Introduction: The Second International Workshop on Epigenetic Robotics

Christopher G. Prince

1. Background

In concluding their summary of papers from the First Epigenetic Robotics workshop, Zlatev and Balkenius (2001) considered a number of questions. First, were the issues of the embodiment, situatedness, and development that are most adequate for epigenetic robotics (and cognitive science). Second, was the question of true animal-like or human-like intelligence in artificial (including robotic) systems (i.e., "Strong AI") versus simulations dependent on external interpretation. Third was whether or not consciousness (with manifestations including emotion, intentionality, awareness, intersubjectivity) is essential for complex cognitive development. These issues are certainly still open questions, less than a year after the first workshop. Our present 15 oral papers, and 12 posters, along with contributions from three invited speakers: Luc Steels, Colwyn Trevarthen, and John Weng contain further inquiry into these and other questions. Like other areas in which working systems are a desirable outcome (e.g., mobile robotics), an important concern in epigenetic robotics is that of component integration. In the present summary of papers from the Second Epigenetic Robotics workshop, I classify the papers as involving: *theoretical and conceptual issues*, *complex skills*, *developmental architectures*, and *sensory-motor integration and imitation*. From my perspective, research in epigenetic robotics involves addressing *theoretical and conceptual issues* by constructing potentially *complex skills* based on *developmental architectures*, and involving *sensory-motor integration and imitation* (and other functions). The papers and posters for the workshop are summarized below under these headings.

2. Theoretical and conceptual issues

Razvan V. Florian discusses the possibility of *artificial science*, where embodied artificial cognitive agents might develop their own science. Florian considers science as the abilities, in a particular environment, to understand the structure of the environment, and to generate both predictions and causal explanations. The following cognitive mechanisms are proposed as important to science: causality detection and mental simulation, coherence detection (comparing a simulation against reality), subitizing (fast discrimination between small quantities), abstraction, projective reasoning (e.g., analogies, metaphors), and symbolic association. The author suggests that "human scientific concepts…crucially depend on…sensorimotor capabilities." For example, causality and coherence detection both will likely depend strongly on the specifics of the particular environment, as provided through sensori-motor interaction.

Carlos Gershenson presents a theoretical categorization of types of behaviors in an effort to show how behavior-based robotics can close the gaps between these systems and knowledge-based systems. Types of behaviors span vegetative, reflexive, reactive, motivated, reasoned, and conscious. Gershenson suggests the following epigenetic stages to reach knowledge from behavior: 1) concept abstraction, 2) grounding of concepts through action, 3) sharing of concepts through social interactions (language), 4) manipulation of concepts (logic), and 5) evolution of concepts (culture). The author indicates that the literature contains theories and implemented models for most of these separate stages.

Jessica Lindblom and Tom Ziemke discuss *social situatedness* in the context of contributions from developmental psychologist Lev Vygotsky. Socially situated agents acquire information about, and may interact with, the social and physical worlds. Additionally, this encompasses "the idea that the development of individual intelligence requires a social (and cultural) embedding." Vygotsky describes transformation processes, from elementary to higher psychological functions as occurring via signification (use of mental mediating 'tools' between stimulus and response) and cultural development or *internalization*. Internalization is the idea that every function in a child's psychological development appears twice—first at the social level, between people, and second, internalized, psychologically inside the child. The authors raise the question, in regards to socially situated robots, as to whether internalization can occur in these robots. Another Vygotzkyian concept addressed is the zone of proximal development, in which a child learns more than they would independently, through the support of adults or more capable peers.

Matthew Schlesinger considers the important issue of innate knowledge in infants. In particular, while it seems clear from developmental psychology that that young infants (e.g., 6 months of age) are at least perceptually sensitive to anomalous visual test scenes (e.g., a toy car appearing to pass through a solid object), the kind of mental representation underlying this sensitivity is less than clear. Two possible explanations involve a representational account with explicit knowledge (e.g, Baillargeon, 1999) on the one hand, and a perceptual-processing account with implicit knowledge (e.g., Smith, 1999) on the other. Schlesinger presents a feed forward neural network model for oculomotor control that is designed to mimic one of the anomalous visual test studies with infants. Simulated input is provided to the model in a training phase, followed by possible and impossible (anomalous) test cases. The model is found consistent with a perceptual-processing account, and infants possessing implicit innate knowledge.

In a poster, **Uri Hershberg** argues that a focus of study for cognitive modeling should be the statistical distribution of environments and, how this affects the behavior and development of the artificial system. Based on this and the concept that optimal environmental examples are important to cognitive systems, Hershberg proposes that immune systems are cognitive systems, because they adapt, creating the immune system's receptor repertoire, based on examples presented by the environment. In another poster, **S. Itakura, A. Izumi, M. Myowa, M. Tomonoga, M. Tanaka, and T. Matsuzawa** present a study of infant chimpanzees which uses techniques and a hypothesis from the study of human infants by Rochat and colleagues. The hypothesis is that an early sense of self arises in young human infants via a propensity to engage in self-perception and systematic exploration. The authors found that infant chimpanzees showed more rooting responses following external (than self) stimulation, which seems aligned with Rochat's hypothesis. In two final posters in this section, **Claudia Uller** considers the psychology of young infants. In the first, Uller considers the representation of *number* by young infants. Two current proposals for how infants represent number are an "object file" account and analog representations. The former is characterized by a small set size (e.g., groups of 1, 2 or 3 items), whereas set size is not a concern with analog representation. Uller rallies data that support the object file account. In Uller's second poster, she presents research on young infants' capacity to perceive the desires of others. These children are familiarized with adults saying "wow" or "yuk" to particular foods (broccoli, cracker), and then tested, in a habituation paradigm, to determine if they show surprise when the experimenter grabs the food to which she previously expressed disgust ("yuk").

3. Complex skills

Naoto Iwahashi discusses an implemented robotic system which models language acquisition through auditory and visual modalities, and reinforcement. Words are learned for objects and for motions. The system focuses on particular objects by both tracking the pointing gesture direction of the human experimenter, and through object motion—"When attention is given to objects and the person speaks, the observations of those objects are associated with that speech." The language structures that were learned included phoneme sequences (comprising lexical items), grammar (probabilities of word type ordering in 1, 2, or 3 word sentences, with three word types), and mutual beliefs (comprehension of sentences by the robot). Visual object images are represented using multivariate normal probability density functions. Learning techniques included hidden Markov models (HMM's), likelihood maximization, and Bayesian learning.

Giorgio Metta and Paul Fitzpatrick consider active robotic strategies for gaining experience with "objects" in the world, i.e., "parts of the environment [that] are physically coherent ensembles … [that] move together, and which are more or less independent." They work from the hypothesis "that action is required [for] object recognition in cases where an agent has to develop categorization autonomously." The authors use correlations of changes in optical flow, and changes in the robot's arm position to localize the position of the robot's arm. Specifically, with a reversal in direction of the robot arm, the optic flow at that instant will change sign. With the robot's head kept steady to simplify vision processing, and knowledge of what part of the visual scene is comprised of the robot's arm (e.g., using optic flow as above), image differencing is used to specify the parts of a scene that move together when "poked" by the robot's manipulator, and hence which parts of the scene form a coherent "object." Their experimental platform is the Cog robot (Brooks et al., 1999).

Paulina Varshavskaya presents a paper extending Kismet's (Breazeal & Scasellatti, 2000) behavioral repertoire to include proto- (pre-grammatical) language development. This protolanguage module includes vocal behaviors, regulatory drives (including speech and exploration), and learning algorithms. The goal of this project was to produce vocal output modeling a child of age 10-12 months, including emotive grunts, canonical babbling, and formulaic speech. Vocal behaviors ("Grunt", "Babble", and Concept/word behaviors) compete to establish the mode of vocal expression, and the phonemic string that is to be produced. Behaviors of type *Concept* represent specific words that have either been learned or initially programmed into the system. The *ConceptMap* behavior creates a new instance of a Concept behavior when speech input arrives that does not match the speech associated with existing words (Concepts).

In a poster, **Paul Dickerson** considers conversational analysis (CA) and its relevance for the design and understanding of interacting robots. CA entails detailed analysis of interactions in communicative sequences, and emphasizes language pragmatics. Dickerson considers the implications of CA for humanoid robotics, and analyzes some interactions with Kismet in these terms.

4. Developmental architectures

In an invited paper, **Juyang (John) Weng and Yilu Zhang** present a clearly argued statement justifying why robotics needs developmental algorithms—because the full set of tasks facing a robot are unknown and cannot be predicted by the programmer. Indeed, when robots venture into a new environment, by its very nature, the new environment involves unknown tasks. A developmental algorithm should be task nonspecific. The authors discuss the evolution of developmental robotics at their lab at MSU from *Creseptron* (automatic generation of networks for recognizing images), *SHOSLIF* (which uses linear discriminant analysis and principle component analysis for video-based navigation), to *DAV* (a humanoid robot) and *SAIL*, their latest projects. SAIL is based on IHDR, a general and efficient algorithm for incrementally generating a mapping from high-dimensional input vectors (sensory input and state) to high-dimensional output vectors (new state and motor output). Training bouts can be interleaved with performance and use inputs consisting of current system state, current sensor data, and motor output as shaped by a human trainer, and derives sensory and state features related to the new state and motor output, disregarding unrelated features. In this way, a human trainer modifies the mapping function and hence the behavior of the robot. Innate behaviors (e.g., motion detection) are either explicitly programmed or they are generated using the IHDR algorithm offline before online use of the robot. The authors extend the traditional notion of agent by suggesting that agents should have both internal and external sensors and effectors. An example of an internal effector is attentional selection.

Xiao Huang and John Weng report on developments of the SAIL robot and the Incremental Hierarchical Discriminant Regression (IHDR) method (Hwang & Weng, 1999). IHDR provides a basis within which to design task non-specific developmental programs that run on robots and enable development of robot-skills through real-time, online environmental interactions. In these experiments, the authors have added a value system to signal the occurrence of salient sensory inputs. The goal in this research was to integrate novelty detection (salience) and reinforcement learning to model a habituation effect. Novelty was considered to be low in the case when the robot had a high degree of success in predicting the outcomes of its actions. IHDR was used to predict the sensory input and value after taking an action. Experiments were conducted both via simulations and with the SAIL robot, and included environmental exploration and altering novelty levels by showing the robot a new object.

Andrea Kulakov and Georgi Stojanov discuss two developmental robot architectures: Petitagé and Vygovorotsky, based respectively on developmental theories by Piaget and Vygotsky. Petitagé innately has a value system, schemas (action sequences), and tries to execute portions of its schemas, hence enabling the schemas in the context of certain percepts. Higher-level structures are created by detecting regularities (cycles) in the stream of enabled schemas, and yet higher-order structures can be created by detecting regularities in the regularities. Vygovorotsky is based on the use of analogy for problem solving. Knowledge is represented by a conceptual network, and percepts activate portions of this network. Three "innate" drives: hunger, affect, and curiosity are used. The latter drive, for example, causes new parts of the conceptual network to receive high activation levels.

Ulrich Nehmzow suggests the use of *learning controllers* because "fixed behavioural strategies…will usually be brittle in practice, due to the noisy and partly unpredictable nature of the real world." This research uses a collaborative learning process with multiple mobile robots and a Physically Embedded Genetic Algorithm (PEGA). PEGA is an augmented genetic algorithm with genetic strings stored in each robot, and communicated to other robots, along with their fitness measures. In this way, each robot can contribute to the exploration of the task sensory-motor space. Experiments were conducted to acquire individual competences of phototaxis (light seeking), obstacle avoidance, and robot seeking, and also to acquire multiple competences.

Akitoshi Ogawa and Takashi Omori report on a functional parts combination model. The goal of this project was to develop a method to acquire a problem solving strategy, depending on a task. While some research involves methods designed to solve particular tasks (e.g., a specific application of reinforcement learning), this research intends to design a "method that acquires a problem solving strategy itself depending on a task." Indeed, this is one general feature of development itself, to provide a period of environment-organism interaction in which the organism is, as a result, better adapted to the environment. Their model involves a number of fixed functional parts (e.g., state recognition, reinforcement learning, action selection) which are combined in specific ways through use of a genetic algorithm and interaction with a specific environment. Two grid world simulation experiments are conducted involving the navigation problems of locating a goal and avoiding obstacles.

In a poster, **Andrzej Buller** presents a psychodynamic robot model that emphasizes psychic tension as a source for a robot's motivation. This tension is taken to be a variable characterizing working memory, and is computed based on the difference in quantities of satisfaction and dissatisfaction content items in working memory. The system is hardwired with a goal of maintaining a low level of tension.

5. Sensory-motor integration and imitation

Pierre Andry, Philippe Gaussier, and Jacqueline Nadel emphasize an online learning process of random exploration of motor dynamics and physics to solve the problem of having a robot generate appropriate movement of an arm. Their perception-action loop control architecture consists of perception paths from arm proprioception and vision merged in a sensory-motor map, consisting of clusters of neural units learning visuo-motor associations. Visual information triggers Neural Field maps, which are used to compute motor commands. The authors indicate their method works independent of the number of degrees of freedom of the robot arm, and they experimented with a supervised form of pointing, and imitative behaviors.

Max Lungarella and Luc Berthouze present a study focused on the hypothesis "that physical limitations inherent to body development could be beneficial to the emergence of stable sensorimotor configurations and allow for more tolerance to environmental interaction." The authors experimented with a small-sized humanoid robot, and the control of swinging behavior in this robot. Three degrees of freedom (DOF) of the robot's legs were used, and proprioception for joint position was simulated by a camera. Experiments were conducted with 1- and 2-DOF exploratory control, and bootstrapped 2-DOF exploratory control. The bootstrapped 2-DOF experiments corresponded to developmental release of a second DOF after the system had stabilized with 1-DOF control. In this case, the authors found "that independently of the choice of control parameters, the system converges into a unique smooth, in-phase swinging behavior with maximal amplitude."

Yuval Marom, George Maistros and Gillian Hayes discuss robot learning by imitation and temporal attention using the idea of mirror neurons. Their hypothesis is that a socially situated agent with a mirror system and a set of innate skills can develop new motor skills exhibited by a demonstrator. The software architecture of their system involved a perceptual classification system that clusters the sensory-motor experience of the robot, and an ("innate") inverse model that can convert perceptual states into motor commands. A self-organizing feature map augmented to enable habituation and dishabituation is used for the perceptual classification system. Simulation experiments were conducted with a "drinking" behavior, as was a robot experiment involving learning of wall-following by tracking a human demonstrator (learning occurred offline).

Toru Nakata reports on simulations of life-like body motion using an algorithm based on somatic theories. Somatic theories are conventionally used to evaluate human psychological and developmental states through analysis of body movements. Nakata's algorithm was based on anatomical constraints, distribution of tension signals, and application of rhythm stereotypes (wave forms; e.g., undulate, burst). Experiments were conducted in which adult human subjects were shown displays of the simulated motion, and asked to answer questions about the type of motion displayed. Questions involved calmness vs. excitation of the motion and whether the subject considered the motion to involve pleasure or displeasure.

In a poster, **Ben Choi and Yanbing Chen** propose a humanoid motion description language including motion description layers (joint angle, path, motion primitive, and motion sequence), an egocentric reference system, progressive quanitized refinement (akin to foveated vision), and indicate a need for automatic constraint satisfaction. They also outline methods to learn motor skills, using a motion description language, based on non-deterministic finite state automata. In additional posters, **Hideki Kozima, Kokoro Nakagawa, and Hiroyuki Yano** describe a model of the emergence of a mirror system for imitative learning. A mirror system is an intermodal mapping between someone else's action (as seen) and one's own action (as executed). The authors claim that this mapping is not a geometric transformation, but a learned functional correspondence. They outline how this mapping might be used in imitation. **Lawrence Warnett and Brendan McGonigle** describe the learning of navigation behavior based on only inbuilt compass information, short-range sensing, and primitive state. The authors' algorithm includes identifying the place in the environment where the robot is located, generating potential headings from this place, iteratively following the headings and returning to the original place, and selecting the best of the headings. **Krister Wolff and Peter Nordin** describe three humanoid robot platforms for use in robotics research. Because of the difficulties of inverse kinematics with bio-inspired robots, and because of environmental complexities, the authors suggest adaptation mechanisms for these complex robots. They also propose an anthropomorphic principle that aims to explore humanoid robots built with close correspondence to human morphology and motion. **Yuichiro Yoshikawa, Hiroyoshi Kawanishi, Minoru Asada, and Koh Hosoda** discuss the problem of how different sensory modalities of a robot may be associated to construct a body scheme for a robot. A body scheme is needed when performing tasks where body parts and external space must be related. The authors propose a cognitive development body scheme model consisting of a cross-modal map among tactile, vision, and proprioceptive spaces, which is acquired by experiences of self-touching. **Tomoyuki Yamamoto and Yasuo Kuniyoshi** propose a concept of global dynamics in which a detailed map of body dynamics is made, separating stable from unstable regions of phase space. The authors state that it is primarily in the unstable regions of the phase space that control inputs are necessary, and that constructing and using such a map can reduce the cost of adaptation.

6. Conclusions

I would like to emphasize two themes evident in a variety of the contributions. First, *What behaviors, skills, or mechanisms should be provided innately in an epigenetic robot?* **Schlesinger** addresses this question, with the current debate in developmental psychology as to innate knowledge in infants. Other authors install particular innate features in their robots (e.g., **Andry, Gaussier, and Nadel** and a perception-action loop architecture; **Marom, Maistros, and Hayes** and an inverse model). Second, *How can the various components of an epigenetic robot be combined?* **Gershenson** touches on this, and work on developmental architectures considers the issue more explicitly (e.g., contributions by **Weng** and colleagues). This question has another form—how can we, as researchers, best combine our efforts to lead the way to these working epigenetic robots? Certainly, we need contributions ranging from (but not in anyway limited to) psychologists (such as **Uller** and **Itakura, Izumi, Myowa, Tomonoga, Tanaka, and Matsuzawa**), to engineers, and including developmental linguists.

In closing this introduction, we thank the Communications Research Laboratory, Japan (CRL) for their generous support of this workshop. We also thank the program committee members for their efforts in reviewing submissions.

References

Baillargeon, R. (1999). Young infants' expectations about hidden objects: A reply to three challenges. *Developmental Science*, *2*, 115-132.

Breazeal, C. & Scassellati, B. (2000). Infant-like social interactions between a robot and a human caregiver. *Adaptive Behavior*, *8*, 49-74.

Brooks, R. A., Breazeal, C., Marjanovi_, M., Scassellati, B., & Williamson, M. M. (1999). The Cog project: Building a humanoid robot. In C. L. Nehaniv (Ed.), *Computation for Metaphors, Analogy and Agents* (pp. 52-87). Volume 1562 of Springer Lecture Notes in Artificial Intelligence. Springer-Verlag.

Hwang, W. & Weng, J. (1999). Hierarchical discriminant regression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *22*, 1277-1293.

Smith, L. B. (1999). Do infants possess innate knowledge structures? The con side. *Developmental Science*, *2*, 133-144.

Zlatev, J. & Balkenius, C. (2001). Introduction: Why "epigenetic robotics"? *Proceedings of the First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems* (pp. 1-4). Lund University Cognitive Studies, Volume 85.