In Berthouze, L., Kozima, H., Prince, C. G., Sandini, G., Stojanov, G., Metta, G., and Balkenius, C. (Eds.)
Proceedings of the Fourth International Workshop on Epigenetic Robotics
Lund University Cognitive Studies, 117, ISBN 91-974741-3-4

Developmental Learning: A Case Study in Understanding "Object Permanence"

Yi Chen and Juyang Weng
Department of Computer Science and Engineering
Michigan State University
East Lansing, MI, 48824
{chenyi1,weng}@cse.msu.edu

Abstract

The concepts of muddy environment and muddy tasks set the ground for us to understand the essence of intelligence, both artificial and natural, which further motivates the need of *Developmental Learning* for machines. In this paper, a biologically inspired computational model is proposed to study one of the fundamental and controversial issues in cognitive science – "Object Permanence." This model is implemented on a robot, which enables us to examine the robot's behavior based on perceptual development through real-time experiences. Our experimental result shows consistency with prior researches on human infants, which not only sheds light on the highly controversial issue of object permanence, but also demonstrates how biologically inspired developmental models can potentially develop intelligent machines and verify computational modeling that has been established in cognitive science.

1. Introduction

Object permanence is the understanding that objects continue to exist, even when they cannot directly be perceived. It is an important theoretical construct that has been widely researched with infants. However, the understanding of the underlying principles that enable the development of cognitive and behavioral capabilities demonstrated in object permanence goes beyond the issue per se. Above all, it bears upon a fundamental issue of the origin of knowledge that has divided scientists and great thinkers ever since the time of classical antiquity.

Where does our knowledge come from? Two kind of sources are ususally considered: what is given us by virtue of our nature, or what we know as a consequence of our nurture. This debate of "nature vs. nurture" was originally started by Plato (428-347 BC), who believed that all knowledge is innate, and his student Aristotle (384-322 BC), who argued that mind is only a *tabula rasa* at birth. In modern times, this debate has resulted into two academic schools in developmental psychology. On the one side, the constructivists, represented by Jean Piaget, believed that infants make sense of sensory information through interaction with environment. On the other side, the nativists believed that some physical knowledge, one facet of which is the sense of object permanence, are innately available for perception.

More recently, however, advances in brain science understanding are reshaping this "nature vs. nurture" debate closer to a "nature AND nurture" consensus. The fact we know now is that learning involves changes in synaptic connections, and these changes are effected by the products of specific genes which are expressed only under certain environmental conditions. In other words, there must be something intrinsic that is innate to generate behaviors under certain environmental conditions, however, this little something might not be the knowledge itself, but certain intrinsic learning mechanisms.

Experimentally, this idea has been supported by recent research work in both psychology and neural science. For example, in a language study by Saffran et al. in 1996, infants of 8 month old were found out to be able to segment words solely based on statistical learning, indicating that babies may have an innate statistical learning mechanism which allows them to carry out such tasks instead of innate understanding ability (Saffran et al., 1996). In another experiment conducted on neonatal ferrets by Mriganka Sur in 2000, auditory cortex shows orientational selectivity after rewired into visual pathway (Melchner et al., 2000). This demonstrated that orientation selection is not embedded in the visual system as innate knowledge, but rather a principle developed from sensory inputs through general neural mechanisms, such as lateral inhibition and hebbian learning rule, etc.

In this paper, we focus on "Object Permanence" as one of the most fundamental issues in this general debate. By applying our computational model on a robot as a unique test-bed, we try to find out what potential innate mechanisms may be possible for the development of object permanence. This methodology provides an important complement to the observational and descriptive nature of developmental psychology.

Unfortunately, developing such computational models with general purpose is quite difficult. Schlesinger and Mareschal et al. have proposed some computational models to study the object permanence issue (Mareschal et al., 1995) (Schlesinger, 2002). However, their models are either embedded with prior knowledge, such as temporal contiguousness of objects, or limited with the capability to interact with the environment. In contrast, we propose a new experience-based learning mechanism called *Developmental Learning*. That is, no function of a specific task is predefined, and the internal representation is autonomously generated from all possible interactions with the environment. In addition, our model conducts incremental, real-time computation so that the association between real world experience and consequences of actions can be sensed and learned right away while the physical events carry on.

The result of this work is the first that we know to provide detailed developmental analysis for the study of object permanence on an autonomous robot. Specifically, our contributions in this paper include: (a) Implementation of the Developmental Learning mechanism into a task-independent developmental program on a robot; (b) Establishment of a general computational theory for novelty detection and novelty-based value system; (c) Training the robot in general environments (d)Testing the robot in a simulated setting (e)Testing the robot in the real experimental environment; (f) Comparison of our experimental data on 12 developed robot "brains" with prior results on human infants and its implications.

2. Background and Related Work

The term object permanence and its measure was first introduced by Jean Piaget in 1954 (Piaget, 1954). In his classic A-not-B task experiment, Piaget discovered that seven-to-twelve-month-old infants failed to retrieve a completely hidden object, showing insensitivity to object permanence when "out of sight is literally out of mind."

On the other hand, contemporary researchers have suggested that Piaget's manual search task was too conservative as a test for object permanence since it required sophisticated motor skills and efficient memory. Baillargeon and Spelke et al., then proposed a new type of experiment named "drawbridge", which only requires looking times as a measure, making it possible to test very young infants (Baillargeon et al., 1985). In their experiments, infants as young as 5 months old showed well-documented tendency to look longer at the impossible event than the possible event. According to Baillargeon and her colleagues, this is because babies were surprised or puzzled by the violation of physical laws in the impossible event, indicating that they already possessed an early physical knowledge of object permanence. However, recent studies by Schilling, Bogartz, Shinskey, Cashon, etc. argued that Baillargeon's results may reflect infants' perceptual capacities instead of conceptual understanding of "object permanence," which can be the preference for either familiarity, novelty or larger movement (Roder et al., 2000) (Rivera et al., 1999) (Cashon and Cohen, 2000) (Bogartz et al., 2000).

The open question is then: what mechanism if it is not physical knowledge that gives rise to such early perceptual capabilities? Although it is extremely difficult to give an exact answer, recent neuroscientific findings provides supporting evidence of novelty preference since a population of novelty neurons in monkey brain has been identified active during an occluding experiment in (Baker et al., 2001). These studies inspired us to use our robot as a computational model to test its perceptual novelty preference in the "drawbridge" experiment.

3. System Architecture

Fig. 1 shows the architecture implemented on our SAIL robot, which includes a sensorimotor system, motor mapping units and a value system. Sensations from both external and internal sensors are first sent to the sensorimotor system for learning. Here, the external sensors are two cameras to capture the visual images and the internal sensor is a sensor to sense the neck position from the neck motor. At the same time, external sensations from the keyboard are sent directly into the motor mapping unit, where the action output from the sensorimotor system is also captured and processed. The reason that only neck position is fed back as internal sensation is for experimental control. Just as mothers holding their babies in the original experiments, we try to prevent the robot from moving other than turning the neck during the experiment in order to avoid distraction. At last, the value system is employed to determine which neck action will be executed – stay, turn left or turn right, based on the individual Q-value of each action.

Different from the traditional computational modeling, our approach aims at enabling the robot to autonomously explore the world without much predefined physical knowledge. This is demonstrated in the following two ways:

We avoid modeling the world. The architecture does not need to include explicit knowledge about what the environmental world is like. Instead, it autonomously builds representation of the world, in a distributed and implicit way, through interactions with the world.

We avoid modeling the agent behavior. The architecture does not decompose the behaviors into any hierarchy, in the sense that one layer takes care of one behavior and another layer takes care of another. Although a number of basic behaviors are programmed, the amount is very limited compared with the behaviors that will be developed autonomously.

3.1 Representation

Autonomous generation of internal representation is the essential capability that is required to fulfill our goal. Therefore, the internal representation in our architecture

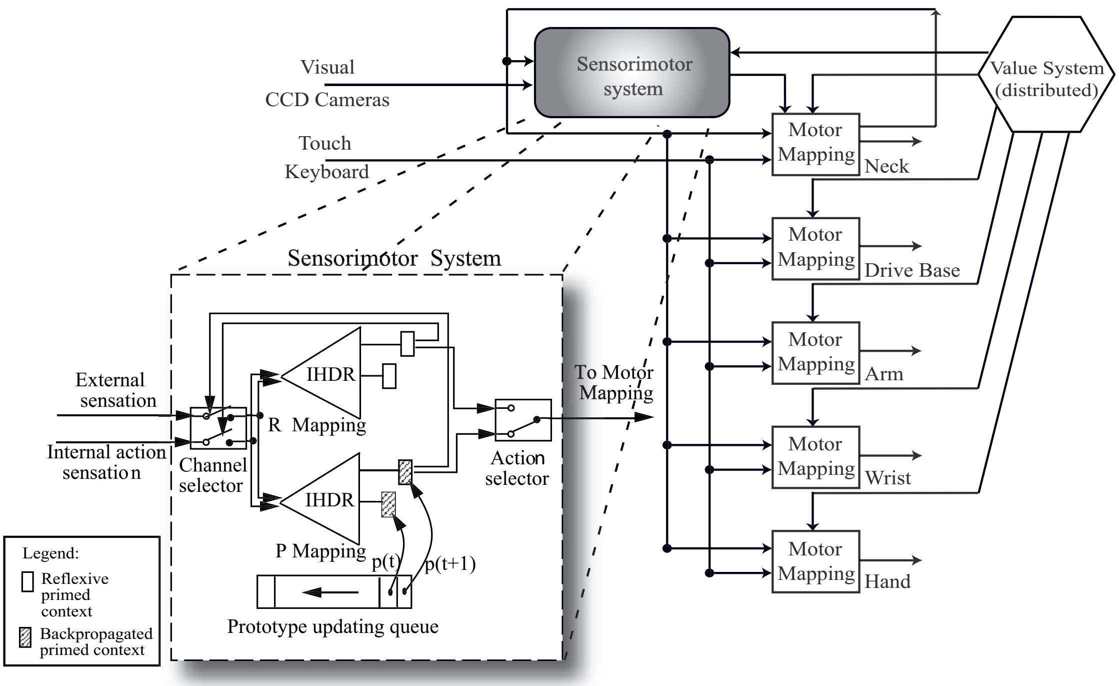

Figure 1: The architecture of SAIL robot mainly consists of three components: sensorimotor system, motor mapping unit and value system. The central unit of this architecture, as magnified, is the sensorimotor system, a distributed design to simulate neural activities and structures in the brain.

has to meet the following requirement.

Firstly, internal representation includes both external sensations from external sensors, such as cameras, microphone, etc. as well as those from internal ones such as motor sensors. Take humans as an example, when pulling a door, we know exactly how much force each muscle exerts, although often unconsciously. Similarly, when an agent enters a room, it should not only remember what it sees through cameras, but also the physical information such as its neck position at that time.

Secondly, the internal representation in our model is highly distributed since no single neuron in a human brain corresponds with any particular symbol or object. Unlike the traditional models which typically uses symbolic representation, our model uses high dimensional numerical vectors for internal representation.

With the realization of above distinctions, there comes the question as how such internal representation can be applied to high-level extension, especially conceptual development. We suggest that this can be achieved by multimodel integration at the representation level since each of the sensory inputs is so closely tied to a unique subjective impressions. Foe example, when we use a word to refer to an object we see or touch, we categorize the object together with others to which we might apply the same term. More specifically, objects in our model are internally represented as high-dimensional numerical vectors. When it is referred by a term, the corresponding auditory, visual and somatosensory sensation will all be combined together. Since the reference term implies an underlying conceptual mechanism, the concept is then a hypothetical construct which consists of all the representation and individual possesses about a category of objects or events.

Although our current architecture allows such valuable extension, the uncertainty that how integration takes place in human brain still prevents us from the actual application. Moreover, a baby's early intellectual development is largely non-verbal, rather it is more concentrated on learning to coordinate purposeful movements through sensory information. Therefore, in this paper, our architecture only considers two sensory modalities, vision and somatosensory (neck position). This sensory-motor representation resembles one of the most fundamental phenomena of the natural world, although in a much simpler way.

3.2 *Sensory-motor Subsystem*

Built upon a regression engine called incremental hierarchical discriminant regression (IHDR) (Hwang and Weng, 2000), the sensori-motor subsystem magnified in Fig. 1 is the central unit of the architecture. Because of space limit, we will not go into details of IHDR. Statistically, it is a hierarchical tree structure of (nested) partitions on the input space, with the boundary of each partition determined by Bayesian estimation. Thus, when an input vector comes in, its projection onto the most-discriminating feature subspace will be conducted in a coarse-to-fine fashion until the closest matched prototype be retrieved.

It should be noticed that there are two IHDR trees in the sensori-motor subsystem. The upper one is called the reality tree or R-tree, and the bottom one the priming tree or P-tree. The P-tree is only different from R-tree by having a prototype updating queue (PUQ), which will be discussed in the following section.

3.3 Priming Mechanism

Representation is not only to remember, but in some cases to anticipate based on experiences. This ability to retrieve the past and predict the future is also called priming, one of the most elementary forms of learning.

In order to implement priming in our architecture, it is necessary to keep a doer and a predictor at each state in the IHDR tree. The doer $l(s(t))$, which is called 'last context,' includes the last sensation and last action, while the predictor $p(s(t))$, also 'primed context,' includes the primed sensation and primed action. At each time frame, the agent uses the last context of the next state to update the current primed context in order to predict the future. The rule is shown in the following equation based on Q-learning[9]:

$$p^{(n)}(s(t)) = \tfrac{n-1-a(n)}{n} p^{(n-1)}(s(t)) + \tfrac{1+a(n)}{n} \gamma l^{(n-1)}(s(t+1)), \quad (1)$$

where, $p^{(n)}(s(t))$ is the primed context at time instance t, n represents the number of times $p^{(n)}(s(t))$ has been updated, $l^{(n-1)}(s(t+1))$ is the last context of the next state, and γ is a time-discount rate. $a(n)$ is an amnesic parameter used to give more weight on the newer data points. Generally, the last context of the next state will be used to sharpen the primed representation of the current state and help the agent to reliably predict the future next time when it encounters the same visual experience.

In order to have a farther prediction capability, especially when an agent is required to prime several steps before the real seeing, we need a strategy called context updating. This is done by the prototype updating queue(PUQ), which maintains a list of pointers to the primed contexts that have been recently retrieved by IHDR. At every time instance, a pointer to a newly retrieved primed context enters the PUQ while the oldest one moves out. When the pointers are kept in PUQ, the primed contexts they point to will be updated with a recursive model based on the same learning rule as above:

$$p^{(n)}(s(t)) = \tfrac{n-1-a(n)}{n} p^{(n-1)}(s(t)) + \tfrac{1+a(n)}{n} \gamma p^{(n-1)}(s(t+1)). \quad (2)$$

3.4 Value System

Value system signals the occurrence of salient sensory inputs, modulates the mapping from sensory inputs to action outputs, and evaluates candidate actions. A developmental robot should be able to perform this task since otherwise it is just a passive receptor without knowing what it needs. In our model, the value system is based on two parts, namely, internal preference to novelty as well as external reinforcements. Since early-stage learning such as "drawbridge" experiment does not include much reward and punishment effects, we will only concentrate on internal preference to novelty in this paper, readers are referred to (Huang and Weng, 2002) for reinforcement learning and its reward system.

In 1998, Brown & Xiang found novelty neurons in ventral temporal lobe, which respond strongly to the first presentation of a novel stimulus and only weakly to its repeated presentation some minutes later (Brown and Xiang, 1998). It is suggested that priming, which causes a sharpening of stimulus representation for more efficient processing in the cortex, may lead to such neural activities in those novelty neurons. Inspired by this, the novelty in our model is measured based on priming, as the disagreement between what is predicted by the predictor and what is really seen by the doer. Algorithmatically we define novelty as the normalized distance between the selected primed sensation $p^{(n)}(s(t)) = (p_1^{(n)}, p_2^{(n)} ... p_m^{(n)})$ and the last (actual) sensation $l(s(t+1))$ at the next time:

$$n(t) = \sqrt{\frac{1}{m} \sum_{j=1}^{m} \frac{(p_j^{(n)}(s(t)) - l_j(s(t+1)))^2}{\sigma_j^2(t)}}, \quad (3)$$

where m is the dimension of sensory input and σ_j is the time-discounted average of the squared difference $(p_j^{(n)}(s(t)) - l_j(s(t+1)))^2$.

There are two problems associated with the novelty. First, the novelty so defined is sporadic in time since consecutive frames may have very different novelties. Second, the novelty is typically delayed. A baby might not realize a novel stimuli until it is clear and obvious enough. Therefore, novelty values need to be smoothed and look ahead in time.

To solve these problems, we keep a value $Q(s(t), a(t))$ for every possible action a for state s at time t. Since the values are highly distributed in the robot's brain, so is the value system. The action $a'(t+1)$ that maximizes the value $Q(s(t), a(t))$ for the current state s will be chosen as the action for the next state s' with a novelty of $n(t+1)$ to be received. The Q-learning updating expression of this process is as follows:

$$Q(s(t), a(t)) = (1-\alpha) Q(s(t), a(t)) + \alpha (r(t+1) + \gamma max_{a'} Q(s'(t+1), a'(t+1)), \quad (4)$$

where α is the updating rate and γ a time-discounter. The above algorithm shows that Q-values are updated according to the immediate novelty $n(t+1)$ and the value of the next state, which allows delayed value to be back-propagated in time during learning.

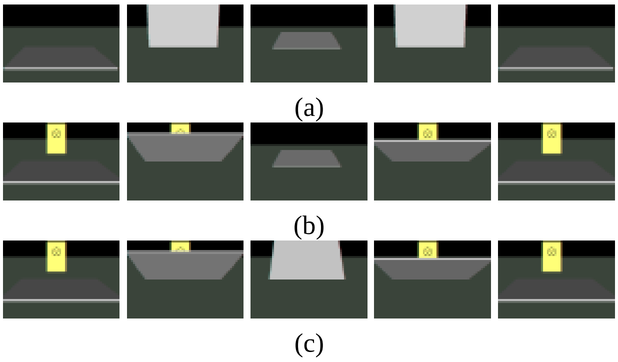

Figure 2: The animated settings used in (a)Habituation Event (b)Impossible Event (c)Possible Event.

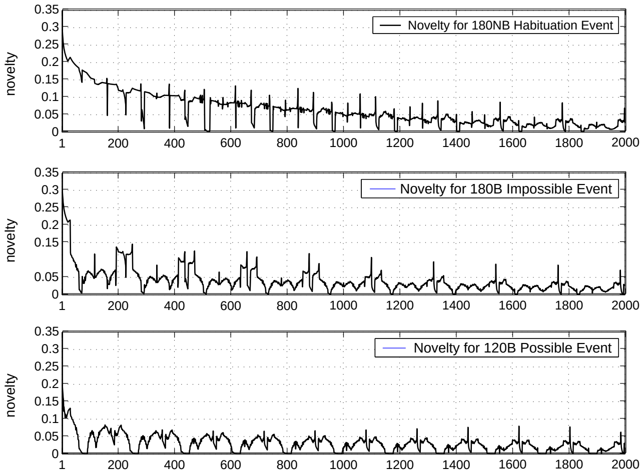

Figure 3: Novelty measured in the simulation experiment.

Figure 4: SAIL in the drawbridge experiment.

4. Simulation Experiment

SAIL (Self-orgnizing Autonomous Incremental Learner) is a human-size mobile robot, house-made at Michigan State University. With a total number of 13 degrees of freedom and a variety of sensors and effectors, SAIL has been serving as a test-bed for psychological and developmental experiments for more than five years. The "drawbridge" experiment was originally conducted in a dimly lit room, where a special setting was set up, including a gray table with black background, a silver screen and a yellow box with a clown face painted on it. Three events were then presented during the experiment, namely, habituation event, impossible event and possible event.

In the habituation event, the yellow box is absent and the screen has a full rotation of 4 sec away from the robot, 1 sec pause in the middle and 4 sec back toward the robot. In the impossible event, the box is presented and the screen moves all the way to $180°$ as if the box were not there, whereas in the possible event, the screen rotates up to the position of the box at $120°$ and then comes back.

For precise control of the experiment, we created a 3-D animated experimental setting (Fig. 2) by using OpenGL. The three events were shown sequentially one after another as consistent with the original "drawbridge" experiment.

Fig. 3 showed the novelty that was detected by SAIL during each of the three events. As is shown, after being habituated, the robot found more novelty in the impossible event than in the possible one, especially during the first 1200 image frames. This primary result reinforced our hypothesis that the value system in our architecture can be used to develop the robot's perceptual capabilities and show behaviors that are consistent with those of human babies in the original experiments (Baillargeon et al., 1985) (Baillargeon, 1987).

5. Online Experiment

The simulation experiment is conducted on SAIL without any pre-experimental experiences. In order to be more consistent with the psychological experiment and to conduct a statistical analysis of the results, we create 12 subjects for the "drawbridge" experiment by developing SAIL under different environments. Totally, we obtain 12 "babies", with an average living time of 15 minutes per day for four days, and the average growing speed of brain size at 9.0MB per day. The settings that these robot "babies" experienced include toys, objects, people, in the room and along the corridor. Fig. 5 shows the actual setting of the online experiment.

All subjects are then tested in two sections separately: Section 1 – the impossible event is presented first and Section 2 – the possible event first. Totally, we obtain a full set of 12×2 samples for each experiment. During the experiment, the robot has seven head positions with three actions for each state: stay, turn left and turn right. At each time frame, the novelty detected by the value system leads the robot to act. If the novelty of staying is higher than that of the other positions, the robot just stays, otherwise, it turns the head. Once the robot turns away, we record the time and then prepare for the display of the next event.

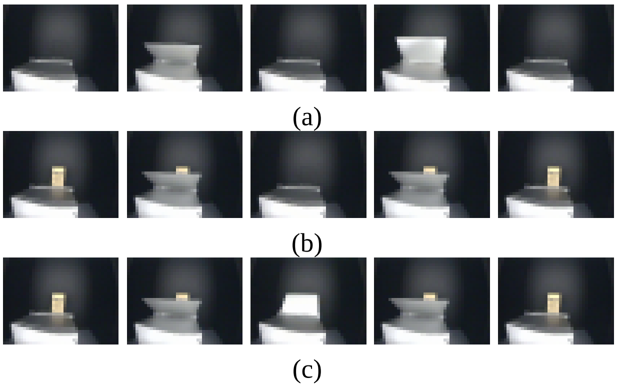

Figure 5: The actual settings of “drawbridge” experiment in (a) Habituation Event (b) Impossible Event (c) Possible Event.

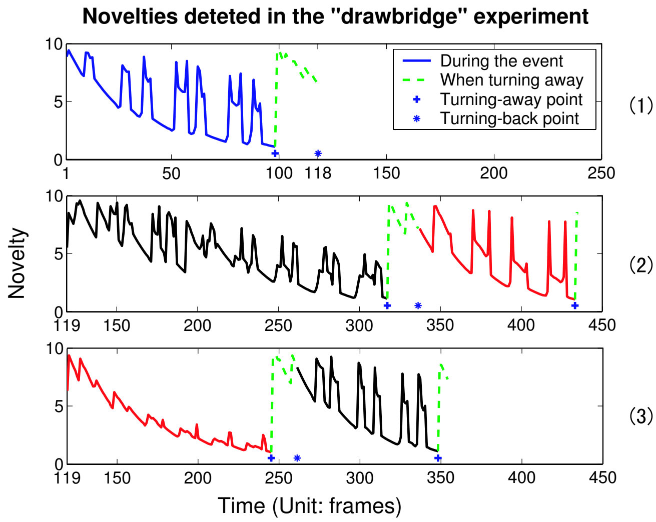

Figure 6: Novelties detected in 1) Habituation event. 2) Impossible and Possible events in Section 1. 3) Possible and Impossible events in Section 2.

Fig. 6 shows the novelties detected in baby 1. In the habituation event, the baby stays still without turning until the 99th time frame (the solid line in Fig. 6-(1)). If six trials of screen rotation has not been finished by the time the robot turns its head, the screen will keep rotating until it is done (the dash line in Fig. 6-(1)). The $+$ signs in the figure mark the points when the robot turns its head away and the $*$ signs represent the points that the robot is turned back to the setting. During all other time, the robot stays still.

After the habituation event, we set up the yellow box and start the two sections separately. In Section 1, we play the impossible event first and then possible event; while in Section 2, the possible event is displayed first. Every time when the robot turns away from the setting after one event, time is recorded and it is turned back when the next event begins.

In Section 1, as shown in Fig. 6-(2), the robot spends 199 (from the 119th to 317th) time frames on the impossible event, and only 97 (from the 337th to 433th) time frames on the possible one. In Section 2, as shown in in Fig. 6-(3), the robot spends 127 (from the 119th to 245th) time frames on the possible event and 98 (from the 262th to 349th) on the impossible one. In this case, baby 1 spend more time looking at the impossible event than the possible one in Section 1 but almost equal time at the events in Section 2.

The above experiment is then carried on all other brains that have been developed on SAIL. First two figures in Fig. 7 shows the experimental results on all 12 subjects during the two sections respectively. It must be noticed that the time measurement in this figure is the real time instead of the frame number. Generally, the robot can process two to three frames per second, considering the variance on the real-time system. Each robots shows consistently longer looking time at the impossible event than the possible event no matter which event is displayed first.

The last figure in Fig. 7 demonstrated that the difference between the impossible and the possible event during Section 1 is obviously bigger than that during Section 2. It indicates that in Section 1, when impossible event is first presented, the robot “babies” look reliably longer than in the possible event. Whereas in Section 2 when possible event parented first, the robot “babies” tend to look equally long at the two events.

Compared with the original experimental results in (Baillargeon et al., 1985), our experimental results are very consistent. Specifically, a significant main effect of event is detected (F(1,47)= 58.21, P<0.0001), which is also found in (Baillargeon et al., 1985) (F(1,83)=13.66, P=0.0004). Although in our experiment, the order is not significant in terms of the looking time of the two events (F(1,47)=0.01, P=0.9300), significance of order in terms of time difference of the two events does emerge (F(1, 23)=4.35, P=0.0488). This significance reflects different looking patterns for the two orders in our experiment. Specifically, during Section 1, when impossible event is first displayed, the robot looks reliably longer at this event (M=47s, SD=6.7823) than at the possible one (M=32.5s, SD=3.4510), whereas when possible event (M=35.25s, SD=6.0772) is presented first and the impossible event the second (M=44.6667s, SD=4.6188), the robots tend to look less differently at both events (fig .8).

5.1 *Conclusion*

By conducting the “drawbridge” experiment on our SAIL developmental robot, we find out that innate mechanisms plus the developmental learning paradigm, especially novelty detection based on experience learning, should have played an important role during the early stage of infant perceptual and cognitive development. Particularly, our computational model of novelty detection enables the robot to show similar behavior as the infants do in the original experiment without prior physical knowledge of object permanence, which indicates that during the first 9

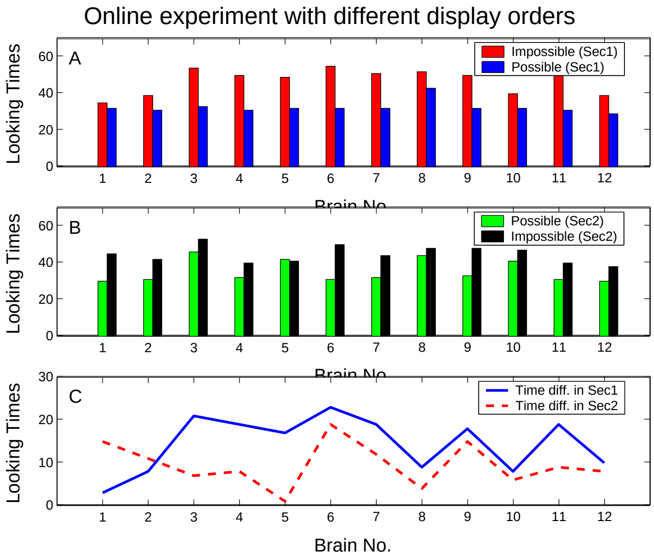

Figure 7: Looking times of twelve babies in both sections and the time difference in each of the two sections.

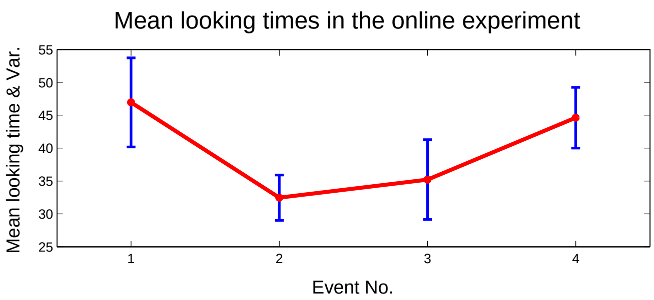

Figure 8: The average looking time and its variance in 1) Impossible event in section 1; 2) Possible event in section 1; 3) Possible event in section 2; 4) Impossible event in section2.

month (as suggested by Jean Piaget), the babies might still use the perceptual mechanism instead of sophisticated conceptual knowledge to generate correspondent behaviors.

There are several reasons to support our hypothesis: First is the influence of the habituation effect. In the original experiments, only fast habituators (those who reached the habituation criterion in six or seven trials) looked significantly longer at the impossible event, while slower ones tended to look about equally at the two events. We doubt whether infants were using possibility and impossibility of the events since the habituation phase should not have such an obvious effect. Second is the influence of the order. In the original experiments, only when impossible event was presented first would the infants look longer at it. Otherwise, the infants spent equal time looking at the two events. Although Baillargeon claimed that “such order effects are not uncommon in infancy research and are of little theoretical interest (Baillargeon et al., 1985),” again, we doubt whether this is really the case. Moreover, we find influence of order in the looking time difference of two events, showing that the order does change the looking patterns. Finally, Baillargeon and her colleges reported in 1995 that they did not find the same results with $6\frac{1}{2}$-month-olds, while Cashon and Cohen successfully replicated the previous findings with those infants who responded based on familiarity.

Because of the above inconsistency, we doubt whether the conceptual point of view holds any water in the “drawbridge” experiment case. Our experiment results show that the phenomenon detected in the original “drawbridge” experiment may be mainly caused by visual preference for novelty, which should be considered as the origin of the conceptual knowledge of object permanence rather than the outcome. More control experiments should be conducted and compared with the original results in order to add more credibility to this work.

Acknowledgments

This work is supported in part by the National Science Foundation under grant No. IIS 9815191, DARPA ETO under contract No. DAAN02-98-C-4025, and DARPA ITO under grant No. DABT63-99-1-0014.

References

Baillargeon, R. (1987). Object permanence in 3.5- and 4.5-month-old infants. *Developmental Psychology*, 23:655–664.

Baillargeon, R., S.Spelke, E., and S.Wasserman (1985). Object permanence in five-month-old infants. *Cognition*, 20:191–208.

Baker, C. I., Keysers, C., Jellema, J., Wicker, B., and Perrett, D. (2001). Neuronal representation of disappearing and hidden objects in temporal cortex of macaque. *Exp. Brain Res.*, 140:375–381.

Bogartz, R. S., Shinskey, J. L., and Schilling, T. H. (2000). Object permanence in five-and-a-half-month-old infants? *Infancy*, 1(4):403–428.

Brown, M. and Xiang, J. (1998). Recognition memory: neuronal substrates of the judgement of prior occurrence. *Progress in Neurobiology*, 55(2):149–89.

Cashon, C. H. and Cohen, L. B. (2000). Eight-month-old infants’ perception of possible and impossible events. *Infancy*, 1(4):429–446.

Huang, X. and Weng, J. (2002). Novelty and reinforcement learning in the value system of developmental robots. In *Proc. Second International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems (EPIROB 2002)*. Edinburgh, Scottland.

Hwang, W. and Weng, J. (2000). Hierarchical discriminant regression. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22:1277–1293.

Mareschal, D., Plunkett, K., and Harris, P. (1995). Developing object permanence: A connectionist model. In *Proc. 17th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ.

Melchner, L., Pallas, S. L., and Sur, M. (2000). Visual behaviour mediated by retinal projections directed to the auditory pathway. *Nature*, 404:871–875.

Piaget, J. (1954). *The construction of reality in the child*. Basic Books, New York, NY.

Rivera, S. M., Wakeley, A., and Langer, J. (1999). The drawbridge phenomenon: Representational resoning or perceptual preference? *Developmental Psychology*, 35:427–435.

Roder, B. J., Bushnell, E. W., and Sasseville, A. M. (2000). Infants' preference for familiarity and novelty during the course of visual processing. *Infancy*, 1(4):491–507.

Saffran, J. R., Aslin, R. N., and Newport, E. L. (1996). Statistical learning by 8-month old infants. *Science*, 274:1926–1928.

Schlesinger, M. (2002). A lesson from robotics: Modeling infants as autonomous agents. In *Second International Workshop on Epigenetic Robotics*, pages 133–140. Sweden.