture that can learn the mapping from $\mathcal{L} \rightarrow \mathcal{P}$, from experience only, will be more practical.

A component called a *regressor* R learns to map $\mathcal{L} \rightarrow 2^{\mathcal{P}}$. The function of the regressor is to learn to output a set of predicted outcomes having high probability given the last context:[1] $\{\mathbf{p}_1(t), \mathbf{p}_2(t), ..., \mathbf{p}_n(t)\} = R(\mathbf{l}(t))$.

B. Multilayer In-place Learning

The regressor used in this study is the Multi-Layer In-Place Learning Network (MILN; see Weng and Luciw 2006 [10] for a more detailed presentation). It is designed to be used in open-ended developmental systems and learns a mapping incrementally, from examples, using a limited neural resource. The mapping may have complicated, non-stationary, and high-dimensional input and/or output spaces, as is often the case with real sensor to motor mappings. MILN is highly biologically motivated. Through mechanisms of Hebbian learning and lateral inhibition, it approximates the manifolds in the sample input and output spaces with a set of neurons called lobe components (Weng et al. 2006 [24]), each of which represents the set of samples in its lobe region nearly optimally [2].

C. Attention Selection

Attention selection allows the regressor to learn mappings while disregarding some irrelevant dimensions of the sensory input. We model attention selection as a mask of ones and zeros, applied to the last context $\mathbf{l}(t)$ to create new context by applying the mask before learning. The masks themselves could be innate or learned. We used several types of innate masks, discussed in the next section. First, we applied masks to S_{e1} and S_{e2} so that the agent attends to only the single closest beacon. The other beacons are ignored (set to zero).

1) Learning Associations With Internal Attention: Without internal attention, MILN will learn the function $\mathbf{l}(t) = R(\mathbf{l}(t))$. Two innate *attention pathways* are used to learn more meaningful associations. To better understand conceptually, consider that there are two parts to the last context called A and B. A is the sensor components only. B is the effector components only. One attention pathway enables learning of the association from sensors (A) to motors (B). The other pathway enables learning of the association from motors (B) to sensors (A). Internal attention is like a switch that chooses what path to learn. The learned associations do not even have to be sensorimotor; all that is required is that A and B are different. The single function that is learned, denoted f, will *approximate* both $A = f(B)$ and $B = f(A)$. Since it learns both mappings, this function can perform automatic "filling in" whenever either part is missing.

Formally, *each* attention pathway is a set of two different masks $\mathbf{m}^{(1)}$ and $\mathbf{m}^{(2)}$, applied to $\mathbf{l}(t)$ so that $\mathbf{l}^{(1)}(t) = \mathbf{m}^{(1)} \otimes \mathbf{l}(t)$ and $\mathbf{l}^{(2)}(t) = \mathbf{m}^{(2)} \otimes \mathbf{l}(t)$. The function now learned is $\mathbf{l}^{(2)}(t) = R(\mathbf{l}^{(1)}(t))$. In terms of associative learning, $\mathbf{l}^{(1)}$ is the conditional context, and $\mathbf{l}^{(2)}$ is the dependent context. In our experiment, we use two pathways. One pathway maps from beacon sensors (input space) to motor sensors (output space). The second pathway is the inverse, from motor sensors to beacon sensors. If the number k of observations within $\mathbf{l}(t)$ is greater than 1, the associations will be over sequences of movements and sensations, instead of for experiences from a single time step.

A two-layer MILN can be used for two modalities (as in our experiments). In this simple case, the bottom-up input to layer-one is the conditional context $\mathbf{l}^{(1)}$. Neurons will self-organize and compete to represent their inputs using their bottom-up, lateral, and top-down weights. "Winning" neurons can fire and will adjust their weights to try to fire more in the future. Non-winning neurons are inhibited and will not fire (their firing rate equals zero). Denote the output firing rate of layer-one as $\mathbf{z}^{(1)}$. The bottom-up input to layer-two is a concatenation of both $\mathbf{z}^{(1)}$ and the dependent context $\mathbf{l}^{(2)}$. Thus, the output firing rate

[1] A value system was not needed for this study, so it is not discussed in detail. The function of the value system is to select the single "best" predicted outcome, thus completing the mapping to $\mathcal{P}$.

[2] The samples are represented optimally if an averaging technique called "CCI Plasticity" is not used. However, this technique allows the system to continue learning as time increases. Normal averaging prevents adaptation at large t.

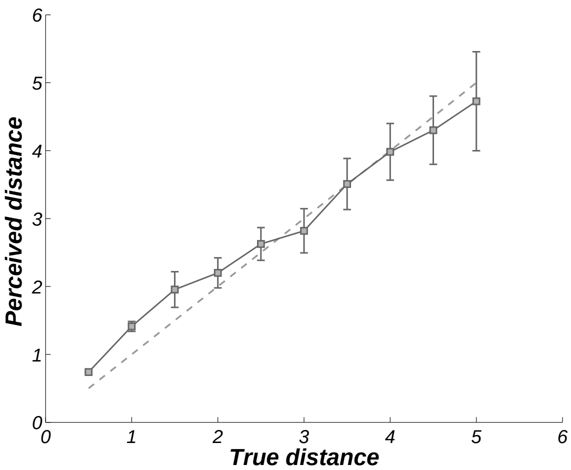

Fig. 2. The agent's perception of distance as it was passively moved in 100 different environments, for 10 time steps. Each square is the average distance perceived after each time, from 1 to 10. The dotted line shows perfect perception.

$\mathbf{z}^{(2)}$ of layer-two contains information about both modalities. There is natural information loss since the number of neurons is much less than the number of observations. This associative area can be induced to fire even when some information from one of the modalities is missing from the last context.

III. The Learning Algorithm

The following parameters must be set for each t: (1) the number k of previous observations to keep in the last context, impacting temporal length of the sequences learned; (2) the currently active attention pathway, affecting the type of associations learned; (3) any number of other application-specific parameters set as a function of time, or when other events occur. Denote all parameters as components of a parameter vector θ and a function $g(t, \mathbf{e}) = \theta$ that sets each parameter during learning, where components of the binary $\mathbf{e}$ are flagged after application-specific events (not specifically temporally dependent).

Initially, set time $t = 0$. Set the initial θ and define g. Do the following in a loop, ending when t has reached a pre-set maximum value. For each t:

1) Set the current parameter vector $\theta = g(t, \mathbf{e})$, including k and the active attention pathway.
2) Generate a randomized exploratory action.
3) The agent observes the current sensory frame $\mathbf{x}(t)$.
4) The last context $\mathbf{l}(t)$ is updated to contain the current frame $\mathbf{x}(t)$. $\mathbf{l}(t) = (\mathbf{x}(t), \mathbf{x}(t-1), ..., \mathbf{x}(t-k))$
5) The current attention pathway specifies the current attention masks that generate $\mathbf{l}^{(1)}(t)$ and $\mathbf{l}^{(2)}(t)$.
6) Update the agent's current cortex $\mathbf{l}^{(2)}(t) = R(\mathbf{l}^{(1)}(t))$, where the role of the module R is filled by a multi-layer, in-place learning network.
7) Increment t.

IV. Experiments and Results

Using the algorithm, the concepts develop through incrementally establishing and refining the mapping from equivalent instances of conditional to dependent context. The equivalence is learned from the agent's active exploration of the sensorimotor manifold. Any mapping can be learned if the sensorimotor experiences present their relationships consistently. In the following, we show how the algorithm can lead to the semi-concrete concept of distance traveled.

A. Distance Reproduction: Environment Invariance

The learning algorithm described in Section III was applied to an agent placed within a 40 by 40 unit environment containing 3 randomly placed beacons. The agent acted in a range of randomized movements. It was allowed to wander for 10,000 time steps; if it reached the edge of the enclosure, its position was reset to the home location (the origin). The internal attention path was switched to the other pathway every 50 time steps. The length k of the last context, initially set to 1, was incremented every 200 time steps. The maximum k was set at 15; if k was increased above this, it was reset to 1.

The network consisted of 500 neurons on two layers (total 1000) to approximate the input and output spaces. On each layer, for each new input, the top-3 winning neurons would update their weights and fire. After 5,000 time steps, only the associative weights were allowed to update.

We hypothesized that, after refinement, the output firing rate of layer-two would encapsulate two types of information: 1) "beacon-flow," (like optic flow) sequences of sensed beacon position during movement, and 2) sensed self-movement sequences. Recall that layer-two will fire whenever a beacon-flow and/or a self-movement sequence that is sufficiently similar to a stored neural prototype is observed. In passive movement, the agent experiences beacon flow sensations, but no self-movement sensations. From our architecture, the agent would perceptually fill in the missing self-movement. The accuracy of this perceived action sequence—which may not be the actual sequence—was tested by inducing the agent to perform that action sequence, while recording the actual distance traveled.

This experiment was performed by moving the agent passively forward in 100 different randomly generated, single-beacon environments, for 10 time steps. After passive movement in each environment, the agent would then actively move using the associated action sequence implied by the firing of layer-two. These "perceived distances" were measured and compared with the true distances as shown in Figure 2.

B. Distance Reproduction: Environment and Speed Invariance

A second experiment tested development of the distance concept invariant to two different speeds. We found that this required the following. (1) The two speeds must be different enough to allow the network to differentiate between them. We used a "slow" (mean .25) and "fast" (mean .75) speed with variances of .05. (2) Speed could change only after a pause of inactivity, since, if the speed shifted suddenly, neurons in

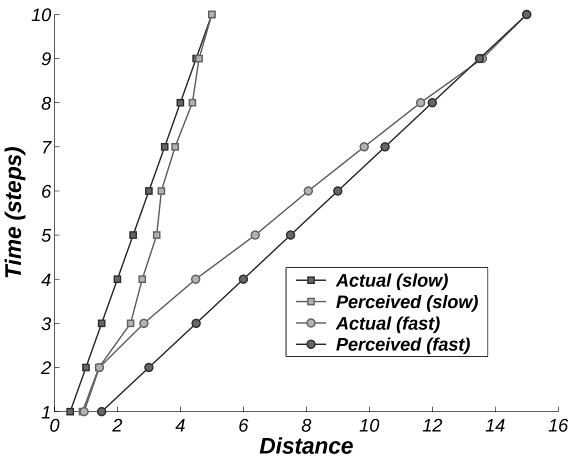

Fig. 3. The agent's perception of distance for two different speeds. It was tested with 100 different beacon placements, over 10 time steps, and the two speeds. Each square is the average distance perceived at each time, from 1 to 10.

the network started to average the two speeds, generating a third intermediate speed. We switched speed when the agent reached the edge of the enclosure and zeroed the last context vector.(3) The beacons must be close enough to the agent for the network to discriminate differences in angle and distance between the two speeds, over time. If the beacon is too far away, the difference in change in angle over a few time steps between the two speeds will be negligible. We constrained the agent to always move toward a beacon and used a smaller enclosure (15 by 15). The network used 2000 neurons each for input and output spaces, but the other parameters were the same.

We performed the same testing procedure as before, except we passively moved the agent at either the slow or fast speed. The results are shown in Figure 3. The agent tends to over-estimate distance traveled at the slow speed and underestimate distance traveled at the fast speed, except with longer paths. However, the results clearly show that 2 distinct pathways developed, corresponding to the 2 speeds. The agent is able to to judge the correct distance with an average error of 0.45 units for the slow speed (traveling 5 units/10 steps) and 0.85 units for the fast speed (traveling 15 units/10 steps).

V. CONCLUSION

We present a computational model for sensorimotor concept development through autonomous associative learning, using a biologically inspired developmental architecture and developmental network. Our results demonstrate how a semi-concrete concept can be learned in a simulation environment, but the principles are applicable to more complex and sophisticated settings, concepts, and skills.

REFERENCES

[1] S. Laurence, and E. Margolis, "Concepts and cognitive science," in E. Margolis and S Laurence, Eds., Concepts: Core Readings, Cambridge, MA: MIT Press, pp. 3-81, 1999.

[2] D. Vernon, G. Metta, and B. Sandini, "A survey of artificial cognitive systems: implications for the autonomous development of mental capabilities in computational agents," IEEE Transactions on Evolutionary Computation, vol. 11, no. 2, pp. 151-180, April 2007.

[3] I. Kupfermann, "Localization of higher cognitive and affective functions: the association cortices," in E. R. Kandel, J. H. Schwartz, and T. M. Jessell, Eds., Principles of Neural Science, 3rd ed., Norwalk, CT: Appleton & Lange, pp. 823-825, 1991.

[4] R. Menzel and M. Giurfa, "Cognitive architecture of a mini-brain: the honeybee," Trends in Cognitive Sciences, vol. 5, no. 2, pp. 62-71, 2001.

[5] R. Menzel and M. Giurfa, "Dimensions of cognition in an insect, the honeybee," Behavioral and Cognitive Neuroscience Reviews, vol. 5, no. 1, pp. 24-40, 2006.

[6] H.M. Wellman, and S.A. Gelman, "Cognitive development: foundational theories of core domains," Annual Reviews of Psychology, vol. 43, pp. 337-375, 1992.

[7] S. Carey, http://www.wjh.harvard.edu/ lds/pdfs/carey2000b.pdf, "The origin of concepts," Journal of Cognition and Development, vol. 1, pp. 37-41, 2000.

[8] J.E. Opfer, and R.S. Siegler, "Revisiting preschoolers' living things concept: a microgenetic analysis of conceptual change in basic biology," Cognitive Psychology, 2004.

[9] J. Weng, and S. Zeng, "A theory of developmental mental architecture and the Dav architecture design," International Journal of Humanoid Robotics, vol, 2, no. 2, pp. 145-179, 2005.

[10] J. Weng, and M.D. Luciw, "Optimal in-place self-organization for cortical development: limited cells, sparse coding and cortical topography," Proc. IEEE 5th International Conf. on Development and Learning (ICDL 2006), May 30 - June 3, 2006, Bloomington, IN, USA.

[11] D. Pierce and B. J. Kuipers, "Map learning with uninterpreted sensors and effectors," Artificial Intelligence, vol. 92, pp. 169-227, 1997.

[12] L. Olsson, C. L. Nehaniv, and D. Polani, "Sensory channel grouping and structure from uninterpreted sensor data," 2004 NASA/DoD Conference on Evolvable Hardware, pp. 153-160, 2004.

[13] D. Philipona, J.K. O'Regan, J.-P. Nadal, and O. J.-M. D. Coenen, "Perception of the structure of the physical world using unknown multimodal sensors and effectors," in Advances in Neural Information Processing Systems 16, S. Thrun, L. Saul and B. Schölkopf, Eds. Cambridge: MIT Press, 2004.

[14] S. Harnad, "The symbol grounding problem," Physica D, vol. 42, pp. 335-346, 1990.

[15] Y. Choe and N. H. Smith, "Motion-based autonomous grounding: inferring external world properties from encoded internal iensory states alone," Proc. 21st Natl. Conf. on Artificial Intelligence and the 18th Innovative Applications of Artificial Intelligence Conf., July 16-20, 2006, Boston, MA, USA.

[16] J.M. Loomis, R.L. Klatzky, R.G. Golledge, and J.W. Philbeck, "Human navigation by path integration," in R.G. Golledge, Ed., Wayfinding behavior: cognitive mapping and other spatial processes, Baltimore: Johns Hopkins University Press, pp. 125-151, 1999.

[17] R. A. P. Roche, M. A. Mangaong, S. Commins,and S. M. O'Mara, "Hippocampal contributions to neurocognitive mapping in humans: a new model," Hippocampus, vol. 15, pp. 622-641, 2005.

[18] A. S. Etienne, and K. J. Jeffery, "Path integration in mammals," Hippocampus, vol. 14, pp. 180-192, 2004.

[19] R. Wehner, B. Michel, and P. Antonsen, "Visual navigation in insects: coupling of egocentric and geocentric information," Journal of Experimental Biology, vol 199, pp. 129-140, January 1996.

[20] F. Bremmer, and M. Lappe, "The use of optical velocities for distance discrimination and reproduction during visually simulated self motion," Experimental Brain Research, vol. 127, pp. 33-42, 1999.

[21] L. P. Kaelbling, M. L. Littman, and A. W. Moore, "Reinforcement learning: a survey," Journal of Artificial Intelligence Research, vol 4, pp. 237-285, 1996.

[22] R. Pfeifer, and C. Scheier, "Sensory-motor coordination: the metaphor and beyond," Autonomous Systems, vol. 20, pp. 157-178, 1997.

[23] R. Vickerstaff, and E.A. Di Paolo, "Evolving neural models of path integration," Journal of Experimental Biology, vol. 208, pp. 3349-3366, 2005.

[24] J. Weng, and N. Zhang, "Optimal In-Place Learning and the Lobe Component Analysis," Proc. World Congress on Computational Intelligence, July 16-21, 2006, Vancouver, Canada.